LightHouse

BY SILENT HARBOR

FIELD GUIDE / VERSION 1.0 / JULY 2026

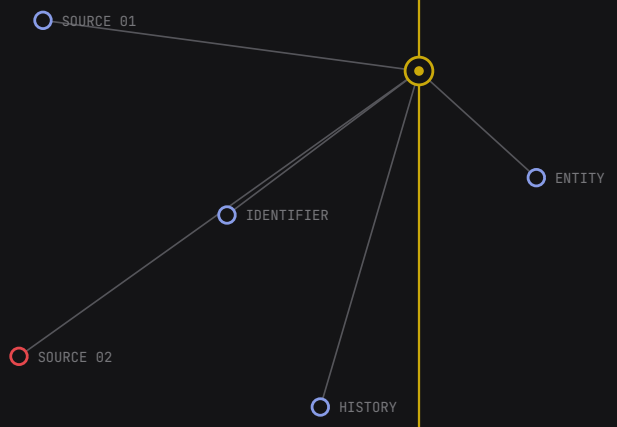

OSINT INVESTIGATIONS WITH LIGHTHOUSE

How to turn open-source information into evidence you can inspect, explain, and act on

Open-source information is abundant. Useful intelligence is not. This field guide explains how LightHouse helps teams move from an authorized question and a fragment of identity to a source-cited, human-reviewed intelligence product.

00 / LIGHTHOUSE OSINT FIELD GUIDE

EXECUTIVE BRIEF

Most organizations do not have an information shortage. They have a resolution problem.

The useful fact may be public or commercially available, yet scattered across corporate records, websites, social platforms, technical infrastructure, archives, news, sanctions data, and internal holdings. The same person may appear under several names. Several people may share the same name. A domain may have changed hands. A copied article may look like independent corroboration. A true fact may be irrelevant because it belongs to the wrong time, entity, or context.

Open-source intelligence, or OSINT, is the disciplined work of turning that environment into an answer to a defined intelligence requirement. The United States Intelligence Community describes OSINT as intelligence derived exclusively from publicly or commercially available information that addresses a specific priority, requirement, or gap. The important word is intelligence. Collection is necessary, but collection alone is not the product.

LightHouse is Silent Harbor's Identity Intelligence Engine. It helps teams plan an OSINT investigation, direct mission-bound agents across permitted sources, resolve the people and entities behind fragmented identifiers, preserve the evidence path, expose contradictions, and produce a decision record for human review.

The operating model is direct:

1. Define the decision, scope, authorities, evidence standard, and review points.
2. Begin with the fragment that matters: a person, company, email, alias, domain, wallet, transaction, photograph, or piece of infrastructure.
3. Resolve what stands behind it and develop the relevant environment around it.
4. Test connections, chronology, source independence, and alternative explanations.
5. Separate supported facts from assessments and name what remains uncertain.
6. Deliver the answer in the form the decision owner can review and use.

LightHouse does not make OSINT trustworthy by automating more searches. It makes the work more inspectable by binding collection, identity resolution, analysis, provenance, and review to one mission.

> The search is not the product. The supported conclusion is.

01 / LIGHTHOUSE OSINT FIELD GUIDE

OSINT BEGINS WITH A DECISION

OSINT is often described through tools, sources, or search techniques. Those matter, but none of them defines the mission.

MISSION BOUNDARY

DECISION	SCOPE	SOURCES	EVIDENCE	REVIEW
What must be decided?	Who, what, and when?	What is permitted?	What will count?	Who retains authority?

A useful investigation begins with a decision that has to be made. Should this transaction be held? Does this counterparty require deeper review? Is an exposed executive profile creating a practical targeting path? Are several anonymous accounts likely controlled by the same actor? Does a weak insider-risk signal warrant escalation? Which finding would materially change the answer?

