

THERE'S A NEW SHERIFF IN TOWN

An Evidence-First Accountability System for Canada's Fentanyl Crisis

Building an auditable source map, question engine, five-deck review system and publication-gated public database capable of showing what happened, what was claimed, what is supported, and what remains unproven.

By Scott Gilbert
with research, drafting and risk-review assistance from ChatGPT

Methodology and Project Status Edition
Current through completion of ED-LIVE-01
August 29, 2026 | Vancouver, British Columbia

Publication and legal status

This paper is a methodology and project-status document. It does not accuse any identified person of a crime or publish any held person-level card, finding, relationship or allegation. The authoritative research package remains private and publication is not authorized. Any later public release naming a person in a materially adverse context must pass source verification, identity resolution, privacy review, defamation review, legal review, proportionality review and right of reply.

The phrase courtroom-ready describes a design objective. It does not mean that the present database, any current card, or the project as a whole is admissible evidence or ready to prove a legal claim. Courts receive underlying records, authenticated exhibits, witness testimony and properly qualified expert opinion. The Evidence Deck is an organizing and analytical system.

Current checkpoint: ED-LIVE-01 completed with PASS_WITH_GAPS on August 29, 2026. A private review database, five private 52-person rosters, card evidence packets, publication ledgers, a sanitized public read model and a minimal Replit handoff were produced and validated. No Replit change or republication occurred. Owner approval remains required.

Core evidentiary rule

Failure is evidence of failure. It is not automatically evidence of fraud.

The visible condition of Vancouver's Downtown Eastside justifies a demanding examination of public policy, spending, research, enforcement and communications. It does not, by itself, identify who caused a result, what any person knew, whether a relationship influenced a decision, or whether an offence occurred.

Authorship and use

This paper was developed from Scott Gilbert's direction with drafting and risk-review assistance from ChatGPT. ChatGPT is not a source, witness, expert or legal authority. Internal project counts are supported by the verified ED-ALBPSN-01 and ED-LIVE-01 completion artifacts, but they remain project accounting rather than independent factual findings about the fentanyl crisis.

This is not legal advice, an expert report, a criminal information, a pleading or a final finding. It may be shared for public-interest discussion with attribution. Third-party records remain subject to their own legal terms, privacy restrictions and publication rules.

Contents

Publication and legal status	1
Core evidentiary rule	1
Authorship and use	1
1 Abstract	4
2 Executive summary	4
2.1 The public-interest proposition	4
2.2 The distinction that protects the project	4
2.3 What exists now	5
2.4 What does not yet exist	5
2.5 The defensible accountability message	5
3 Why the Downtown Eastside demands an evidentiary accounting	6
3.1 What present conditions can establish	6
3.2 What present conditions cannot establish alone	6
4 What the Evidence Deck is	6
4.1 A source registry	6
4.2 A claim ledger	6
4.3 A relationship graph	7
4.4 A contradiction engine	7
4.5 A missing-record generator	7
4.6 A courtroom-readiness layer	7
4.7 A card deck as a navigation layer	7
5 Current state and scope after ED-LIVE-01	7
5.1 Authoritative lineage and verification	7
5.2 The counts that must not be confused	8
5.3 Question engine and finding synthesis	8
5.4 Expanded substantive scope	8
5.4.1 Legal and enforcement records	8
5.4.2 Business, pharmacy and safer supply	9
5.4.3 School and campus naloxone	9
5.4.4 Contaminated supply and drug checking	9
5.5 The five private 52-card rosters	9
5.6 The public read model	10
6 The complete source-record map	10
6.1 The bridge-record concept	11
7 Testing public narratives as falsifiable propositions	11
7.1 A major seizure disrupted supply	12
7.2 Contamination or counterfeit supply caused a death	12
7.3 A program reduced mortality	12
7.4 The science was independently replicated	12
7.5 Public money was used effectively	12
7.6 A school naloxone initiative was effective	12
8 Evidentiary method	13
8.1 One evidence file per proposition	13
8.2 Evidence classes	13
8.3 Independence analysis	13
8.4 Alternative explanations	13
8.5 Falsification	13
8.6 Confidence ladder	14

9 From database to courtroom	14
9.1 Courtroom-readiness test	14
10 Identifying responsibility without prejudging guilt	15
10.1 Responsibility states	15
10.2 Requirements before a serious individual finding	15
11 Accountability pathways	15
12 Governance and publication gates	16
12.1 Separation from campaign messaging	16
12.2 Independent review	16
12.3 Right of reply	16
12.4 Corrections	16
12.5 Release gates	17
13 Live system architecture	17
13.1 Two physically separate databases	17
13.2 Read-only API	17
13.3 Minimal Replit handoff	17
13.4A genuinely living document	18
14 Implementation roadmap from the current checkpoint	18
14.1 Phase 1 - Preserve and review ED-LIVE-01	18
14.2 Phase 2 - Approve neutral public metadata	18
14.3 Phase 3 - Implement the live read model	18
14.4 Phase 4 - Retrieve bridge records	18
14.5 Phase 5 - Reproducible analysis and expert review	18
14.6 Phase 6 - Findings, reply and publication	18
15 Public participation	19
16 What success would look like	19
17 Conclusion: what a new sheriff actually does	19
18 Appendix A: ED-LIVE-01 completion checkpoint	20
19 Appendix B: model finding format	20
20 Selected project records and legal authorities	21

1 Abstract

Governments, police agencies, health authorities, universities, nonprofit organizations, professional regulators, businesses, pharmaceutical interests and news organizations have made thousands of claims about fentanyl, overdose, contaminated supply, treatment, safer supply, drug checking, enforcement and the Downtown Eastside. Many claims may be accurate. Some may be incomplete, outdated, overstated or repeated beyond the strength of their evidence. Some departures from the record may result from honest error, scientific uncertainty, emergency conditions, institutional silos or reasonable policy disagreement. Others may eventually support findings of negligence, reckless disregard, concealment, self-dealing, knowing misrepresentation or other misconduct.

