

FIELD MANUAL · OCTOBER 2026

Legal OSINT Espionage

A 2026 Field Manual for Post-Predatory Reconnaissance &
Sovereign Intelligence

Marie-Soleil Seshat Landry

CEO · LANDRY INDUSTRIES & MARIE LANDRY'S SPY SHOP

Marie Landry's Spy Shop · marielandryspyshop.com · Moncton, New Brunswick, Canada · ORCID
0009-0008-5027-3337

Prepared with AI assistance from Google Gemini and Anthropic Claude — see AI Disclosure.

Legal OSINT Espionage

A 2026 Field Manual for Post-Predatory Reconnaissance & Sovereign Intelligence

Marie-Soleil Seshat Landry

Publisher: Marie Landry's Spy Shop (marielandryspyshop.com), a division of Landry Industries (LandryIndustries.ca)
Moncton, New Brunswick, Canada

First edition, October 2026.

© 2026 Marie-Soleil Seshat Landry. All rights reserved except where otherwise noted. Short excerpts may be quoted for review, teaching and research with attribution.

Not legal advice. This manual describes general practices for lawful open-source research. Laws differ by jurisdiction and change often. Readers should consult a qualified lawyer before undertaking any investigation with legal, employment, or reputational consequences. Neither the author nor the AI systems used in preparation are lawyers.

SEO keywords: OSINT 2026; Legal OSINT Espionage; Marie-Soleil Seshat Landry; Marie Landry Spy Shop; Ethical Intelligence; Post-Predatory Intelligence; AI OSINT Reconnaissance; Sovereign Investigative Framework.

AI Disclosure & Verification State

Item	Detail
Author / Principal Architect	Marie-Soleil Seshat Landry (ORCID 0009-0008-5027-3337), CEO, Landry Industries & Marie Landry's Spy Shop. Responsible for the concept, the 10-step methodology, the chapter architecture and final editorial judgement.
AI assistant — outline stage	Google Gemini (AI-assisted). Used to draft the original outline, executive abstract, applied case study and initial reference list.
AI assistant — manuscript stage	Anthropic Claude (Claude Opus 5.5, via claude.ai), 5 October 2026. Expanded the outline into the full chapter text, wrote the HTML layout and produced this PDF, and checked the reference list against live web sources.
Verification state	Partially verified. Every entry in the reference list was checked by live web search on 5 October 2026, then re-audited line by line the same day. Six entries in the Gemini-stage list were inaccurate and have been corrected or replaced (see "Reference Corrections" at the end of the book). Statements about current law (Canada's Bill C-36, the EU AI Act Digital Omnibus) were confirmed against legal-commentary sources dated June–September 2026. The body text's methods and recommendations are AI-drafted explanatory prose and have <i>not</i> been independently peer-reviewed.
Human responsibility	AI output can contain errors. The author remains responsible for reviewing, accepting or amending all content before redistribution.

NOTE ON THE EARLIER DRAFT

The Gemini-stage draft described itself as "factually cross-referenced against live institutional databases." The live checks for this edition found that claim did not hold for several references, so it has been replaced with the accurate status above. In keeping with this manual's own axiom, the disclosure reports what was actually verified.

Contents

Executive Abstract..... 5

Preface: How to Use This Manual 6

SECTION I — THE FOUNDATIONS OF ETHICAL INTELLIGENCE

1. The Post-Predatory OSINT Paradigm 1

2. Legal Boundaries & Strategic Jurisdictions..... 1

SECTION II — THE 10-STEP INVESTIGATIVE METHODOLOGY

3. Steps 1-3: Problem Formulation & Hypothesis Design..... 1

4. Steps 4-6: Collection Protocols & Analytical Rigor 1

5. Steps 7-10: Decision Loop, Verification & Dissemination 1

SECTION III — ADVANCED TACTICAL OSINT

6. GEOINT & Digital Footprint Reconstruction..... 1

7. Corporate Forensics & Financial Lineage 1

8. AI-Powered Intelligence Integration..... 1

SECTION IV — COUNTER-SURVEILLANCE & OPSEC

9. Digital Hygiene for the Sovereign Researcher 1

10. The Future of Sovereign Reconnaissance (2030 Horizon) 1

BACK MATTER

Applied Case Study: The 10-Step Framework End to End 1

Appendix A: Field Templates 1

Appendix B: Glossary 1

References (Verified 5 October 2026) 1

Reference Corrections Log 1

About the Author 1

Executive Abstract

Between 2024 and 2026 the ground under open-source intelligence (OSINT) shifted. Three forces collided: privacy regulators made clear that "publicly accessible" does not mean "free to harvest"; platforms hardened themselves against automated scraping; and multimodal AI systems made it cheap to collect, summarise and fabricate information at industrial scale. Tactics that were routine five years ago — bulk profile scraping, unattributed screenshots, AI summaries pasted into reports — now carry legal, ethical and evidentiary risk.

Legal OSINT Espionage is a methodological text for people who investigate without state backing: entrepreneurs doing due diligence, citizen researchers and journalists examining corporate conduct, and privacy-conscious analysts who want rigorous results without crossing legal or ethical lines. Its core is a ten-step method adapted from the scientific method: observe an anomaly, define the intelligence requirement, state a falsifiable hypothesis, design the collection, capture facts verbatim, analyse with structured techniques, test the hypothesis, conclude only what the evidence supports, make the work replicable, and report transparently.

The word "espionage" in the title is deliberately provocative and deliberately limited. Nothing in this manual involves deception to obtain non-public information, unauthorised access to systems, covert surveillance of private persons, or circumvention of technical access controls. The "espionage" here is the discipline of the craft — patience, structure, tradecraft and source protection — applied entirely to lawful, open material, with the primary focus on organisations and the exercise of power rather than on private individuals.

Facts are fixed points. Inferences are hypotheses to be tested.
Nothing enters a conclusion that cannot be traced to a source.

Preface: How to Use This Manual

This book is organised in four sections. **Section I** sets the ethical and legal frame; read it first even if you are experienced, because the law changed materially in 2026. **Section II** is the core method, one step at a time, each with a purpose, a procedure, outputs and common failure modes. **Section III** applies the method to specific source domains: geospatial and infrastructure data, corporate and financial records, and AI-assisted analysis. **Section IV** covers protecting yourself and your work, and looks ahead to 2030.

The back matter contains a fully worked case study that walks one hypothetical investigation through all ten steps, ready-to-copy templates, a glossary, and the verified reference list.

Three ground rules

1. **Lawful sources only.** If obtaining a piece of information would require lying about who you are, bypassing a login or paywall you are not entitled to, or using leaked credentials or breached data, it is out of scope.
2. **Organisations over individuals.** The default subject of an investigation is an entity — a company, an agency, a network of companies. Where individuals appear, it is in their public, professional capacity (an officer of a company, a named patent assignee), and you collect the minimum needed.
3. **Every claim carries its receipt.** A finding without a captured, timestamped, hash-verified source is an opinion, and is labelled as one.

A note on terminology

The author's frameworks use several terms of art. *Ethical Pathfinders* are the intended readers: values-driven builders, researchers and activists. *Post-predatory* describes intelligence practice that refuses to profit from exploiting people's data or vulnerabilities. *Sovereign* describes an investigator who controls their own tools, data and conclusions rather than outsourcing judgement to platforms or opaque AI. *Non-state-aligned* means working independently of any government's intelligence apparatus.

SECTION I

The Foundations of Ethical Intelligence

Before collecting anything, decide what you will not do.

The Post-Predatory OSINT Paradigm

The 2026 landscape, what "ethical espionage" means, and the axiom that governs everything else in this book.