That decision establishes the intelligence requirement. It also creates boundaries:

- **Purpose:** Why is the investigation authorized, and who owns the resulting decision?
- **Scope:** Which subjects, identifiers, relationships, time periods, and jurisdictions are relevant?
- **Sources:** Which public, commercial, customer, and technical sources may be used?
- **Evidence:** What must be observed directly, corroborated independently, or assessed by an expert?
- **Time:** When does the answer lose operational value?
- **Review:** Which conclusions require human judgment, legal review, or another control?
- **Output:** What must the decision owner receive: an alert, graph, dossier, assessment, export, or briefing?

This mission definition matters because the open information environment has no natural stopping point. Without a bounded question, an OSINT investigation can accumulate material without becoming more useful. LightHouse uses the mission to decide where agents can look, what they should pursue, how findings should be evaluated, and when the work is ready for review.

The first design decision is therefore not which source to query. It is what the organization needs to know, by when, and what evidence would change the call.

> LightHouse begins with the decision and the fragment, not the tool list.

02 / LIGHTHOUSE OSINT FIELD GUIDE

WHY CONVENTIONAL OSINT WORKFLOWS BREAK

Capable investigators can work around fragmented tools and manual processes. The difficulty is sustaining rigor when the number of sources, identities, cases, and deadlines grows.

WHERE RIGOR IS LOST

FRAGMENT	REPEAT	IDENTITY	CONTEXT	AUTOMATE	REPORT
Sources remain disconnected.	Copies resemble corroboration.	Weak links look resolved.	Time and purpose fall away.	Speed outruns evidence.	Reasoning is rebuilt late.

Fragmentation without resolution

Search results arrive as documents, profiles, screenshots, technical records, registry entries, alerts, and analyst notes. They rarely agree on names, identifiers, time, or ownership. The investigator must determine which observations refer to the same subject, which are merely adjacent, and which are unrelated.

Repetition mistaken for corroboration

Ten pages may repeat one original claim. If the underlying source is the same, the claim has been duplicated, not independently confirmed. A defensible OSINT process must preserve source lineage and distinguish repetition from corroboration.

Identity ambiguity

Names collide. Aliases drift. Companies share officers or addresses. Accounts are abandoned or transferred. Domains expire. A strong-looking connection can fail once chronology, ownership, or a second identifier is examined.

Context lost between tools

Collection, link analysis, screenshots, case notes, and reporting often live in separate systems. Each transfer creates an opportunity to lose the reason a source was collected, the time it was observed, the chain behind a conclusion, or the analyst's unresolved question.

Automation without an evidence standard

Faster collection can increase false confidence if the system does not distinguish direct observation from inference, current association from historical association, or one source from multiple independent sources. Speed is useful only when the result remains reviewable.

Reporting after the fact

When the evidence path is reconstructed at the end, the report becomes a memory exercise. Important contradictions may disappear, weak links may inherit more confidence than they deserve, and reviewers may be unable to reproduce the reasoning.

LightHouse addresses these failures as one operating problem. It keeps the question, subject, sources, identity graph, chronology, assessments, and decision product connected throughout the mission.

03 / LIGHTHOUSE OSINT FIELD GUIDE

THE LIGHTHOUSE OSINT OPERATING SEQUENCE

LightHouse organizes OSINT investigations as a controlled sequence. Agents pursue bounded tasks, evidence rules determine what counts, and human reviewers retain authority over consequential conclusions and actions.

CONTROLLED OPERATING SEQUENCE

One mission. One evidence path. Human review at consequential points.

1. Frame the mission

The customer and mission team define the intelligence requirement, permitted data, target decision, delivery window, collection limits, escalation conditions, and required output. Known facts and working assumptions are entered separately so the investigation does not quietly treat the starting story as true.

2. Start with the fragment

The starting point may be sparse: a name, email address, phone number, username, company, domain, wallet, transaction, image, document, or infrastructure artifact. LightHouse records where it came from, when it was observed, and why it matters to the mission.