Rhetoric cannot reliably separate these possibilities. A structured evidentiary record can.

The Evidence Deck is being developed as an auditable source registry, claim ledger, relationship graph, contradiction engine, missing-record generator and courtroom-readiness layer. It is paired with a question engine that requires each proposition to be traced to its sources, contrary evidence, alternative explanations, unresolved records, reproducible method and publication status.

Version 1.2 records a major change in project maturity. The ED-LIVE-01 controller has converted the authoritative private integration into two physically separate products: a private editorial review database and a sanitized public read model. It also produced five private person-card rosters, each with 52 selected positions, without using organizations to pad the decks. The 260 positions resolve to 193 unique people because 60 people appear in more than one deck under separately documented roles. Every selected packet remains private, human-review-required and supported by one cited source family. No person card is public.

The authoritative project accounting now distinguishes 674 integrated SQLite/IDSUP source-record rows from 1,044 canonical source works, 1,269 source-version/crosswalk records and 15,665 namespace-crosswalk rows. These are different layers, not competing totals. The project also retains 405 research questions: 11 answered, 137 partially answered, one conflicting and 256 unresolved. The present state is therefore substantial but incomplete.

The strongest public-interest promise is not that the project already knows who is guilty. It is that institutions and individuals exercising public authority will be required to produce the receipts, and that the project will publish only what the evidence can support.

2 Executive summary

2.1 The public-interest proposition

The public has financed a vast response to opioid dependency, overdose and fentanyl. It is entitled to know:

- what was purchased and delivered;
- which organizations and businesses received public money;
- which outcomes were promised and measured;
- whether activity counts were presented as outcomes;
- what decision-makers knew when they spoke or acted;
- whether contrary evidence was available, communicated or omitted;
- whether enforcement operations changed price, availability, composition or harm;
- whether apparently independent studies reused the same cohorts, investigators or datasets;
- whether procurement, funding, lobbying, authorship, employment or corporate relationships were disclosed and properly interpreted;
- which assertions are official statements, which are independently established facts, and which remain unresolved;
- and which records would materially change the present assessment.

2.2 The distinction that protects the project

The Evidence Deck must keep the following categories separate:

What the record shows	What it does not show by itself
An institution issued a statement	The statement was accurate
Two outlets used similar language	They reported independently or coordinated
A person held a role	They personally caused every institutional act
A relationship or payment existed	It improperly influenced a decision
A charge or pleading was filed	The allegation was proven
Fentanyl was detected	It caused death or was consumed unknowingly
A program distributed kits or enrolled clients	It produced a durable health outcome
A seizure occurred	It measurably disrupted supply
A policy failed	A person acted dishonestly or criminally

These distinctions are not technicalities. They are the difference between a research system capable of correction and an accusation machine.

2.3 What exists now

As of completion of ED-LIVE-01:

- the authoritative ED-ALBPSN-01 archive passed its required outer SHA-256, ZIP integrity, 64 internal member checksums, integrated-database hash and size, SQLite integrity check and foreign-key check;
- a separate 112 MB private review database preserves the authoritative research and adds roster, packet, release-ledger and omission-log tables;
- a separate 164 KB sanitized public SQLite read model contains only dataset metadata, 21 defined corpus metrics, five methodology entries, seven status definitions, five deck definitions and one release-hold entry;
- public person cards, sources, questions, results, entities and relationships remain at zero;
- five private 52-person rosters were selected from verified candidate pools, with no padding;
- 260 private deck positions resolve to 193 unique people, and overlap is recorded rather than hidden;
- 260 private evidence packets were generated with neutral role text, source locators, inference limits, unresolved questions and publication recommendations;
- 15 read-only API routes were specified and 23 contract, integrity and security tests passed;
- the public and private packages were rebuilt independently, with 13 selected artifacts byte-identical across both runs;
- Replit Agent calls, Replit changes and publication calls were all zero.

2.4 What does not yet exist

The project is not yet a live database-backed public website. The current public application has been treated as a static React presentation with hard-coded aggregate values. The authoritative package has not been demonstrated to operate behind the live site. The ED-LIVE-01 handoff is ready for a minimal implementation pass, but the site must not be described as database-connected until the deployed API returns the exact approved dataset version and parity checks pass.

The project also does not yet have public approval for person cards, questions, findings, source rows, entity profiles or relationship graphs. A completed roster is a private editorial product, not a public accusation list.

2.5 The defensible accountability message

I will make government, police, public-health agencies, contractors, researchers and publicly funded institutions produce the receipts. We will publish what can be proven, correct what cannot, protect innocent people, and pursue lawful accountability wherever the evidence leads.

3 Why the Downtown Eastside demands an evidentiary accounting

The Downtown Eastside is not an abstract policy laboratory. It is where decisions made in legislatures, ministries, police headquarters, health authorities, universities, nonprofits, pharmacies and boardrooms become visible in human lives.

The public sees severe dependency, preventable death, homelessness, violence, exploitation, unsafe housing, exhausted responders, struggling businesses, repeated funding announcements and recurring claims of progress. A credible accountability system must ask:

What objectively improved, for whom, over what period, at what cost, and compared with what alternative?

That question is not anti-science, anti-police, anti-health care or automatically opposed to any particular intervention. It is the ordinary question that should accompany serious public spending and coercive public authority.