1.1 What OSINT is — and what it is not

Open-source intelligence is the disciplined collection, evaluation and analysis of publicly available information to answer a specific question. The NATO *Open Source Intelligence Handbook* (2001) described the process as discovery, discrimination, distillation and dissemination; modern practice adds verification and preservation as first-class steps, a shift codified in the *Berkeley Protocol on Digital Open Source Investigations* published by the UN Office of the High Commissioner for Human Rights and UC Berkeley's Human Rights Center.

Three distinctions matter from the start:

- **Information is not intelligence.** A company filing is information. A tested, sourced conclusion about who controls that company and why that matters to a decision is intelligence.
- **Public is not the same as open.** A social-media profile set to "public" is technically accessible, but privacy regulators in Canada and abroad have stated that personal information remains protected by law even when it is publicly visible (see Chapter 2).
- **OSINT is not hacking.** Accessing systems you are not authorised to access, using stolen credentials, or exploiting vulnerabilities is not open-source research under any definition, and it is a criminal offence in most jurisdictions.

1.2 Defining "ethical espionage"

Classic espionage involves obtaining secrets through deception, theft or betrayal. This manual borrows the *discipline* of the intelligence profession — requirements management, source evaluation, structured analysis, compartmentation, operational security — and discards its *methods* of clandestine acquisition. Ethical espionage, as used here, has five properties:

Property	Meaning in practice
Lawful	Every source is one the investigator is legally entitled to read, and collection respects terms of access and applicable privacy law.
Non-intrusive	No interaction with the subject's systems beyond what any ordinary visitor would do; no pretexting, no fake personas used to extract non-public information; no physical surveillance.
Proportionate	The depth of collection matches the importance of the question. Questions about corporate misconduct justify more than questions about a private person's habits — and the latter are generally out of scope.
Accountable	Methods are documented well enough that a third party could audit and replicate them.
Harm-aware	Before publication, the investigator considers who could be hurt by the findings — including bystanders named in documents — and minimises that risk.

"Target acquisition," in this vocabulary, means selecting the *entity or question* that the investigation will examine. It never means locating a person in order to watch or confront them.

1.3 The 2024-2026 shift

Regulators drew the line on scraping

In August 2023, Canada's Privacy Commissioner and eleven counterparts issued a joint statement on data scraping, followed in Oc-

tober 2024 by a concluding joint statement signed by sixteen authorities. Together they established the regulatory expectation that publicly accessible personal information is still covered by privacy law, that platforms must defend against unlawful scraping, and that mass scraping incidents can amount to reportable breaches. For the independent investigator, the message is direct: bulk harvesting of personal data is not a grey zone.

Canada began replacing PIPEDA

After Bill C-27 died on the order paper, the federal government introduced Bill C-36 on 15 June 2026. It would replace the privacy part of PIPEDA with a new *Protecting Privacy and Consumer Data Act* and a new regulator, the Digital Safety and Data Protection Commission of Canada. As of this edition the bill is still before Parliament and PIPEDA remains the law; Chapter 2 explains what to watch.

The EU's AI rules began to bite

The EU AI Act's transparency obligations under Article 50 have applied since 2 August 2026. The Digital Omnibus on AI (Regulation (EU) 2026/1744), in force since 27 July 2026, postponed most high-risk obligations to December 2027 and August 2028, but it did not cancel them. Analysts who publish AI-assisted output to EU audiences should expect disclosure duties.

AI changed both sides of the table

Language models now extract entities from thousands of documents in minutes. They also hallucinate citations, invent corporate officers and confidently summarise documents they have not read. Meanwhile, synthetic images, cloned voices and AI-generated "evidence" have flooded the information environment. The investigator's job now includes proving that a source is real before analysing what it says.

1.4 From standard OSINT to AI-augmented sovereign intelligence

"Sovereign" intelligence means the analyst retains control over three things: the **data** (captured and stored under their own custody, not just bookmarked on a platform that may delete it), the **tools** (understood well enough to explain how each output was produced), and the **judgement** (AI proposes, the analyst disposes). The practical transition looks like this:

Legacy habit	2026 practice
Bookmark a page	Capture the page, its URL, retrieval time and a SHA-256 hash; request an Internet Archive snapshot
Paste an AI summary into the report	Use AI to extract candidate facts, then verify each against the captured source before it enters the evidence log
Scrape every profile that matches	Define the minimum data required by the intelligence requirement and collect only that
"I found it online"	Full provenance: source, method, time, hash, analyst

1.5 The core axiom

Facts are immutable constants; inferences are testable hypotheses.
There is zero room for hallucination.

Every statement in an intelligence product belongs to one of three categories, and must be labelled accordingly:

- **Fact** — directly observed in a captured source (e.g., "The registry lists Jane Roe as director as of 3 March 2026").
- **Inference** — a reasoned conclusion from facts, stated with a confidence level and the facts it rests on (e.g., "It is likely that

the two companies share management, based on facts F3, F7 and F9").

- **Assessment gap** — an explicit statement of what is unknown.

An AI system's output is never a fact. It is, at best, a pointer to where a fact might be found.

CHAPTER 1 CHECKLIST

- Can I state the question I am answering in one sentence?
- Is my subject an organisation, or a person acting in a public capacity?
- Have I ruled out any method that involves deception, unauthorised access or covert surveillance?
- Do I have a way to capture and hash sources before I start?

Legal Boundaries & Strategic Jurisdictions

Public-domain research versus non-consensual data mapping, cross-border compliance, and how to build a defensible record.

NOT LEGAL ADVICE

This chapter is an orientation map, not a legal opinion. Laws vary by jurisdiction, by the investigator's status (journalist, licensed private investigator, business, individual) and by the purpose of the work. Engage a lawyer before any investigation whose results may be used in litigation, employment decisions, or publication about identifiable people.

2.1 Public-domain research versus non-consensual data mapping

The most important legal line in modern OSINT is not between "public" and "private" information. It is between **consulting** public records to answer a defined question and **assembling** a comprehensive profile of a person without their knowledge or consent. The first is the ordinary work of researchers, journalists and due-diligence analysts. The second — sometimes called non-consensual data mapping or dossier-building — is what privacy regulators target, and in many contexts it can also amount to harassment or stalking.

Generally defensible	High risk / out of scope
Reading a corporate registry entry to identify a company's directors	Compiling a director's home address, family members, daily routine and vehicle
Reviewing a company's public filings, patents and press releases	Bulk scraping social-media profiles of a company's employees

Analysing a domain's historical DNS and certificate records	Probing or scanning the company's servers for vulnerabilities
Searching court records for litigation involving an entity	Using breach dumps or stealer logs to find credentials or personal data
Citing a public official's on-record statements	Creating fake accounts to befriend staff and extract information

2.2 Canada: PIPEDA today, Bill C-36 tomorrow

The *Personal Information Protection and Electronic Documents Act* (S.C. 2000, c. 5) governs the collection, use and disclosure of personal information in the course of commercial activity. If you sell investigation services, you are engaged in commercial activity. Key consequences:

- **Purpose limitation.** Personal information may be collected only for purposes a reasonable person would consider appropriate in the circumstances (s. 5(3)).
- **Consent.** PIPEDA generally requires knowledge and consent; exceptions exist for certain publicly available information, but they are narrow and defined by regulation. A public social-media post is not automatically within the exception.
- **Journalistic, artistic and literary purposes** are excluded from PIPEDA's application where collection is *solely* for those purposes — an important carve-out for genuine journalism, but one that does not cover a commercial due-diligence report.
- **Provincial overlay.** Quebec (Law 25), Alberta and British Columbia have their own private-sector statutes. New Brunswick businesses are generally covered by PIPEDA.

What Bill C-36 would change

Introduced 15 June 2026, Bill C-36 would enact the *Protecting Privacy and Consumer Data Act*, recognise privacy as a fundamental right, transfer private-sector oversight to a new Digital Safety and Data Protection Commission of Canada, add significantly larger penalties, and create a broader private right of action for affected

individuals. It remains a bill, not a law, at the time of writing; monitor its progress through the LEGISinfo service. For investigators, the direction is unambiguous: enforcement risk for careless handling of personal data is rising.