3. Resolve the subject

Mission-bound agents develop candidate people, entities, accounts, organizations, infrastructure, and relationships across authorized sources. LightHouse compares identifiers, chronology, co-occurrence, ownership, geography, behavior, and source reliability. Plausible matches remain candidates until the evidence supports a resolution.

4. Develop the relevant environment

The subject is the anchor, not the boundary. The system follows relevant relationships into corporate structures, aliases, digital assets, exposed information, historical records, counterparties, technical infrastructure, and other mission-specific context. Collection expands only where it can answer the requirement or test an alternative explanation.

5. Test the apparent story

LightHouse returns to weak links, searches for disconfirming evidence, exposes contradictory records, separates current from historical associations, and checks whether apparent corroboration is actually independent. The investigator can redirect agents, add a hypothesis, narrow the mission, or require a stronger evidence threshold.

6. Produce the decision record

The final product distinguishes observed facts, analytical assessments, confidence, unresolved uncertainty, and recommended human review. Every consequential conclusion retains a path back to its supporting evidence. The result can be delivered as a graph, dossier, written assessment, structured export, alert, or briefing.

7. Monitor what can change

Where the mission requires continued awareness, approved entities and conditions can move into monitoring. LightHouse looks for material changes against the established baseline and returns those changes to the same evidence and review model.

This sequence makes the investigation repeatable without pretending that every mission is the same. The workflow can change by use case. The evidentiary discipline does not.

04 / LIGHTHOUSE OSINT FIELD GUIDE

IDENTITY INTELLIGENCE IS THE ENGINE

OSINT questions frequently arrive disguised as search problems. Underneath, many are identity problems.

IDENTITY-CENTERED EVIDENCE GRAPH

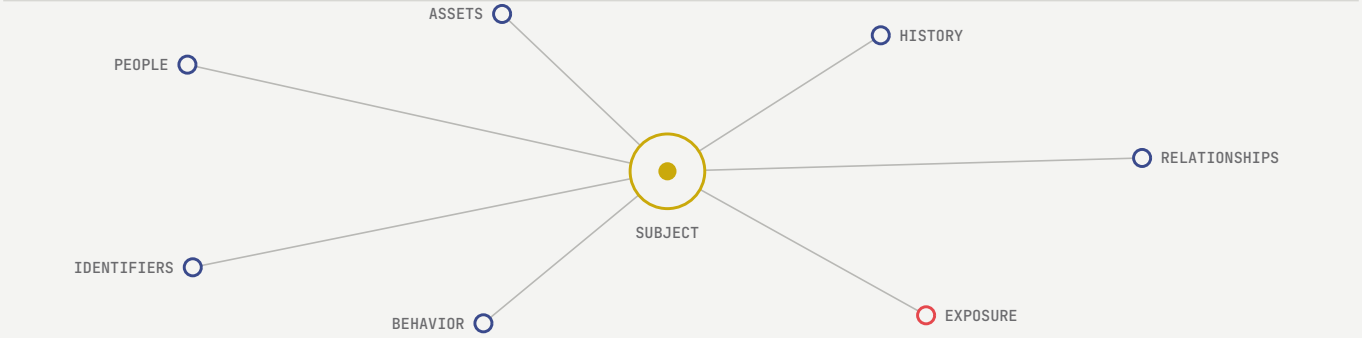

The investigator is not merely trying to find an email address or company record. The investigator is trying to determine whether several fragments belong to the same person, whether an account is controlled by the apparent user, whether an entity is connected to a hidden principal, whether technical infrastructure belongs to the same operating cluster, or whether a historical relationship remains relevant now.