3.1 What present conditions can establish

Present conditions can support the narrower conclusions that:

- the crisis has not been adequately resolved;
- major harms persist;
- announced interventions have not collectively produced an acceptable outcome;
- measurement focused only on activity may be inadequate;
- claims of success require scrutiny;
- transparency, comparative evaluation and fiscal accountability are justified;
- and alternative strategies deserve serious testing.

3.2 What present conditions cannot establish alone

Present conditions do not identify the person responsible for a particular outcome. They do not establish that a specific program made conditions better or worse. They do not prove that someone knew a statement was false, that a financial relationship influenced a decision, that institutions coordinated dishonestly, or that an offence occurred.

The Downtown Eastside is a smoke alarm, not a fingerprint. It establishes the need for investigation. It does not supply every link in the chain of responsibility.

4 What the Evidence Deck is

The project began with a card-deck metaphor: identify the people, institutions, policies and claims that shaped the crisis. Its mature form is a set of connected evidence systems.

4.1 A source registry

Every study, report, news item, statement, court record, contract, invoice, dataset, laboratory result, archived webpage, policy and communication receives a stable identifier. The registry preserves title, author, date, publisher, original locator, archive locator, acquisition date, version, hash, access restrictions, evidence state and relationships to other records.

4.2 A claim ledger

Each material proposition is separated from the document containing it. Examples include:

- a seizure significantly disrupted supply;
- a substance was contaminated or represented as another drug;
- a program reduced mortality;
- a percentage of clients entered detox;
- a quantity represented a stated number of potentially lethal doses;
- a study independently replicated an earlier result.

Each claim records its exact wording, speaker, date, audience, cited basis, uncertainty language, later repetitions, corrections and current evidence assessment.

4.3 A relationship graph

The graph records employment, board membership, funding, grants, contracts, procurement, lobbying, authorship, media sourcing, regulatory responsibility, corporate ownership, pharmacy chains, shared datasets and other relevant relationships. Each edge must have a source, time period, evidence state and inference ceiling.

The graph does not convert a relationship into influence, motive, control, coordination or wrongdoing.

4.4 A contradiction engine

The system can compare:

- public statements with internal records;
- promised outcomes with measured outcomes;
- police descriptions with laboratory and court records;
- media language with its original source;
- study conclusions with protocols, data and shared-cohort lineage;
- procurement claims with contracts, purchase orders, invoices and payments;
- seizure announcements with subsequent price and availability indicators;
- and program participation with longer-term outcomes and opportunity costs.

4.5 A missing-record generator

Uncertainty is converted into a retrieval task. The system can identify the exact record required to strengthen or falsify a proposition: a toxicology report, laboratory certificate, court exhibit, unredacted contract, evaluation, briefing note, raw spectrum, data dictionary, analysis code, charge disposition, procurement matrix or communication-approval chain.

4.6 A courtroom-readiness layer

For each significant proposed finding, the system identifies the underlying exhibits, authenticity issues, chain of custody, hearsay concerns, records custodians, transformations applied to data, required expert disciplines, competing interpretations, unresolved limitations and admissibility risks.

4.7 A card deck as a navigation layer

The Card Deck is not a second source of truth. It is a presentation and navigation layer over the Evidence Deck. A card should link a reader to a person's documented public role, relevant questions, sources, qualifications and publication status. It should never become an unsupported score of guilt.

5 Current state and scope after ED-LIVE-01

5.1 Authoritative lineage and verification

ED-LIVE-01 used the private ED-ALBPSN-01 authoritative integration archive as its parent. The required archive and integrated-database SHA-256 values matched; the 112,345,088-byte SQLite database returned integrity ok and zero foreign-key violations.

```
Parent archive SHA-256
bea0f3a7ad2c4ccc2006422dc2dc771f40e4b12bd0af38ce93d9a43773883de8
Integrated database SHA-256
a2f2067811714f72b023c092ab00346ce59803a3606e7d921695f412ddab71a3
```

The original archive remained immutable. ED-LIVE-01 built derived private and public products from a verified copy.

5.2 The counts that must not be confused

Metric	Current count	Meaning
Integrated SQLite/IDSUP source-record rows	674	Rows in the IDSUP source master; not the total corpus
Canonical source works	1,044	Deduplicated intellectual or documentary works
Source-version/crosswalk records	1,269	Versions and lineage-preserving manifestations
Namespace-crosswalk rows	15,665	Links among project identifiers; not source records
Source-family rows	126	Evidence-family groupings for independence analysis

The former white paper cited the earlier ED-P24 checkpoint at 1,174 source rows, 970 canonical works, roughly 13,130 graph objects and roughly 14,300 crosswalk rows. Those figures described an older checkpoint and partially different definitions. Version 1.2 uses the later ED-ALBPSN-01 and ED-LIVE-01 accounting. The earlier 89 unresolved identifiers were not documented as having a one-to-one reconciliation with the current unresolved queues and therefore should not be presented as a current total.

5.3 Question engine and finding synthesis

The private ED-P25QE question engine contains 405 questions:

Answer state	Count
Answered	11
Partially answered	137
Conflicting	1
Unresolved	256

ED-ALBPSN-01 attached new evidence to 81 questions but moved none across its answer-state threshold because the new material did not satisfy the relevant missing-primary-record, identity, independence, transaction or outcome requirement. That is a useful result: more material did not automatically become a stronger conclusion.

The ED-P25QE synthesis preserved 546 atomic finding candidates, 296 proposed findings, 250 do-not-state items and 173 contradictions. These remain private analytical states, not public findings. Ninety-nine high-defamation-risk candidates were held in the earlier synthesis. No current paper should collapse those categories or publish them as conclusions.