2.3 European Union: GDPR and the AI Act

The General Data Protection Regulation (Regulation (EU) 2016/679) applies to anyone processing personal data of people in the EU in connection with offering goods or services to them or monitoring their behaviour — which can include a Canadian analyst profiling EU residents. Core principles: lawfulness, fairness and transparency; purpose limitation; data minimisation; accuracy; storage limitation; integrity and confidentiality; accountability (Art. 5). Any lawful basis, typically "legitimate interests" for private investigators, requires a documented balancing test.

The EU AI Act adds a second layer. Article 50 transparency duties, applicable since 2 August 2026, require disclosure in certain AI interactions and labelling of AI-generated content. If you publish AI-generated imagery or text to EU audiences, label it.

2.4 United States: CCPA/CPRA and computer-misuse law

California's Consumer Privacy Act as amended by the California Privacy Rights Act gives California residents rights over personal information held by covered businesses, with a partial exclusion for "publicly available" information that is narrower than many assume. Separately, the federal Computer Fraud and Abuse Act and state equivalents criminalise unauthorised access to computers. The rule of thumb for OSINT: if you would have to defeat a technical barrier — a login, a rate limit designed to block you, an IP ban — stop.

2.5 A cross-border compliance matrix

Question	Canada (PIPEDA)	EU (GDPR)	California (CCPA/CPRA)
Does "public" data escape the law?	Only narrow, regulation-defined categories	No — still personal data	Partially; definition is limited
Need a documented purpose?	Yes	Yes, plus lawful basis	Yes, notice at collection for covered businesses
Minimise data?	Yes (limiting collection)	Yes (Art. 5(1)(c))	Yes (reasonably necessary and proportionate)
Individual rights to access/delete?	Access, correction	Extensive	Access, deletion, correction, opt-out

2.6 Operational legal defence

Non-state-aligned provider status

Independent investigators should be able to demonstrate, if challenged, that they are not acting as agents of any government. Keep clear engagement letters, record who commissioned each piece of work and for what purpose, and decline tasks that would amount to surveillance on behalf of a state or its proxies. Check whether your province requires a private-investigator licence for the services you offer; in many Canadian provinces, investigating individuals for a fee requires one.

Evidentiary chain of custody

Courts and tribunals scrutinise digital evidence for authenticity and integrity. The Berkeley Protocol recommends that every captured item be accompanied by information showing where, when

and how it was obtained and that it has not been altered since. A minimal chain-of-custody record contains:

1. A unique item ID
2. The exact source URL or registry reference
3. UTC date and time of capture
4. The capture method and tool (with version)
5. A cryptographic hash (SHA-256) of the captured file
6. The analyst's name and every subsequent transfer or access

Store originals read-only. Work on copies. Never "clean up" a capture.

PRE-COLLECTION LEGAL GATE (COMPLETE BEFORE STEP 4)

- Purpose written down and approved by the client or editor
- Applicable jurisdictions identified (yours, the subject's, the data subjects')
- Lawful basis / exception noted for any personal information
- Licensing requirements checked
- Retention and deletion date set

SECTION II


The 10-Step Investigative Methodology

The scientific method, rebuilt for open-source investigation.

Steps 1-3: Problem Formulation & Hypothesis Design

Most failed investigations fail before the first search: the question was vague, or the answer was decided in advance.

Overview of the method



The method is iterative. A rejected hypothesis is not a failure; it is information, and it sends you back to Step 3 with a better question.

Step 1 — Anomaly Observation

Purpose: notice something that does not fit and that matters.

Investigations begin with a deviation from the expected: a company that should be filing disclosures but is not; a sudden cluster of newly registered domains imitating a brand; a supplier whose registered address is a residential mailbox; a press release that claims certifications no registry confirms. Anomalies fall into three families:

- **Systemic outliers** — numbers, dates or patterns that deviate from a baseline (a two-week-old company winning a large tender).

- **Disinformation signals** — coordinated, inauthentic or synthetic content (dozens of near-identical reviews posted the same hour).
- **Corporate opacity** — deliberate obscuring of ownership, location or funding (nominee directors, privacy-proxied domains, offshore parents).

Output: an *anomaly statement* — two or three sentences describing what was observed, where, when, and why it deviates from expectation. Record the observation itself as your first evidence item, captured and hashed.

Failure mode: treating a hunch as an observation. "Company X seems shady" is not an anomaly. "Company X's website claims ISO 14001 certification but the certifying body's public register lists no such certificate as of 2 October 2026" is.

Step 2 — Intelligence Requirement Definition

Purpose: convert the anomaly into precise questions that, once answered, satisfy the decision-maker.

Intelligence professionals call these Priority Intelligence Requirements (PIRs). A good PIR is specific, answerable from lawful open sources, and tied to a decision. Break each PIR into Essential Elements of Information (EIs) — the discrete facts needed to answer it.

Weak requirement	Strong PIR	EIs
Find out about Company X	Who are the ultimate beneficial owners of Company X, and do they control any competitor bidding on the same contract?	Registered directors; shareholders; parent entities; registered agents; overlap with competitor's officers

Is this campaign fake?	Were the 40 accounts promoting Product Y created and operated in a coordinated way during August 2026?	Account creation dates; posting times; shared text; shared images; linked domains
------------------------	--	---

Scope and proportionality. Write down what you will *not* collect. If the PIR concerns corporate control, you do not need officers' home addresses, family members or personal social-media activity. This written exclusion list is also your best evidence of proportionality if your work is ever challenged.

Output: a requirements sheet listing PIRs, EEIs, the decision they support, exclusions, deadline and retention date.

Step 3 — Hypothesis Formulation

Purpose: commit, in writing and before collection, to claims that the evidence could prove wrong.

The single most effective defence against confirmation bias is to state your hypotheses before you search. Each hypothesis must be **falsifiable**: it must be possible to describe evidence that would refute it. Always formulate at least one competing hypothesis, and a null hypothesis.

WORKED EXAMPLE

H₁: Company X is controlled by Company Y.

H₂: Company X is controlled by an unrelated investor group.

H₀: Company X is independently owned by its listed founders.

Indicator variables: V₁ shared officers or registered agent; V₂ shared infrastructure (hosting, certificates, analytics IDs); V₃ shared addresses; V₄ capital flows documented in filings.

Refutation criteria for H₁: no shared officers, agents, addresses or infrastructure across all available registries and records; documented independent funding.

Indicators versus proof. Variables are indicators. Shared legal counsel, for instance, is weak evidence — large firms represent thousands of unrelated clients. Assign each indicator a diagnostic weight before collection: how much more likely is this observation under H_1 than under H_2 ? An indicator equally likely under all hypotheses has no diagnostic value, however suggestive it feels.

Output: a hypothesis register: each hypothesis, its indicators, the diagnostic weight of each, and its refutation criteria, timestamped before Step 4 begins.

Steps 4-6: Collection Protocols & Analytical Rigor

Plan quietly, capture faithfully, analyse with structure.

Step 4 — Operational Design

Purpose: blueprint the collection so it is lawful, minimal, low-footprint and repeatable.

For each EEI, identify the source, the access method and the capture method. Prefer **passive** sources — those that do not touch the subject's own systems — over active ones.

EEI	Source	Access	Capture
Directors	Official corporate registry; OpenCorporates	Public web / API	PDF + JSON, hashed
Domain history	ICANN RDAP lookup; certificate transparency logs; archive snapshots	Public lookups	Screenshot + raw response
Patents	WIPO PATENTSCOPE; national patent offices	Public search	Record export
Sanctions exposure	OpenSanctions	Public search / API	Raw JSON

Zero-footprint principles

- **Use third-party records first.** Certificate transparency logs, archives and registries describe a subject's infrastructure without your ever visiting it.
- **When you must visit, visit as an ordinary reader.** Load public pages normally. Do not scan ports, enumerate hidden directories, brute-force URLs or test for vulnerabilities — that is security testing, which requires the owner's written authorisation.