LightHouse builds an identity-centered evidence graph around six classes of context:

- **People and entities:** subjects, counterparties, officers, employers, companies, beneficial relationships, and other relevant actors.
- **Accounts and identifiers:** emails, phone numbers, usernames, aliases, addresses, domains, wallets, device or technical identifiers, and customer-provided references.
- **Infrastructure and assets:** websites, hosting, certificates, repositories, technical services, digital assets, and other mission-relevant artifacts.
- **History and chronology:** first and last observation, changes in control, past affiliations, record dates, activity windows, and sequence.
- **Relationships and behavior:** shared identifiers, interactions, organizational ties, transactions, recurring patterns, and operating overlaps.
- **Exposure and change:** publicly discoverable information, compromised or disclosed material available through authorized sources, new records, deleted assets, changed associations, and other developments that affect the decision.

This graph is not a claim that every nearby node belongs to the subject. It is the working environment in which connections can be proposed, tested, accepted, rejected, or left unresolved.

A useful identity resolution therefore carries more than a match score. It shows the evidence supporting the match, conflicting indicators, relevant dates, source relationships, alternative candidates, and the reason a reviewer should accept or challenge the conclusion.

Identity is the anchor. Evidence determines the boundary.

05 / LIGHTHOUSE OSINT FIELD GUIDE

WHERE TEAMS USE LIGHTHOUSE FOR OSINT

The following mission patterns are illustrative. They show how LightHouse can be configured around different decisions; they are not representations of specific customer matters or guaranteed outcomes.

MISSION PATTERNS

- **INSIDER RISK** — Signal to supported context
- **RECON** — Public surface to exposure
- **FRAUD + TRUST** — Transaction to surrounding actors
- **ATTRIBUTION** — Artifacts to assessed origin
- **DUE DILIGENCE** — Record to practical control
- **ONBOARDING** — Claim to coherent history

Insider risk

An organization begins with an authorized signal involving an employee, contractor, or trusted user. LightHouse can resolve relevant external identities and entities, develop permitted relationships and exposure, preserve chronology, and help investigators distinguish an explainable anomaly from negligence, coercion, divided loyalty, impersonation, or deliberate harm.

The decision product is not a character judgment. It is a source-cited account of what was observed, which associations are supported, which explanations remain plausible, and what requires authorized human review.

Adversary reconnaissance and digital exposure

Security teams can examine the organization as an outside actor would. LightHouse can assemble publicly discoverable information across executives, employees, vendors, infrastructure, exposed credentials or records available through authorized services, forgotten assets, and operational dependencies.

The result is a prioritized view of what an adversary could plausibly connect and use, with the evidence needed to remove exposure, adjust monitoring, or test controls.

Transaction fraud and trust

A transaction can be authenticated and still be unsafe. LightHouse can help fraud teams examine the customer, recipient, associated accounts or entities, historical context, and signs of impersonation, coaching, coercion, mule activity, or coordinated behavior.

The purpose is to give the reviewer more than a risk signal. It is to show which facts support intervention, which observations are inconclusive, and what can be checked within the operating window.

Threat actor and infrastructure attribution

Analysts can use LightHouse to connect aliases, accounts, domains, certificates, repositories, infrastructure, behavior, and organizational relationships. The system preserves competing hypotheses and helps distinguish shared tooling or hosting from stronger evidence of common control.

The output is an assessed attribution with explicit confidence and limitations, not an unsupported declaration of identity.

Due diligence and counterparty investigations

An official record can establish that an entity exists without showing who controls it, how it has changed, or which relationships create practical risk. LightHouse can develop ownership and management, related entities, litigation or regulatory context where lawfully available, public claims, digital operations, adverse information, and inconsistencies across time.

The resulting intelligence product helps the customer decide what requires clarification, enhanced diligence, contractual protection, monitoring, or refusal.

Hiring, onboarding, and trust decisions

Where authorized and appropriate, LightHouse can help investigators examine whether a claimed professional or organizational history is coherent, whether material contradictions require review, and whether relevant external exposure changes the trust decision.

This work should be governed by applicable law, policy, proportionality, and human review. LightHouse supports the evidence process; it does not determine employment, eligibility, or legal rights.