5.4 Expanded substantive scope

5.4.1 Legal and enforcement records

The current integration contains:

- 34 adjudicative legal cases;
- seven separate charge or enforcement records;
- 34 case-document records;
- 34 party-allegation rows;
- 68 court finding, holding or disposition rows;

- 38 laboratory-method rows;
- and 19 seizure, exhibit or continuity rows.

Charges, police statements and party allegations remain separate from judicial findings. Missing dockets, exhibits, laboratory certificates and chain records remain queued.

5.4.2 Business, pharmacy and safer supply

The current integration contains:

- 211 pharmacy census records;
- 193 drug-product and DIN-history rows;
- 14 safer-supply or related program rows;
- 47 contract, grant or procurement records;
- 12 corporate-registry records;
- and 36 lobbying, public-relations or communications records.

No pharmacy in the 211-record census is directly confirmed by the present package as a prescribed-safer-supply dispenser. That is a missing-primary-record state, not proof that none exists. The classification is 29 PSS_NOT_ESTABLISHED, 32 SEARCHED_NO_PSS_EVIDENCE and 150 UNRESOLVED.

No newly obtained purchase order, invoice or payment confirmed a program-level transaction chain. Product availability does not prove purchase. A DIN holder is not automatically a manufacturer, wholesaler, dispenser or program supplier. OAT services do not by themselves establish prescribed-safer-supply dispensing. Funding does not by itself establish control.

5.4.3 School and campus naloxone

The register contains 24 program records across Canadian and United States jurisdictions and multiple sectors, including provincial, territorial, district, campus, residence, curriculum and emergency-response programs. B.C. sources explicitly reported 1,600 post-secondary kits.

The register deliberately separates kit quantity, sites, staff training, student reach, demonstrations, competence, administrations, clinical outcomes and effectiveness. The current package has zero confirmed program-level product/DIN links, zero confirmed school-program purchase orders and zero independently evaluated effectiveness records. Distribution and potential reach should not be described as proven effectiveness.

5.4.4 Contaminated supply and drug checking

The event-level contaminated-supply module contains 27 events. Its tables include 55 sources, 69 substance findings, 38 test-method records, eight first-reference records, five death records and an explicit unresolved-record queue. The scope includes early Canada and United States references, counterfeit or fake pill claims, mixtures described as laced or contaminated, seizures, charges, deaths and drug-checking statements.

The system separates what a person intended to obtain, what a seller or package represented, what a screening instrument indicated, what confirmatory analysis established, what toxicology detected and what a coroner identified as causal or contributory.

5.5 The five private 52-card rosters

ED-LIVE-01 replaced the earlier four-suit concept with five separate person-card views:

Deck	Candidate pool	Eligible	Selected	Reserve	Excluded
Research and Science	60	59	52	7	1
Law Enforcement	60	60	52	8	0
Public Institutions	60	59	52	7	1
Business and Pharmacy	269	269	52	8	209

Deck	Candidate pool	Eligible	Selected	Reserve	Excluded
PR, Communications and Advocacy	60	54	52	2	6

Organizations were not used to fill person slots. The 260 selected positions resolve to 193 unique people. Sixty people appear in more than one selected deck, creating 67 overlap positions. A cross-deck appearance is permitted only where the candidate records a separately documented role for each deck.

Each selected packet contains a stable person and candidate ID, documented role, organization, applicable date, reason for relevance, exact source IDs and locators, source-state limitation, proposed neutral profile text, prohibited inferences, unresolved questions, legal/privacy review state and publication recommendation.

Every selected packet currently has one cited source family. That can support a role classification. It is not independent corroboration of conduct. All 260 person cards remain private and owner-approval required.

5.6 The public read model

The sanitized public read model is intentionally small. It contains:

- one dataset-version row;
- 21 defined corpus metrics;
- five methodology entries;
- seven status definitions;
- five deck definitions;
- one release-hold ledger entry;
- and zero source, question, result, entity, relationship or person-card rows.

This is not an empty failure. It is proof that the publication gate worked. The schema and aggregate definitions can be deployed without exposing held people, findings or private editorial material.

6 The complete source-record map

The database becomes powerful only when it connects public propositions to the records that can support, qualify or disprove them.

Record family	Core questions	Main caution
Coroner and medical-examiner files	What was identified as causal or contributory?	Detection is not automatically causation; privacy law applies
Postmortem toxicology	Concentrations, metabolites, co-occurring substances, methods	Interpretation requires toxicology expertise and tolerance/timing context
Hospital and EMS records	Presentation, naloxone, diagnosis, outcome	Individual records require lawful authority; aggregate data is generally preferable
Drug-checking raw data	Instrument, spectrum, possible identifications and exclusions	Screening is not equivalent to confirmatory laboratory analysis
Instrument validation and QC	Sensitivity, specificity, detection limits, calibration and failure	Public summaries may omit limitations
Police seizure files	What was seized, where, when and under what authority?	Operational sensitivity and redactions are likely
Forensic certificates and bench records	Identity, concentration, sampling, uncertainty and date	A tested sample may not represent an entire seizure