- **Respect rate limits and robots directives** for any automated retrieval, and prefer official APIs and bulk-data offerings.
- **Pre-register the plan.** Save the collection plan with a timestamp. It becomes part of the replication package (Step 9).

Step 5 — Verbatim Fact Collection

Purpose: preserve sources exactly as found, with provenance, so later analysis rests on something fixed.

Collection is not note-taking. It is evidence capture. For each item:

1. Capture the full artefact (page as PDF/WARC, API response as raw JSON, document as original file).
2. Record URL, UTC timestamp, tool and version.
3. Compute and record a SHA-256 hash immediately.
4. Request an independent archive snapshot (e.g., the Internet Archive's Save Page Now) for third-party corroboration.
5. Log the item in the evidence register with a unique ID.
6. Write any observations in a separate analyst-notes field — never in the artefact itself.

```
# Minimal hashing and logging on Linux/macOS
sha256sum registry_extract_2026-10-05.pdf
# Append to the evidence register (CSV):
# id,utc_time,source_url,method,tool,sha256,analyst,notes
```

Verbatim means verbatim. Quote, do not paraphrase, in the evidence log. Translation, OCR and AI summarisation are transformations; store the original alongside any transformed version and label which is which.

Archived states. Websites change and filings get amended. Archived versions often reveal what an entity said before it learned someone was looking. Capture both the current and historical states, noting the archive's own capture timestamp.

Step 6 — Systematic Analysis

Purpose: turn a pile of facts into tested judgements, using techniques that expose bias.

Structured Analytic Techniques (SATs) externalise reasoning so it can be checked. The most useful for OSINT:

Analysis of Competing Hypotheses (ACH)

Build a matrix: hypotheses across the top, evidence items down the side. For each cell, mark whether the evidence is Consistent (C), Inconsistent (I) or Not applicable (N/A). Then focus on *inconsistencies* — the hypothesis with the fewest serious inconsistencies survives, not the one with the most consistencies.

Evidence	H ₁ (controlled by Y)	H ₂ (other investor)	H ₀ (independent)
E1: Same registered agent as Y	C	C	C
E2: Director also director of Y's subsidiary	C	I	I
E3: TLS certificate lists both companies' domains	C	I	I
E4: Funding round led by unrelated fund	I (weak)	C	I

Note that E1 is consistent with everything and therefore tells you nothing — exactly the kind of evidence that feels persuasive and is not.

Other core techniques

- **Timeline alignment** — order every dated fact. Sequences reveal causality claims that cannot be true (an officer appointed after the decision they supposedly made).

- **Entity resolution** — decide, with explicit criteria, whether two records refer to the same entity. Names are not identifiers; combine name with registry number, address, date of birth only where lawfully public and necessary, or other stable attributes.
- **Link analysis** — graph entities and relationships, labelling each edge with its evidence ID and strength.
- **Key assumptions check** — list every assumption the analysis depends on and ask what happens if each is false.
- **Devil's advocacy / red team** — have someone (or a separately prompted AI, see Chapter 8) argue the strongest case against your leading hypothesis.

Source evaluation

Grade each source on reliability and each item on credibility, separately. A reliable source (an official registry) can still contain a wrong entry; an unreliable source can occasionally be right. The common Admiralty-style scale (A-F for source reliability, 1-6 for information credibility) is a good default.

Steps 7-10: Decision Loop, Verification & Dissemination

The discipline of accepting what the evidence says, even when it disappoints you.

Step 7 — The Decision Loop

Purpose: test each hypothesis against its pre-registered refutation criteria.

Return to the hypothesis register from Step 3. For each hypothesis, ask: has any refutation criterion been met? If so, the hypothesis is **rejected** — not "weakened," not "complicated." Record the rejection and the evidence that caused it. If the evidence is mixed, the hypothesis may be **revised**, which means writing a new hypothesis and returning to Step 4 for targeted collection. If no refutation criterion is met and diagnostic evidence supports it, the hypothesis is **provisionally supported**.

THE SUNK-COST TRAP

The more time spent on an investigation, the stronger the pull to find what you expected. Pre-registered refutation criteria exist precisely to protect you from yourself at this moment. Moving the goalposts after seeing the data invalidates the method.

Step 8 — Verified Conclusion Building

Purpose: state only what the evidence supports, with calibrated confidence.

Each conclusion must cite its evidence IDs and carry a confidence expression. Use a fixed vocabulary so readers can compare statements:

Term	Approximate probability
Almost certain	95-99%
Highly likely	80-95%
Likely	55-80%
Roughly even chance	45-55%
Unlikely	20-45%
Highly unlikely	5-20%

Separate the conclusion into three layers: **what is established** (facts with sources), **what is assessed** (inferences with confidence), and **what remains unknown** (gaps). Never let an inference migrate into the "established" layer through repetition.

Bad: "Company X is a front for Company Y."

Good: "As of 1 October 2026, Company X and Company Y shared a director (E2) and appeared on the same TLS certificate (E3). We assess it is *likely* that the two are under common control. We found no documentation of share ownership; this is the principal gap."

Step 9 — Replication

Purpose: make it possible for an independent analyst to reach the same findings.

The replication package contains:

- The pre-registered requirements sheet, hypothesis register and collection plan
- Every search query, API endpoint and parameter used, with timestamps
- The evidence register, with hashes, and the captured artefacts (or archive links where redistribution is restricted)
- Analytical working: ACH matrix, timeline, link chart

- A log of AI usage: model, date, prompts and how each output was verified
- Known limitations: sources that were unavailable, records that may have changed since

Ask a colleague to attempt replication from the package alone before publication. Every place they get stuck is a gap in your documentation. Redact personal information of bystanders from any package shared outside your team.

Step 10 — Technical Reporting

Purpose: deliver findings that a reader can act on, verify and trust.

Standard dossier structure

1. **Bottom line up front** — two to four sentences answering the PIR, with confidence.
2. **AI disclosure & verification state** — which systems were used, for what, and how output was verified.
3. **Key judgements** — numbered, each with confidence and evidence IDs.
4. **Background and scope** — the anomaly, the requirement, what was excluded.
5. **Findings** — organised by EEI.
6. **Alternative hypotheses considered** — and why they were rejected.
7. **Gaps and limitations.**
8. **Methodology summary** — pointing to the replication package.
9. **Source list** — full citations, archive links, hashes.

Before publication

- **Right of reply.** If you will publish allegations about an organisation, give it a fair opportunity to respond before publication and report its response.

- **Defamation review.** Statements of fact about identifiable entities must be true and provable; statements of opinion must be clearly framed as such and based on disclosed facts. Have a lawyer review any publication alleging wrongdoing.
- **Harm review.** Remove personal details of bystanders and junior employees unless strictly necessary.
- **Proper channels.** Where findings suggest illegality, consider reporting to the relevant regulator, law enforcement or a journalist rather than publishing directly.

An "indictment" in this manual's vocabulary means a well-evidenced public-interest case presented to the proper authority or audience — never a private act of punishment.

SECTION III

Advanced Tactical OSINT

Techniques and toolkit, applied within the limits set in Section I.

GEOINT & Digital Footprint Reconstruction

Where things are, and how an organisation's online infrastructure fits together — learned from records, not from probing.

6.1 Geospatial intelligence from open sources

Satellite and aerial imagery

Freely available imagery (Sentinel-2 via the Copernicus programme, Landsat via the USGS, and the historical imagery layers of mainstream map platforms) supports questions such as: Was a facility actually built where a company claims? Has a site expanded, been cleared or been abandoned? Does a "manufacturing plant" listed in investor materials show any industrial activity?