06 / LIGHTHOUSE OSINT FIELD GUIDE

WHAT A SUPPORTED FINDING CONTAINS

A conclusion is inspectable when a qualified reviewer can understand what was observed, how the pieces relate, why the interpretation is reasonable, and where uncertainty remains.

ANATOMY OF A SUPPORTED FINDING

A reviewer can move from the conclusion back to the underlying evidence.

LightHouse structures consequential findings around seven elements:

1. **Claim:** The narrow statement the evidence is being used to support.
2. **Observations:** The underlying records, artifacts, or events, kept distinct from interpretation.
3. **Provenance:** Source, collection time, method, and any relevant preservation record.
4. **Source relationship:** Whether sources are independent, derivative, duplicated, or unknown.
5. **Chronology:** The time boundaries that make an association current, historical, or unresolved.
6. **Assessment:** The analyst's interpretation, confidence, rationale, and alternative explanations.
7. **Review status:** Who reviewed the finding, which controls applied, and what action remains with the authorized owner.

This structure helps prevent several common errors. A repeated claim does not become corroborated merely because it appears on more pages. A past address does not silently become a present residence. A shared provider does not establish common control. The absence of a record does not prove the absence of conduct. An automated inference does not become a fact because it appears in a polished report.

LightHouse can present the same underlying evidence through different decision products:

- **Identity graph:** for relationships, clusters, competing candidates, and source inspection.
- **Subject dossier:** for a coherent view of the person or entity, identifiers, chronology, exposure, and unresolved questions.
- **Narrative assessment:** for findings, alternatives, confidence, limitations, and implications.
- **Structured export:** for integration into an existing case, fraud, security, or intelligence workflow.
- **Alert:** for a defined material change against an approved baseline.
- **Briefing:** for a decision owner who needs the answer, evidence path, and remaining uncertainty in operational language.

The interface can change. The evidence standard should not.

07 / LIGHTHOUSE OSINT FIELD GUIDE

HUMAN AUTHORITY BY DESIGN

Agentic systems can pursue more lines of inquiry than a person can manage manually. That makes governance more important, not less.

HUMAN AUTHORITY

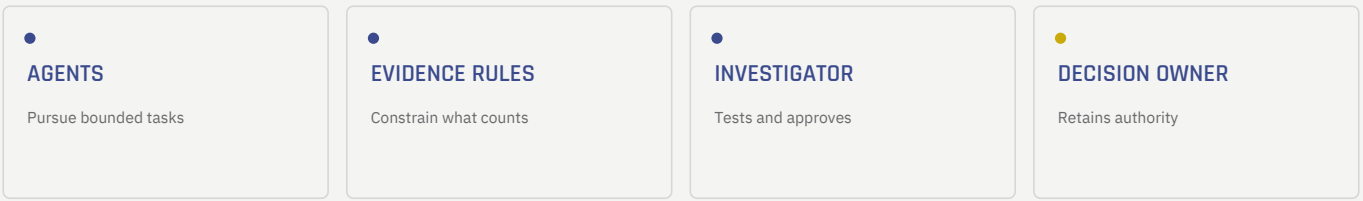

Every LightHouse mission should begin with an authorized purpose, permitted sources, collection boundaries, defined user roles, review points, retention expectations, and a named decision owner. Those controls should reflect the customer's legal, contractual, regulatory, security, and policy obligations.

Within that frame, agents can decompose questions, select approved capabilities, pursue leads, revisit thin areas, and assemble candidate findings. Evidence rules constrain what can be promoted into the decision product. Human investigators direct deeper inquiry, test alternative explanations, resolve ambiguity, and approve consequential assessments. The authorized owner remains responsible for action.