Record family	Core questions	Main caution
Chain of custody	Can the tested sample be connected to the alleged exhibit?	Continuity defects can materially reduce evidentiary weight
Charges, pleadings and informations	What was alleged and against whom?	An allegation is not a finding
Judgments and dispositions	What was admitted, proven, withdrawn, stayed or rejected?	Silence in a disposition may not resolve every factual issue
Research protocols and registrations	What was planned, measured and changed?	Deviations may be legitimate but should be explained
Data dictionaries and analysis code	How were variables, exclusions and denominators defined?	Small coding choices can alter results
Funding, authorship and conflicts	Who financed, designed, analysed, wrote and approved?	Funding and authorship are relationships, not proof of purchased results
Shared-cohort lineage	How many publications represent genuinely independent evidence?	Publication count can exaggerate independent replication
Program contracts and evaluations	Deliverables, targets, budgets, adverse events and outcomes	A contract shows an obligation, not performance
Procurement and payments	What was purchased, from whom, at what price and in what quantity?	Product availability does not prove purchase or use
Corporate, lobbying and PR records	Ownership, influence efforts and message development	Lawful business and lobbying are not misconduct
Nonprofit governance and finance	Directors, related parties, compensation and restricted funds	Status and compensation do not establish improper motive
Complaints and incident records	What harms were reported and what corrective action followed?	Complaint volume depends partly on reporting access
Archived policy and communications	What staff were instructed to do and what changed over time?	A missing page can result from ordinary web maintenance
Witness and lived-experience accounts	How did systems operate and which records should exist?	Memory and perspective require corroboration

6.1 The bridge-record concept

The most important record is often the one linking an outcome to a person's knowledge and decision. Suppose a program is publicly described as successful. An internal evaluation shows no improvement on the primary outcome. Senior officials receive it. Communications staff are directed to omit the primary outcome and emphasize participation. A revised release is approved, and further funding is obtained using that description.

The release alone does not prove dishonesty. The evaluation alone does not prove dishonesty. The funding alone does not prove dishonesty. The receipt, instruction, approval chain, timing, stated duty, later representation and response may provide the bridge. Alternative explanations still have to be tested.

7 Testing public narratives as falsifiable propositions

7.1 A major seizure disrupted supply

The public proposition should generate observable predictions: price may increase, reported availability may fall, acquisition time may rise, composition may change, replacement suppliers may emerge, and effects should occur in a plausible time window.

Required records include exact composition and quantity, the date material left circulation, distribution stage, market-share evidence, contemporaneous prices, availability observations, drug-checking data, replacement-supplier intelligence and a pre-specified comparison period.

Permitted outcomes include established disruption, temporary local disruption, displacement without net reduction, no measurable disruption, insufficient data, or a public statement stronger than the available evidence.

7.2 Contamination or counterfeit supply caused a death

The inquiry should distinguish:

- what the person intended to obtain;
- what the seller, package or communication represented;
- whether residue or a retained sample was tested;
- what toxicology detected;
- what the coroner identified as causal or contributory;
- route of administration and other substances;
- and the origin of words such as laced, contaminated, counterfeit or fake.

Fentanyl detection does not by itself establish that the person believed they were consuming another drug or that fentanyl alone caused death.

7.3 A program reduced mortality

The file should identify eligibility, enrollment denominator, comparison group, loss to follow-up, death ascertainment, cause classification, treatment exposure, concurrent interventions, analysis code, adverse outcomes and cost.

No deaths at a program site does not necessarily mean the program reduced mortality. The relevant outcome may be what happened to participants over an appropriate period compared with a credible alternative.

7.4 The science was independently replicated

The project should test whether papers use the same cohort, investigators, registry, trial or data platform; whether a guideline cites several publications from one trial; whether negative or null results were published; and whether the primary outcome changed.

Ten papers derived from one cohort are ten publications. They are not necessarily ten independent confirmations.

7.5 Public money was used effectively

The project should calculate direct, administrative, capital, pharmaceutical and device cost; cost per eligible participant; cost per completed intervention; cost per durable outcome; downstream health and justice costs; and comparative performance against credible alternatives.

High cost is not proof of waste. Low cost is not proof of value. The question is whether the claimed benefit is supported and whether a better alternative was reasonably available.

7.6 A school naloxone initiative was effective

The project should distinguish kits purchased, kits placed, sites covered, people attending training, competence demonstrated, real administrations, clinical outcomes and longer-term effectiveness. A launch release or distribution total establishes activity, not educational mastery or outcome.

8 Evidentiary method

8.1 One evidence file per proposition

A significant proposition should have:

1. a stable proposition ID and neutral wording;
2. the public narrative and material speakers;
3. primary and secondary supporting records;
4. material contrary records;
5. source-independence analysis;
6. missing records and retrieval tasks;
7. alternative explanations;
8. a falsification condition;
9. required expert disciplines;
10. method, assumptions, exclusions and code;
11. confidence and evidence-state assessment;
12. privacy, legal and publication status;
13. right-of-reply status;
14. and a statement of what would change the finding.

8.2 Evidence classes

Primary record. Created as part of the event, transaction, test or decision, such as a laboratory certificate, signed contract, invoice, original dataset, contemporaneous email or court order.

Official assertion. A public body's statement. It proves the statement was made, subject to authentication. It does not automatically prove the stated facts.

Independent corroboration. Evidence that does not inherit the same press release, cohort, data source or unverified premise.

Derived analysis. A calculation or inference from records. The formula, assumptions, exclusions and code should be available.

Allegation. A claim in a complaint, charge, pleading, interview or unresolved account.

Judicial or regulatory finding. A finding by a competent body, recorded with scope, disposition and appeal status.

Unresolved. A proposition with incomplete, inaccessible or materially conflicting evidence.

8.3 Independence analysis

Two records are not independent merely because they appear in different places. The system asks whether reports arose from the same release, papers used the same cohort, experts repeated the same unverified premise, agencies exchanged wording, or apparent consensus resulted from one dataset being published repeatedly.

8.4 Alternative explanations

Before drawing an adverse inference, the project should test ordinary error, ambiguous language, outdated information, privacy or legal limits, scientific uncertainty, different definitions, emergency decision-making, resource constraints, staff turnover, sincere policy disagreement, statistical noise, confounding and unforeseen market adaptation.