- **Change detection.** Compare imagery of the same coordinates across dates. Record the capture date of each image (not the date you viewed it) and the platform's attribution.
- **Resolution limits.** Free medium-resolution imagery shows fields, buildings and roads, not individuals. Do not over-interpret pixels; note the ground resolution in your report.
- **Licensing.** Many platforms restrict reproduction. Cite and link rather than republish where terms require.

Geolocation and chronolocation of published media

When an organisation publishes a photo or video making a claim ("our new facility," "our cleanup operation"), you can verify the claim by matching visible features — terrain, signage, architecture, road layouts — against maps and imagery (geolocation), and by estimating time from shadows, weather records and dated events (chronolocation). This is the verification work the Berkeley Protocol describes, and it is the most powerful defence against synthetic or recycled media.

BOUNDARY: PLACES, NOT PEOPLE

Geolocation techniques are used here to verify claims made by organisations and to locate facilities, not to find where a private individual lives, works or travels. Applying these techniques to locate a person is out of scope for this manual and can constitute stalking or harassment.

Aggregate demographic context

Census and statistical-agency data (Statistics Canada, Eurostat, the US Census Bureau) supply area-level context: population, industry mix, income distributions. Use this at the aggregate level — for instance, to assess whether a claimed local workforce of 500 is plausible in a municipality of 2,000 people. Do not attempt to re-identify individuals from small-area statistics.

6.2 Reconstructing corporate digital infrastructure

An organisation's websites, domains and certificates form a footprint that often reveals relationships it does not advertise. The techniques below rely on public, third-party records.

Domain registration and history

ICANN's Registration Data Lookup service (lookup.icann.org) queries current registration data via the RDAP protocol. Since GDPR, registrant names are usually redacted, but the registrar, creation date, name servers and status codes remain informative. Historical registration data (from commercial services) can show who registered a domain before privacy redaction or proxy services were applied.

DNS records

A domain's public DNS records — name servers, mail exchangers, SPF and verification TXT records — are published so that the internet can function. Shared mail providers, matching verification

tokens, or identical SPF includes across two domains can indicate common administration. Passive DNS databases record historical resolutions without querying the target.

Certificate transparency

Publicly trusted TLS certificates are logged in append-only certificate transparency logs, searchable through services such as crt.sh. A certificate's Subject Alternative Name (SAN) field lists every domain it covers. When one certificate covers both company-x.example and parent-y.example, that is a strong indicator of shared technical control — though not, by itself, of ownership (shared hosting providers and agencies can produce the same pattern).

Subdomains

Certificate logs and passive DNS frequently reveal subdomains (staging., partners., investor.) that hint at business relationships or unannounced projects. Discover them from records; do not brute-force subdomain names against the target's DNS servers.

Shared identifiers

Public web pages often embed analytics and advertising identifiers, favicon files, or distinctive code. Identical identifiers across sites suggest a common operator. Several commercial and community tools index these; the method only requires viewing public page source as any browser would.

Internet-exposure search engines

Services such as Shodan index banners and metadata from internet-connected devices. They are useful for *understanding an organisation's public attack surface for defensive reporting*. Viewing an index entry is passive; connecting to, logging into, or testing a listed device without authorisation is not, and is illegal in most jurisdictions.

6.3 Exposed cloud storage: a hard boundary

Misconfigured cloud storage buckets sometimes expose files to the open internet. The fact that a file is reachable does not make downloading it lawful or ethical, particularly where it contains personal information. Canadian, EU and US regulators treat exposed personal data as a breach, and accessing it can expose the investigator to liability.

IF YOU ENCOUNTER EXPOSED DATA

- Do not download, copy or analyse the contents beyond what is needed to recognise the exposure.
- Record the minimum metadata (URL, time) and stop.
- Notify the organisation through its security contact or vulnerability-disclosure programme, or a national CERT (in Canada, the Canadian Centre for Cyber Security).
- Do not publish the location of the exposure until it is fixed.

6.4 Toolkit summary

Need	Open resource	Passive?
Domain registration	ICANN Lookup (RDAP)	Yes
Certificates / subdomains	Certificate transparency search (e.g., crt.sh)	Yes
Historical pages	Internet Archive Wayback Machine	Yes
File / URL reputation	VirusTotal	Yes (submitting files makes them visible to others — never upload client documents)
Exposure index	Shodan	Viewing yes; interacting no
Imagery	Copernicus / Sentinel Hub; USGS EarthExplorer	Yes

Tool directory	OSINT Framework (osintframework.com)	—
----------------	--	---

Corporate Forensics & Financial Lineage

Following control and money through registries, filings and public ledgers.

7.1 Why ownership is hidden — and why that is a signal

Layered ownership structures have legitimate uses: tax planning, liability isolation, joint ventures. They are also the standard tool for concealing conflicts of interest, sanctions exposure and corruption. The investigator's task is not to assume wrongdoing from complexity, but to establish who ultimately controls an entity and whether that control matters to the question.

7.2 Mapping parent-subsidiary structures

Primary registries

Start with the official registry of the jurisdiction where the entity is incorporated: Corporations Canada for federal corporations, Service New Brunswick's corporate registry for New Brunswick, Companies House in the UK, the relevant secretary of state in each US state, and so on. Capture the official extract, not a third-party summary.

Aggregators

OpenCorporates aggregates company records from many jurisdictions, making cross-border officer and address searches practical. Treat it as an index that points you to the official record, and verify against the primary registry before relying on an entry.

Beneficial-ownership registers

Many jurisdictions now require companies to record individuals with significant control. Canada's federal register of individuals with significant control became publicly searchable in January 2024 for federally incorporated corporations; the UK's People with Significant Control register is public. Availability elsewhere varies and changes; check the current rules for each jurisdiction.

Securities and regulatory filings

Public companies file disclosures (SEDAR+ in Canada, EDGAR in the US) that list subsidiaries, related-party transactions and material contracts. Exhibits and footnotes are where structures are revealed.

Sanctions and risk lists

OpenSanctions consolidates sanctions lists, politically exposed persons data and related entities from many official sources. A match is a lead, not a finding: confirm against the original listing and resolve identity carefully before reporting.

7.3 Building the structure chart

1. List every entity with its jurisdiction and registry number (the number, not the name, is the identifier).
2. For each, record officers, registered agent, registered address and shareholders as filed, with dates.
3. Draw ownership edges only where documented; draw "shared officer," "shared address" and "shared agent" edges in a different style, because these indicate association, not ownership.
4. Flag nominee indicators: officers holding hundreds of directorships, mass-registration addresses, corporate service providers.
5. Mark the point where documentation runs out. That boundary is a finding in itself.

7.4 Patents and intellectual property

Patent assignments and reassignments are public. WIPO's PATENTSCOPE searches international (PCT) applications and many national collections; national offices publish assignment records. Patterns worth noting include a newly formed entity acquiring a cluster of patents in one technical field, assignments routed through intermediaries, and patents acquired and then allowed to lapse — which can be consistent with defensive acquisition intended to keep a technology off the market. As always, that last interpretation is a hypothesis to test, not a conclusion.

7.5 Public blockchain analysis

Public blockchains such as Bitcoin and Ethereum record every transaction in a ledger anyone can read through a block explorer. This makes them open sources — but pseudonymous ones.

- **Starting points must be published.** Begin from addresses an organisation has publicly declared — donation addresses, treasury wallets disclosed in filings, smart-contract addresses on a project's website.
- **Flows, not identities.** You can trace value moving between addresses. Attributing an address to a real-world entity requires independent evidence (the entity's own disclosure, a regulator's statement, an exchange's public attestation).
- **Smart contracts.** Verified contract source code, deployer addresses and administrative keys are public on chains that support them. Who can upgrade or pause a contract is a question of control, and therefore of governance.
- **Decentralised organisations.** DAO governance proposals, token-holder votes and treasury movements are often fully on-chain and are ideal material for accountability research.
- **Heuristic risk.** Clustering heuristics used to group addresses are probabilistic and can be wrong. State them as assumptions and avoid naming individuals on the strength of a heuristic alone.