Practical controls can include:

- mission-specific source and capability permissions;
- access controls and separation of customer data;
- logged collection and analytical activity;
- provenance and collection timestamps;
- required review for sensitive findings or actions;
- explicit distinction between facts and assessments;
- visible contradictions and unresolved identity candidates;
- retention and deletion rules;
- evaluation against representative mission cases;
- escalation when evidence falls below the required threshold.

The Berkeley Protocol on Digital Open Source Investigations offers a useful reference for professional, legal, ethical, and security-aware handling of digital open-source material in its investigative context. The NIST AI Risk Management Framework offers a complementary structure for governing, mapping, measuring, and managing risks in systems that use artificial intelligence. LightHouse engagements translate relevant principles into the customer's mission, authorities, and operating environment.

No platform can define the customer's lawful authority or eliminate the need for expert judgment. The goal is controlled scale with visible accountability.

> Agents pursue the work. Evidence rules decide what counts. Humans remain responsible for action.

08 / LIGHTHOUSE OSINT FIELD GUIDE

MEASURE THE INVESTIGATION, NOT THE THEATER

An OSINT pilot should be evaluated against the mission it was built to improve. Source count, graph size, and search volume may describe activity, but they do not establish that the intelligence is useful.

MEASURE THE MISSION

01	02	03	04	05
TIME	EFFORT	EVIDENCE	QUALITY	VALUE
First supported finding	Analyst hours per case	Provenance coverage	Resolution + rejection	Decision changed

Start with a baseline from the current workflow and measure a small set of operational and evidentiary outcomes:

- **Time to first supported finding:** How quickly does the team reach a fact that materially advances the requirement?
- **Decision latency:** How long does it take to move from the initial question to a reviewable answer?
- **Analyst effort per case:** Which manual collection, reconciliation, and reporting tasks were reduced or redirected?
- **Provenance coverage:** What share of consequential claims retains a complete source and collection record?
- **Independent corroboration:** Can reviewers distinguish true source independence from repetition?
- **Identity resolution quality:** How often are candidate matches confirmed, rejected, or left unresolved after review?
- **Contradiction handling:** Are conflicting records visible and addressed rather than averaged away?
- **Review efficiency:** Can a qualified reviewer understand and challenge the evidence path without reconstructing the case?
- **Operational value:** Did the product change a decision, trigger a useful follow-up, reduce avoidable exposure, or establish that the evidence was insufficient?
- **Monitoring quality:** Do alerts represent material change with acceptable noise for the mission?

The right measures will differ across fraud, security, due diligence, insider risk, and attribution. What should remain constant is the connection between the metric and the decision.

09 / LIGHTHOUSE OSINT FIELD GUIDE

THREE WAYS TO OPERATE LIGHTHOUSE

The same OSINT capability can enter an organization in different ways.

OPERATING MODEL

CUSTOMER-OPERATED

Your analysts run the mission.

FORWARD-DEPLOYED

We build beside your team.

INTELLIGENCE PRACTICE

Our investigators deliver the product.

Customer-operated

Customer analysts operate LightHouse within their existing authorities, evidence rules, review points, and workflows. Silent Harbor supports configuration, integration, evaluation, and tradecraft transfer.

This model fits teams that already own the investigative mission and want an identity intelligence engine to make the work more coherent, repeatable, and inspectable.

Forward-Deployed Intelligence Engineering

A Forward Deployed Engineer and Intelligence Mission Lead work with the customer's decision owner and practitioners. Together they turn a consequential question into a bounded LightHouse workflow, connect the necessary systems and sources, define evaluations, and adapt the product to the operating environment.

This is not generic implementation support. It is engineering and intelligence work conducted close to the mission until the capability performs under real conditions.

Silent Harbor Intelligence Practice

Silent Harbor investigators operate LightHouse for customers that need the intelligence outcome without building or expanding an internal investigative function. The practice directs collection, tests alternatives, reviews evidence, and delivers the agreed decision product under the customer's authorities and controls.

Customers can begin with one model and change as the mission matures. The constant is the evidence standard and the customer's authority over the decision.