8.5 Falsification

Every strong conclusion should state what evidence would weaken or reverse it. A proposition that cannot identify a falsification condition is probably too vague, too political or too broad for a defensible finding.

8.6 Confidence ladder

	Level	Label	Minimum meaning
	0	Unassessed	Collected but not reviewed
	1	Reported	One or more sources made the claim
	2	Source documented	Accurately attributed to an identifiable source
	3	Primary-record supported	A primary record supports the scoped fact
	4	Independently corroborated	Genuinely independent sources support it
	5	Strong inference	The total record substantially favours one explanation
	6	Alternatives tested	Principal reasonable alternatives were materially tested
	7	No reasonable competing explanation presently identified	Reproducible, internally consistent and resistant to identified alternatives
	8	Adjudicated	A court or competent tribunal made a final finding, with scope and appeal status recorded

The phrase beyond a reasonable doubt belongs to criminal adjudication. The database can seek rigorous support for narrow propositions, but it should not imitate a verdict.

9 From database to courtroom

The Evidence Deck will ordinarily be an organizing system. The underlying records and witnesses are what may become evidence.

Sections 24 and 30 of the Canada Evidence Act provide routes for certain certified public documents and business records, subject to the statutory conditions and the other rules of evidence.¹ They do not make every download, screenshot or database entry automatically admissible.

The project should preserve the original file, acquisition source and time, metadata, hash, certification where available, custodian information, completeness, version and transformation history.

Canadian courts use the Mohan framework for expert evidence: relevance, necessity, absence of an exclusionary rule and a properly qualified expert.² The Supreme Court has also emphasized the expert's duty to be fair, objective and non-partisan.³

An independent expert should be able to inspect source-selection rules, review supporting and contrary records, reproduce calculations, test assumptions, identify limits and disagree with the provisional conclusion.

9.1 Courtroom-readiness test

A proposed finding should be able to answer:

¹Justice Laws Website, Canada Evidence Act, s 24 and full text including s 30, consulted August 29, 2026. Recheck statutory currency before legal use.

²Supreme Court of Canada, R v Mohan, [1994] 2 SCR 9.

³Supreme Court of Canada, White Burgess Langille Inman v Abbott and Haliburton Co, 2015 SCC 23.

- Can every material number be traced to a record?
- Can the record be authenticated?
- Is the original preserved?
- Is the dataset complete enough for the proposition?
- Were exclusions and transformations disclosed?
- Can the analysis be rerun?
- Are reasonable alternative methods presented?
- Is the expert qualified and independent?
- Is contrary evidence included?
- Is the conclusion narrower than the evidence?
- Can the evidence survive hostile cross-examination without relying on Scott Gilbert's credibility or ChatGPT's authority?

10 Identifying responsibility without prejudging guilt

The phrase bad actor should remain an investigative category, not a verdict. A person should not be classified adversely merely because they worked for a health authority, police service, university, nonprofit, pharmacy or pharmaceutical company; received a salary or grant; advocated a policy; appeared in a release; served on a board; made an error; or knew someone in the graph.

10.1 Responsibility states

The project should keep dimensions separate rather than use an automated culpability ladder:

- factual accuracy on the information then available;
- completeness and material context;
- applicable duty or professional standard;
- evidence of actual receipt and understanding;
- evidence of reasonable reliance, ambiguity or changed information;
- materiality and harm;
- purpose, benefit and state of mind;
- correction opportunity and response.

Permitted descriptive outcomes include supported and materially complete; accuracy unresolved; later shown inaccurate but reasonably supported when made; inaccurate with basis unresolved; incomplete with duty and materiality unresolved; or contrary information demonstrably received while understanding or reliance remains unresolved.

Contradiction is not automatically a lie. Should have known is not knew. Public availability does not prove receipt or understanding. Silence is not an admission. Institutional failure does not establish individual culpability. Benefit is not guilt.

10.2 Requirements before a serious individual finding

The file should identify the person, role and duty; the specific act or omission; date; information available; representation or decision; materiality; contrary information; evidence of knowledge or recklessness; any relevant benefit; causal connection; alternative explanations; corroboration; response; and legal review.

No automated system should characterize conduct as criminal. The database can prepare an issue file for qualified counsel and a competent authority.

11 Accountability pathways

Not every adverse conclusion belongs in criminal court. The pathway should match the evidence.

- **Correction and clarification:** unsupported or inaccurate public language.
- **Access review:** missing records, delay or contested redactions.
- **Audit and value-for-money review:** spending, procurement, controls and outcomes.

- **Professional or regulatory complaint:** conduct within a professional duty.
- **Civil advice or proceeding:** where facts may support a recognized cause of action and procedural requirements can be met.
- **Police or prosecution referral:** where admissible evidence reasonably indicates an offence.
- **Special-prosecutor request:** where the statutory and policy conditions are engaged and public confidence requires independence.
- **Private prosecution advice:** only after careful legal review, element-by-element analysis and disclosure of exculpatory material.

British Columbia's current charge-assessment standard asks whether there is a substantial likelihood of conviction and, if so, whether the public interest requires prosecution.⁴ Section 7 of the Crown Counsel Act authorizes the Assistant Deputy Attorney General to appoint a special prosecutor when considered in the public interest.⁵ Section 507.1 of the Criminal Code provides judicial screening for privately laid informations and a role for the Attorney General.⁶

These pathways should never be used as campaign threats or leverage. The Evidence Deck should be capable of preparing a neutral referral. It should not promise charges before the evidence is obtained and reviewed.