Decentralized Identifiers (DIDs), standardised by the W3C as a Recommendation in July 2022, are increasingly used to anchor organisational identity claims. When an entity publishes a DID, its DID document can be resolved to see which keys and service endpoints it controls — another public record of control.

AI-Powered Intelligence Integration

Using language models as tireless junior analysts — and never as witnesses.

8.1 Where AI genuinely helps

Task	AI role	Human check
Entity extraction from long filings	Pull names, registry numbers, dates, amounts into a table	Spot-check every row against the source page
Translation	First-pass translation of foreign filings or media	Verify key passages with a fluent reader; keep the original
Clustering narratives	Group thousands of posts by theme or near-duplicate text	Sample each cluster manually
Drafting queries	Suggest search operators and alternative spellings	Review for scope creep
Red-teaming	Argue against your leading hypothesis	Evaluate each argument on the evidence
Report drafting	Structure and prose	Every factual sentence traced to an evidence ID

8.2 Building a custom extraction agent

A reliable AI workflow constrains the model to the documents you supply and demands evidence for every output. A practical design:

1. **Ingest** captured documents (with their evidence IDs) into a local store.
2. **Retrieve** only the passages relevant to each question.

3. **Instruct** the model to answer only from the supplied passages, to quote the exact supporting text, and to cite the evidence ID and page for each claim — and to reply "not found in sources" otherwise.
4. **Validate automatically**: a script checks that each quoted string actually appears verbatim in the cited document. Any claim whose quote fails the check is discarded.
5. **Log** the model name and version, date, prompt and output in the replication package.

```
SYSTEM: You extract facts ONLY from the documents provided.
For each fact output JSON:
  {"fact": ..., "quote": "<exact text>", "evidence_id": ...,
  "page": ...}
If the documents do not contain the answer, output {"fact": null}.
Never use outside knowledge. Never infer ownership from similarity.
```

8.3 Counter-hallucination protocol

Language models generate plausible text; plausibility is not truth. The protocol below treats every AI output as unverified until it passes a mechanical check.

- **Verbatim grounding.** No claim is accepted without an exact quote that can be string-matched to a captured source.
- **Citation verification.** Every reference an AI suggests is looked up independently. (For this edition, six of twenty references in the AI-generated draft list failed this check — see the Corrections Log.)
- **Cross-model disagreement.** Running the same extraction through two different models and comparing outputs flags items for human review; agreement is not proof, but disagreement is a useful alarm.
- **No AI-only facts.** If a fact exists only in an AI's output, it does not exist.
- **Disclosure.** State in the report which systems were used and how.

8.4 Disinformation and synthetic-media analysis

AI can help detect coordinated inauthentic behaviour by surfacing near-duplicate text, synchronised posting times and recycled images across accounts. Treat automated "AI-generated" detectors with caution: their error rates are significant, and a detector score is a weak indicator, not proof. Stronger evidence of synthetic or manipulated media comes from provenance (content credentials where present, earliest known appearance via reverse-image search and archives), internal inconsistencies (lighting, shadows, text rendering), and contradiction with verified geolocation or chronolocation.

Sentiment analysis can describe the tone of a body of public content at aggregate level. It should not be used to profile or score individual private people.

8.5 Local versus cloud models

Sending client documents to a cloud AI service is a disclosure to a third party. Review the provider's data-use terms, use enterprise or no-training settings where available, and for sensitive material prefer locally run open-weight models on hardware you control (see Chapter 9). Record the choice in the replication package.

SECTION IV

Counter-Surveillance & Operational Security

Protect the analyst, the sources and the integrity of the work.

Digital Hygiene for the Sovereign Researcher

Operational security for open-source work is defensive: it keeps your identity, your clients and your evidence safe.

9.1 Why OPSEC matters even for lawful research

Organisations under scrutiny sometimes notice who is looking at them. Visiting a small company's website from your personal connection can reveal your location and employer in its server logs; viewing a professional profile can notify its owner; a careless search history can tip off a subject before a regulator is informed. OPSEC protects the investigation, your clients' confidentiality and your own safety — especially when researching powerful or hostile entities.

9.2 Threat-model first

Before choosing tools, answer four questions: What am I protecting (identity, client, evidence, sources)? From whom (the subject, its contractors, opportunistic attackers)? How likely is it they will look? What happens if they succeed? A due-diligence check on a large public company needs modest precautions; research into organised fraud networks needs much more.

9.3 Hardening the investigative environment

Compartmentalised workstations

- **Virtual machines.** Conduct research inside a dedicated VM (e.g., a Linux guest in VirtualBox or similar), separate from personal email, banking and social accounts. Snapshot a clean baseline and revert after each engagement.

- **Dedicated browser profiles** per investigation, with extensions limited to those you need and understand.
- **Separate research accounts** — email addresses and, where a platform's terms permit, accounts used only for passive viewing. These protect your identity; they must never be used to deceive people into giving you non-public information, to join private groups under false pretences, or to interact with subjects.

Network layer

A reputable VPN or Tor masks your IP address from sites you visit. Understand the trade-offs: VPN providers can see your traffic metadata; Tor is slower and blocked by some sites; neither protects you if you log into a personal account in the same session. Commercial VPN use is legal in Canada, but check the terms of the sites you use and the laws of any country you are in.

Burner infrastructure, properly understood

"Burner" infrastructure means disposable, isolated environments — a VM, a research phone number from a lawful provider, a separate cloud workspace — that can be discarded after an engagement so that investigations do not cross-contaminate. It does not mean anonymous infrastructure for evading accountability. Keep internal records tying each burner to the engagement it served; your chain of custody depends on it.

Local LLM deployment

Running open-weight language models locally keeps client material off third-party servers. Practical points: choose a model whose licence permits your use; run it on an offline or firewalled host when processing sensitive documents; log the model name, version and quantisation in your replication package; and remember that smaller local models hallucinate at least as readily as large cloud ones — the counter-hallucination protocol of Chapter 8 applies unchanged.

Evidence security

- Encrypt evidence stores at rest (full-disk encryption plus encrypted archives for transfer).
- Keep at least two backups, one offline.
- Use a password manager and hardware-key two-factor authentication on every account that touches the investigation.
- Delete personal information on the retention date set in Chapter 2's legal gate.

9.4 Reducing your signal while you collect

The outline for this chapter called this "signal-to-noise manipulation." In a lawful practice, it reduces to a set of habits that keep your presence unremarkable during *passive* collection:

- **Prefer third-party records** (registries, archives, certificate logs) to the subject's own servers. Records do not log your visit to the subject.
- **View archived copies** of a subject's site rather than the live site where the archive suffices.
- **Disable "who viewed your profile" disclosure** on professional networks, or use a research account that does not identify you.
- **Pace automated retrieval** politely; aggressive request patterns stand out and may breach terms.
- **Avoid distinctive search terms** on services the subject might monitor, such as alerts on their own name.

WHAT THIS MANUAL DOES NOT TEACH

This manual does not cover covertly watching, following or tracking people, installing monitoring software or tracking devices, intercepting communications, or using pretexts to extract information. These activities are illegal or tightly regulated in Canada and most other jurisdictions, and they are incompatible with a post-predatory practice. "Active target surveillance" in the sense of watching a person is out of scope; the investigator's attention stays on documents, records and organisational conduct.