10 / LIGHTHOUSE OSINT FIELD GUIDE

START WITH ONE CONSEQUENTIAL QUESTION

The strongest starting point is a real mission with a clear owner, available baseline cases, and an answer that matters.

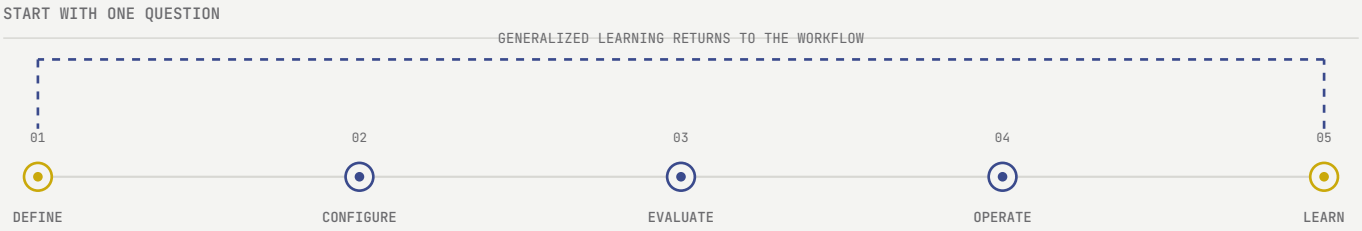

Define

Name the decision, subject, users, authorities, permitted sources, operating window, evidence threshold, human review points, and output. Record what the current process costs in time, effort, delay, and unresolved uncertainty.

Configure

Translate the mission into LightHouse agents, source permissions, identity and evidence rules, integrations, decision products, and logging. Use representative historical or synthetic cases that reflect the actual difficulty of the mission without exposing unauthorized information.

Evaluate

Run known and ambiguous cases. Test correct resolution, rejection of false links, chronology, source independence, contradiction handling, provenance, analyst control, security, and the quality of the final decision product.

Operate

Move into a bounded production mission with named owners and escalation paths. Measure the investigation against its baseline. Review misses, weak evidence, analyst interventions, and unexpected behavior.

Learn

Return generalized learning to the workflow, evaluation set, and core engine. Do not allow one case's sensitive facts to become another customer's data. Improve the capability without weakening the boundary.

The first objective is not an enterprise-wide OSINT transformation. It is proving that one consequential question can be answered with greater speed, stronger evidence, and clearer human accountability.

ABOUT LIGHTHOUSE

LightHouse is Silent Harbor's Identity Intelligence Engine. It resolves people and entities, develops the relevant environment around them, tests contradictions, and shows what the evidence supports.

Silent Harbor builds, deploys, and operates agentic intelligence systems for high-consequence decisions. Customers can operate LightHouse directly, work with a Forward-Deployed Intelligence Engineering team to build a mission-specific capability, or engage the Silent Harbor Intelligence Practice to run the investigation and deliver the intelligence product.

To discuss an OSINT mission, visit **https://silentharbor.com/Contact.html#lighthouse**.

Selected references

1 Office of the Director of National Intelligence, **The IC OSINT Strategy 2024–2026**. https://www.dni.gov/files/ODNI/documents/IC_OSINT_Strategy.pdf

2 Office of the United Nations High Commissioner for Human Rights and Human Rights Center, University of California, Berkeley, **Berkeley Protocol on Digital Open Source Investigations**. https://www.ohchr.org/sites/default/files/2022-04/OHCHR_BerkeleyProtocol.pdf

3 National Institute of Standards and Technology, **Artificial Intelligence Risk Management Framework (AI RMF 1.0)**. https://doi.org/10.6028/NIST.AI.100-1

Scope note: This paper describes a general operating approach. Available sources, capabilities, and controls depend on the customer's authorities, deployment, source rights, mission design, and applicable obligations. Illustrative mission patterns are not customer claims or promised outcomes.