12 Governance and publication gates

12.1 Separation from campaign messaging

Because Scott Gilbert intends to seek public office, the project requires a firewall among private research, legal review, publishable findings, campaign commentary, fundraising, opposition research and personal disputes.

A politically useful conclusion must not receive a lower threshold. An inconvenient correction must not be suppressed. If campaign messaging uses a project conclusion, it should link the source package, limitations and date.

12.2 Independent review

Major findings should be reviewed by people with relevant expertise who are not required to agree with the project's political objectives. A panel could include an evidence lawyer, statistician, toxicologist, public-health researcher, forensic accountant, investigative journalist, lived-experience representative and privacy or ethics specialist. Conflicts should be disclosed.

12.3 Right of reply

Before a material adverse finding about an identifiable person or organization:

1. provide the exact proposed finding and its supporting records;
2. ask focused questions and allow a reasonable response period;
3. preserve the complete response;
4. correct errors and publish material rebuttal fairly;
5. record non-response without treating it as an admission.

12.4 Corrections

The project should publish the original wording, corrected wording, date, reason, impact on linked findings and version history. A system demanding corrections from institutions must correct itself visibly.

⁴Province of British Columbia, BC Prosecution Service and Charge Assessment Guidelines, consulted August 29, 2026.

⁵Crown Counsel Act, s 7 and Province of British Columbia, About Special Prosecutors.

⁶Justice Laws Website, Criminal Code, s 507.1, consulted August 29, 2026.

12.5 Release gates

No serious adverse item should be public until it passes:

- identity resolution;
- source and provenance verification;
- evidence-family and independence review;
- contradiction and alternative-explanation review;
- privacy and security review;
- defamation and legal review;
- right of reply;
- proportionality and public-interest review;
- explicit owner approval;
- and a release-ledger event that can be audited.

The ED-LIVE-01 public model contains zero person cards because these gates have not been completed. That is the correct result.

13 Live system architecture

13.1 Two physically separate databases

The private editorial model contains the authoritative evidence graph, unresolved identities, candidate scores, evidence packets, reviewer notes, release recommendations and omissions. It must never be placed in a public/static directory or exposed by an endpoint.

The public read model contains only approved and sanitized records. Its table families cover dataset versions, defined counts, methodology, status, sources, questions, results, entities, roles, relationships, decks, cards, corrections, right of reply and release history. Empty content tables remain empty rather than weakening a gate.

13.2 Read-only API

The handoff specifies 15 GET routes for status, corpus, sources, questions, results, entities, relationships, decks, cards, corrections and methodology. The contract requires bounded pagination, allowlisted filters and sort fields, parameterized queries, safe errors, versioned responses, no arbitrary SQL, no raw database download and no private graph traversal.

All 23 current contract and security checks passed locally. This verifies the handoff, not the live site.

13.3 Minimal Replit handoff

The remaining implementation should be one narrow pass in the existing app titled Evidence deck:

1. verify the public handoff ZIP and its checksums;
2. create a recoverable source checkpoint;
3. add the smallest compatible server-side SQLite adapter;
4. keep the database outside public/static assets;
5. wire the supplied read-only API contract;
6. replace hard-coded data with API calls;
7. add the routes and explicit held states;
8. run the supplied tests and ordinary application tests;
9. verify that no private artifact is in the deployment;
10. republish once and confirm the live dataset version.

Replit should not research, rank candidates, select cards, rewrite narratives or make publication decisions.

13.4 A genuinely living document

A PDF can be authoritative as a versioned snapshot, but it is not live. The live document should be a database-backed /project-status or /methodology/project-status page generated from the approved public read model.

The page should draw its counts, question states, deck release states, dataset hash, generated date, limitations, corrections and milestone status directly from the current approved dataset. Every new checkpoint would create:

1. a new immutable private archive;
2. a newly generated sanitized public read model;
3. parity, privacy and contract tests;
4. an explicit approval event;
5. a new live page version;
6. and archived PDF snapshots such as Version 1.2.

Aggregate project status can update after automated QA and an approval rule. Person-level or materially adverse content should never auto-publish merely because the research database changed.

14 Implementation roadmap from the current checkpoint

14.1 Phase 1 - Preserve and review ED-LIVE-01

- Store the public and private ZIPs, sidecars and verification reports.
- Review the five private rosters and evidence packets.
- Check cross-deck overlaps and role dates.
- Keep scores private and correct any misclassification.

14.2 Phase 2 - Approve neutral public metadata

- Review the neutral source-metadata ledger separately from person cards.
- Approve count definitions, methodology and non-adverse source entries in auditable batches.
- Leave adverse, ambiguous and person-level material held.

14.3 Phase 3 - Implement the live read model

- Upload the public handoff once to the existing Evidence deck app.
- Wire the read-only API and data-driven frontend.
- Verify the deployed dataset version, route behaviour, counts and privacy boundary.

14.4 Phase 4 - Retrieve bridge records

Prioritize records most likely to change a question state: court exhibits and laboratory certificates; chain of custody; program contracts and evaluations; procurement transactions; toxicology methodology; communications approval chains; shared-cohort data; and governance records.

14.5 Phase 5 - Reproducible analysis and expert review

Preserve raw data, document cleaning, publish code where lawful, run sensitivity analyses, record null results and give neutral packages to independent experts.

14.6 Phase 6 - Findings, reply and publication

Prepare narrow findings with supporting and contrary records, confidence, limitations, falsification conditions, accountability route and full right-of-reply record. Release in layers: source catalogue, methodology, neutral timelines, claim lineage, financial maps, tested findings, unresolved questions, responses and corrections.