9.5 Counter-surveillance for the analyst

Investigators of powerful entities are sometimes targeted in return. Reduce your own exposure: remove personal addresses from data-broker sites where the law allows opt-out, keep business and personal identities separate, lock down social-media privacy settings, and agree on a response plan (lawyer, publisher, police contact) before publishing findings likely to provoke a reaction. If you believe you are being harassed or physically surveilled, document it and contact police.

The Future of Sovereign Reconnaissance (2030 Horizon)

Law, accountability and civil rights as autonomous systems take over collection.

10.1 The legal trajectory

Three trends are clear from developments up to October 2026, and are likely to continue to 2030:

- **Privacy law is converging on the GDPR model.** Canada's Bill C-36, if enacted, would move federal private-sector law toward stronger rights, a dedicated regulator and meaningful penalties, while Quebec's Law 25 already operates at that level. Investigators should expect documentation of purpose, minimisation and retention to become a routine compliance expectation rather than best practice.
- **AI transparency is becoming mandatory.** The EU AI Act's Article 50 obligations already apply; high-risk obligations follow in December 2027 and August 2028. Disclosure of AI involvement in published intelligence products — already a principle of this manual — will increasingly be a legal requirement.
- **Beneficial-ownership transparency is spreading,** though unevenly and with periodic reversals as courts and legislators balance transparency against privacy.

10.2 Autonomous collection and its dangers

Agentic AI systems can now plan and execute multi-step research with little human input. This creates three risks the independent investigator must manage: **scope creep** (an agent told to "find everything about Company X" may wander into employees' personal lives), **unauthorised access** (an agent may attempt actions

— form submissions, logins, scans — that cross legal lines), and **fabrication at scale** (an agent that summarises its own summaries compounds errors). The answer is the same discipline this book teaches, applied to machines: written requirements, explicit exclusions, tool permissions limited to passive retrieval, and verbatim grounding before anything is accepted.

10.3 Public accountability and civil rights

The same capabilities that let citizens scrutinise corporations let states and corporations scrutinise citizens. A credible independent OSINT community strengthens public accountability only if it holds itself to higher standards than the actors it investigates. That means refusing work that targets activists, journalists, vulnerable people or private individuals; supporting legal limits on mass surveillance by any actor; and publishing methods openly so others can check the work.

10.4 Open intelligence and the Organic Revolution

The author's broader programme — described elsewhere as the Organic Revolution — envisions economies that reward sustainability and transparency over extraction. Open intelligence contributes in concrete, lawful ways:

- **Verifying sustainability claims** against certifying bodies' public registers, exposing greenwashing.
- **Mapping market concentration** through ownership and patent records, supplying evidence to competition regulators (in Canada, the Competition Bureau).
- **Tracing supply chains** from customs data, corporate disclosures and imagery to test claims about origin and labour practices.
- **Documenting lobbying and influence** using public lobbyist registries and political-finance disclosures.

"Dismantling predatory monopolies," in this frame, means producing rigorous public-interest evidence and delivering it to regulators, courts, legislators and the public — the institutions with legitimate authority to act on it.

10.5 A closing word

The investigator of 2030 will have more powerful tools than any intelligence agency had in 2000. What will distinguish trustworthy work is not capability but restraint: the willingness to collect less than is possible, to conclude only what is proven, and to show every step. That is the whole of legal OSINT espionage.

The 10-Step Framework End to End

A hypothetical investigation into an undisclosed acquirer of green-technology assets. All entities are fictional.

1. Observe anomaly

An unfamiliar entity, "Verdant Holdings Ltd." (fictional), has acquired several patents and domains associated with bio-based materials over four months. Industry observers can find no announcements, and the entity's website is a single holding page.

2. Define the intelligence requirement

PIR: Who ultimately controls Verdant Holdings, and is it affiliated with any established petrochemical company? **EElS:** incorporation details; officers; registered agent and address; shareholders or significant-control filings; patent assignment chain; domain and certificate relationships; funding sources disclosed in filings. **Excluded:** officers' personal addresses, family, personal social media.

3. Hypothesise

H₁: Verdant is a wholly owned vehicle of a legacy petrochemical company, acquiring patents defensively. **H₂:** Verdant is an independent venture-backed startup. **H₀:** Verdant is a patent-assertion entity with no operating affiliation. Indicators: shared officers (V₁), shared infrastructure (V₂), documented capital flows (V₃), patent maintenance behaviour (V₄). Refutation of H₁: no shared officers, infrastructure or capital links with any petrochemical company in available records.

4. Design the collection

Query the official registry and OpenCorporates; check significant-control filings; pull patent assignment records via PATENTSCOPE and the national office; search certificate transparency logs and

RDAP for the acquired domains; review archived versions of the holding page; check OpenSanctions. No contact with Verdant's servers beyond loading its public page once inside a research VM.

5. Collect verbatim facts

Twelve evidence items captured and hashed, including: the registry extract showing two directors and a corporate-services registered agent; patent assignment records showing transfers from three university spin-outs; a certificate whose SAN field lists Verdant's domain alongside a domain registered to a subsidiary of a large energy group; an archived 2025 version of the holding page carrying the same analytics identifier as that subsidiary's site.

6. Analyse

The ACH matrix shows: the shared registered agent is consistent with all hypotheses (non-diagnostic); the shared certificate and analytics identifier are inconsistent with H_2 and H_0 ; one director also sits on the subsidiary's board (inconsistent with H_2 and H_0). The timeline shows Verdant's incorporation two weeks after the subsidiary announced a strategic review of bio-materials.

7. Decision loop

H_1 's refutation criteria are not met. H_2 and H_0 each face multiple diagnostic inconsistencies and are rejected. However, no document establishes share ownership, so H_1 is *provisionally supported* as "common control," not "wholly owned." The "defensive acquisition" element of H_1 remains untested — patents are still being maintained — and is set aside as an open question.

8. Verified conclusion

"As of 1 October 2026, Verdant Holdings shared a director (E4), TLS certificate (E7) and analytics identifier (E9) with a subsidiary of an energy group. We assess with *high likelihood* that the two are under common control. We found no filing establishing ownership, and no evidence on whether the acquired patents will be used or shelved."

9. Replication

The package includes all queries with timestamps, the twelve artefacts with hashes and archive links, the ACH matrix, the timeline, and a log of AI use (entity extraction from the patent records, verified line by line).

10. Technical report

A six-page dossier with a bottom-line-up-front statement, AI disclosure, key judgements with confidence levels, alternative hypotheses, gaps, and sources. Before any publication, Verdant and the energy group are offered a right of reply, and the report is reviewed by counsel.

Field Templates

A.1 Requirements sheet

Case ID:

Date opened (UTC):

Client / editor:

Purpose (one sentence):

Anomaly statement:

PIR-1:

EETs:

PIR-2:

EETs:

Decision supported:

EXCLUDED from collection:

Jurisdictions:

Lawful basis / exception:

Licence required?

Retention / deletion date:

A.2 Hypothesis register

ID	Hypothesis	Indicators (weight)	Refutation criteria
Registered (UTC)			
H1			
H2			
H0			

A.3 Evidence register

ID	UTC time	Source URL / registry ref	Method	Tool+ver	SHA-256
	Archive link	Analyst	Reliability (A-F)	Credibility (1-6)	

A.4 AI usage log

Date	Model + version	Task	Input evidence IDs	Prompt file
Output file	Verification method	Items rejected		

A.5 Report AI disclosure block

AI Disclosure & Verification State

- Author / analyst:

- AI systems used (name, version, date, purpose):

- Verification: [what was checked, how, and what failed]

- Status: [Verified / Partially verified / Unverified draft]

Glossary

Term	Definition
ACH	Analysis of Competing Hypotheses: a matrix method that tests evidence against multiple hypotheses, emphasising inconsistencies.
Certificate transparency	Public, append-only logs of issued TLS certificates.
Chain of custody	Documented record of an evidence item's capture, storage, transfer and access.
DID	Decentralized Identifier, a W3C-standardised identifier controlled by its subject.
EEI	Essential Element of Information: a specific fact needed to answer a PIR.
Ethical Pathfinder	The author's term for values-driven entrepreneurs, researchers and activists.
GEOINT	Geospatial intelligence derived from imagery and location data.
Non-state-aligned	Operating independently of any government's intelligence apparatus.
PIR	Priority Intelligence Requirement: a precise question tied to a decision.
Post-predatory	Intelligence practice that refuses to profit from exploiting people's data or vulnerabilities.
RDAP	Registration Data Access Protocol, successor to WHOIS for domain registration data.
SAN	Subject Alternative Name: the field of a TLS certificate listing the domains it covers.
SAT	Structured Analytic Technique.
Sovereign researcher	An analyst who retains custody of data, understanding of tools and ownership of judgements.
UBO	Ultimate Beneficial Owner: the natural person(s) who ultimately own or control an entity.

REFERENCES

References (Verified 5 October 2026)

Each entry below was checked by live web search on 5 October 2026. Entries marked † were corrected or replaced from the Gemini-stage draft; details are in the Corrections Log. Database and tool entries are cited by their 2026 on-line versions. A second, line-by-line audit of every entry, URL and statutory citation was completed the same day (see Corrections Log).

- Bazzell, M. (2023). *OSINT Techniques: Resources for Uncovering Online Information* (10th ed.). IntelTechniques.com. †
- Canada. (2000). *Personal Information Protection and Electronic Documents Act*, S.C. 2000, c. 5. Justice Laws Website. <https://laws-lois.justice.gc.ca/eng/acts/P-8.6/>
- European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). *Official Journal of the European Union*. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- European Union. (2026). Regulation (EU) 2026/1744 (Digital Omnibus on AI), amending Regulation (EU) 2024/1689 (Artificial Intelligence Act). In force 27 July 2026.
- Government of Canada. (2026). Bill C-36, *An Act to enact the Protecting Privacy and Consumer Data Act...* First reading 15 June 2026. 45th Parliament, 1st Session. LEGISinfo: <https://www.parl.ca/legisinfo/en/bill/45-1/c-36>
- Hassan, N., & Hijazi, R. (2018). *Open Source Intelligence Methods and Tools: A Practical Guide to Online Intelligence*. Apress. <https://doi.org/10.1007/978-1-4842-3213-2>
- Internet Archive. (2026). *Wayback Machine*. <https://web.archive.org/>
- ICANN. (2026). *Registration Data Lookup Tool*. <https://lookup.icann.org/>
- Landry, M.-S. S. (2026). *Marie Landry's Spy Shop* [Website]. <https://www.marielandryspyshop.com/>
- Landry, M.-S. S. (2026). ORCID record 0009-0008-5027-3337. <https://orcid.org/0009-0008-5027-3337>
- NATO. (2001, November). *NATO Open Source Intelligence Handbook* (v1.2, January 2002). North Atlantic Treaty Organization. †

- Office of the Privacy Commissioner of Canada et al. (2024, October 28). *Concluding joint statement on data scraping and the protection of privacy*. https://www.priv.gc.ca/en/opc-news/speeches/2024/js-dc_20241028/ †
- Office of the Privacy Commissioner of Canada et al. (2023, August 24). *Joint statement on data scraping and the protection of privacy*. https://www.priv.gc.ca/en/opc-news/speeches/2023/js-dc_20230824/
- OpenCorporates. (2026). *The Open Database of the Corporate World*. <https://opencorporates.com/>
- OpenSanctions. (2026). *OpenSanctions*. <https://www.opensanctions.org/>
- OSINT Framework. (2026). <https://osintframework.com/>
- Shodan. (2026). <https://www.shodan.io/>
- United Nations Office of the High Commissioner for Human Rights & Human Rights Center, UC Berkeley School of Law. (2022). *Berkeley Protocol on Digital Open Source Investigations* (HR/PUB/20/2; advance version first released December 2020). United Nations. †
- United Nations Office on Drugs and Crime. (2011). *Criminal Intelligence: Manual for Analysts*. United Nations. https://www.unodc.org/documents/organized-crime/Law-Enforcement/Criminal_Intelligence_for_Analysts.pdf †
- U.S. Department of State, Bureau of Intelligence and Research. (2024, May). *Open Source Intelligence Strategy*. <https://www.state.gov/wp-content/uploads/2024/05/INR-Open-Source-Intelligence-Strategy.pdf> †
- VirusTotal. (2026). <https://www.virustotal.com/>
- World Intellectual Property Organization. (2026). *PATENTSCOPE*. <https://patentscope.wipo.int/>
- World Wide Web Consortium. (2022, July 19). *Decentralized Identifiers (DIDs) v1.0* (W3C Recommendation). <https://www.w3.org/TR/did-core/> †

Reference Corrections Log

What the live checks on 5 October 2026 found in the Gemini-stage reference list.

Draft entry	Finding	Action
Bazzell (2023), <i>Open Source Intelligence Techniques</i> , 10th ed., IntelTechniques Press	The 10th edition (2023) exists but was retitled <i>OSINT Techniques: Resources for Uncovering Online Information</i> .	Title corrected.
International Criminal Court (2021), Berkeley Protocol	Authored by OHCHR and UC Berkeley's Human Rights Center, not the ICC; released December 2020, UN print edition 2022.	Authors and date corrected.
NATO Standardization Office (2020), "AIntP-03: HUMINT and OSINT Doctrine"	No such publication found. NATO's OSINT doctrine is AJP-2.9 (2019), which is restricted; the public reference is the 2001 OSINT Handbook.	Replaced with the NATO OSINT Handbook.
OPC (2024), "Guidance on Scraped Data...", URL .../ gd_scraped_2024/	No document by that title or URL found. The OPC's 2024 publication on this topic is the concluding joint statement of 28 October 2024.	Replaced with the actual statement; 2023 initial statement added.
UNODC (2023), "Manual on Open Source Intelligence in Criminal Investigations"	No such manual found. UNODC's published analyst manual covering open sources is the 2011 <i>Criminal Intelligence: Manual for Analysts</i> .	Replaced.

US DHS (2022), "Open Source Intelligence Strategy"	No DHS document by that title found. A public US OSINT strategy is the State Department INR strategy (May 2024).	Replaced.
W3C (2023), DIDs v1.0	Became a W3C Recommendation on 19 July 2022.	Date corrected.

Second audit (5 October 2026)

Every reference, URL and statutory section cited in the text was re-checked. Results: the GDPR, PIPEDA (including ss. 4(2)(c), 5(3) and 7(1)(d)), both OPC joint statements (2023: twelve authorities; 2024: sixteen), the W3C DID Recommendation URL, the Hassan & Hijazi DOI, Bill C-36's LEGISinfo record and all database and tool URLs were confirmed. One unconfirmed detail was removed: the draft attributed the NATO Handbook to Supreme Allied Commander Atlantic, which could not be verified, so the entry now credits NATO generally. The Bill C-36 entry gained its official LEGISinfo link.

Additions: the EU Digital Omnibus on AI and Canada's Bill C-36 were added because the manual's legal chapters rely on them. Both were confirmed against multiple legal-commentary sources dated June–September 2026.

ABOUT THE AUTHOR

Marie-Soleil Seshat Landry

Marie-Soleil Seshat Landry is the founder and CEO of Landry Industries, a Moncton, New Brunswick incubator whose ventures span open-source intelligence, materials science, AI governance and organic sustainability. Through Marie Landry's Spy Shop (marielandryspyshop.com) she has published OSINT research, frameworks and field guides for about a decade, serving a readership she calls Ethical Pathfinders. She works in English and French and publishes her research openly under ORCID 0009-0008-5027-3337.

Contact: ceo@marielandryspyshop.com