15 Public participation

No single candidate, researcher, journalist or AI system can complete the record. Contributors may include people who use or formerly used drugs, families, residents, current and former public servants, police officers, paramedics, clinicians, pharmacists, researchers, nonprofit workers, journalists, lawyers, auditors, procurement specialists, archivists and whistleblowers.

Useful contributions include original documents, stable URLs, archived pages, contracts, invoices, policies, meeting minutes, datasets, research papers, court file numbers, correction histories, regulator findings, firsthand accounts identifying records, alternative explanations and evidence contradicting the project's theories.

The project should welcome evidence that proves Scott Gilbert wrong. That is a test of seriousness.

The intake system must not promise anonymity it cannot guarantee. It should use clear consent, security controls, metadata warnings, separation of identity from content, access limits and legal advice concerning source protection. It should not solicit privileged or unlawfully obtained information.

16 What success would look like

The Evidence Deck will have succeeded when it can answer, with traceable records:

- which fentanyl claims began with laboratory evidence and which began with communications language;
- which deaths are supported by primary toxicology and causal findings, and which remain based on suspicion or repetition;
- which seizures measurably affected price or availability;
- which programs produced durable outcomes rather than activity counts;
- which studies represent independent evidence families;
- which public statements exceeded the evidence available when made;
- which contracts delivered what was promised;
- which organizations and businesses received public money;
- which relevant interests were disclosed;
- which errors were corrected;
- which unsupported claims persisted;
- which conduct may warrant audit, regulation, civil advice or independent criminal investigation;
- and which allegations fail when the complete evidence is examined.

Some people viewed with suspicion may be cleared. Some apparently alarming relationships may prove ordinary. Some institutions may be shown to have overstated their evidence. Other records may become more serious when knowledge, decision, concealment, benefit and harm are linked. The project must be willing to produce all of those results.

17 Conclusion: what a new sheriff actually does

"There's a New Sheriff in Town" is not a promise of vigilante justice. It is a promise that public authority will be asked to produce the record behind its claims.

A new sheriff does not declare people guilty from a podium. A new sheriff asks:

- What was claimed?
- What was known?
- What was measured?
- What was purchased?
- Who approved it?
- Who benefited, and what can that benefit actually establish?
- Who objected?
- What was omitted?
- What happened afterward?
- Can the conclusion be reproduced?

- Can an independent expert defend it?
- Can the underlying evidence survive court?

Where the record shows a good-faith mistake, it should be corrected. Where it shows poor policy, the policy should change. Where it shows waste, spending should be reviewed or stopped. Where it shows professional misconduct, the regulator should act. Where it supports a civil wrong, those affected should obtain legal advice. Where it credibly indicates criminal conduct, the matter should be independently investigated according to law.

And where the evidence does not support an accusation, the Evidence Deck must say so clearly, even when that conclusion is politically inconvenient.

No accusation without evidence.
No public claim without its source.
No spending without measurable accountability.
No finding without contrary evidence.
No person condemned merely by association.
No institution exempt from producing the receipts.

18 Appendix A: ED-LIVE-01 completion checkpoint

Item	Verified state
Controller result	PASS_WITH_GAPS
Public/private database separation	PASS
Public API contract tests	23/23 PASS
Deterministic repeat build	13/13 selected artifacts byte-identical
Public handoff ZIP	Verified; full hash below
Private review ZIP	Verified; full hash below
Replit Agent calls	0
Live site changed or republished	No
Public person cards	0
Owner approval required	Yes

Public handoff ZIP ED-LIVE-01_replit_handoff_2026-08-29.zip
SHA-256 ec730d0c36315fca86e79f72908b0ff9de5b9af5a10ce0136ddccfc2423d2f81
Private review ZIP ED-LIVE-01_private_roster_review_2026-08-29.zip
SHA-256 344303ac5415ab7016d44c1a4789618439e787891a4c8cd66ba2e8c3878f3698

19 Appendix B: model finding format

Finding ID: Stable identifier.

Neutral proposition: One factual proposition capable of being true or false.

Public narrative: Exact quotation or accurate summary of the representation being tested.

Narrow finding: Conclusion no broader than the record.

Evidence state and confidence: Defined classification, not an impression.

Supporting primary records: Stable IDs and locators.

Independent corroboration: Genuine independence analysis.

Contrary evidence: All material contrary records.

Missing records: Exact records that remain unavailable.

Alternative explanations: Each principal explanation and the result of testing it.

Method: Calculations, assumptions, exclusions, code and sensitivity analysis.

Expert review: Reviewer, qualifications, instructions and conflicts.

Right of reply: Questions, response and resulting revisions.

Legal characterization: Factual, analytical, regulatory, civil or issue flagged for counsel. No automated criminal label.

Publication status: Private, held, revise, remove, approved or published.

Falsification condition: Evidence that would materially change the conclusion.

20 Selected project records and legal authorities

1. Evidence Deck Project, ED-ALBPSN-01 authoritative integration package, private verified integration and its completion receipt, integration report and quality-control report.
2. Evidence Deck Project, ED-LIVE-01 completion receipt, integration report, quality-control report, project-state/lineage record, public read model, private review model and verified handoff packages, August 29, 2026.
3. Canada Evidence Act, RSC 1985, c C-5, especially ss 24 and 30.
4. R v Mohan, [1994] 2 SCR 9.
5. White Burgess Langille Inman v Abbott and Haliburton Co, 2015 SCC 23, [2015] 2 SCR 182.
6. BC Prosecution Service, Charge Assessment Guidelines and public description of the two-part charge-assessment standard.
7. Crown Counsel Act, RSBC 1996, c 87, s 7, and BC Prosecution Service information on special prosecutors.
8. Criminal Code, RSC 1985, c C-46, s 507.1.