

Shoreline Framework

Status, Stewardship, and Refusal in Information Movement

Christopher Knorr

Preprint — June 2026

Abstract

This article introduces the Shoreline Framework, a status-transition account of information movement under conditions of stewardship, authority, refusal, mediation, and repair. Contemporary information systems often treat availability as sufficient for use: visible material as public, stored material as rightly held, preserved material as legitimate, summarized material as sourced, and algorithmic mediation as a substitute for record-paths. Shoreline rejects those equivalences. It distinguishes objects from statuses, statuses from access, access from authority, preservation from legitimacy, and mediation from recordhood.

The framework operates across two registers. Record ecology concerns conditions that allow objects to remain traceable, checkable, contestable, and repairable. Relational knowledge governance concerns knowledge whose movement depends on authority, relation, place, role, consent, protocol, or refusal. Shoreline does not define relationally governed knowledge from the inside. It specifies an interface problem: what outside systems owe when they encounter material whose authority conditions they do not control.

The paper develops the framework's vocabulary—object, status, condition, transition, authority ecology, protocol position, stewardship, record resistance, holding legitimacy, refusal, exposure, routing, and synthetic closure—and applies it to four public Register 1 or mediation-layer calibration cases: the 2023 Roald Dahl edition controversy, 2025 federal data removals, CVE-2024-3094 involving XZ Utils, and Google AI Overviews in 2024. These cases demonstrate preliminary discriminating capacity but do not validate the companion instrument suite. The contribution is a boundary discipline for preserving answerability in information movement without collapsing access into extraction.

Keywords: information governance; archives; data governance; AI mediation; refusal; Indigenous data sovereignty; record ecology; synthetic closure; public records; knowledge governance; stewardship; provenance; authority; data ethics

Contents

Author Note and Positional Boundary	2
Core Proposition	2
I. Introduction: The False Comfort of Availability	2
II. The Problem: Status Transfer Without Conditions	3
III. The Formal Core: Objects, Statuses, Conditions, and Transitions	4
IV. Two Registers and the Interface	8
V. Record Resistance Is Not Holding Legitimacy	10
VI. Refusal as a Complete Governance Operation	12
VII. Mediation, AI, and Synthetic Closure	13
VIII. Calibration Cases: What the Framework Can Currently Show	14
IX. The Companion Operator Packet	16
X. Rival Explanations and Why Shoreline Is Needed	17
XI. Limits and Validation Pathway	19
XII. Conclusion: The Shoreline as Boundary Discipline	20
Status and Limitations	21
Appendix A — Proposed Next Stages of Development	21
Stage 1 — Manuscript Stabilization	22
Stage 2 — External Scoring Preparation	22
Stage 3 — Independent External Scoring	22
Stage 4 — Divergence Repair and Instrument Revision	23
Stage 5 — Register 2 / Standing Protocol Development	23
Stage 6 — Public-Facing Translation	23
Stage 7 — Maintenance, Misuse Monitoring, and Version Governance	24
Stage 8 — Second Article or Technical Note	24
Development Summary	25
References	25

Author Note and Positional Boundary

This paper is an independent framework proposal. It does not claim to speak for any tribe, nation, community, family, institution, archive, museum, or authority ecology. Its Register 2 claims are intentionally bounded: Shoreline does not define relationally governed knowledge from within, but addresses responsibilities of outside systems when they encounter material whose authority conditions they do not control.

The accompanying operator packet is a draft companion artifact. It has been internally calibrated on selected public cases and prepared for future external scoring. It has not been externally reviewed, validated, certified, community-reviewed, or authorized for Register 2 application.

Core Proposition

Information does not become knowable, holdable, shareable, actionable, or refusable merely because it exists, circulates, is stored, appears public, or is summarized. It takes status only under maintained conditions: provenance, custody, context, authority, protocol, answerability, repair, and refusal.

The paper's governing distinctions are compact: **access is not authority; exposure is not openness; preservation is not legitimacy; mediation is not the record; refusal is not erasure.**

I. Introduction: The False Comfort of Availability

A great deal of contemporary information practice rests on a quiet assumption: availability settles the question. If a document can be found, it is treated as usable. If a dataset is online, it is treated as public. If an archive holds a record, the holding is treated as legitimate. If an AI system summarizes a body of material, the summary is treated as a sufficient substitute for the path back to sources. If a webpage disappears, the disappearance is treated either as erasure or as a technical loss that can be repaired by locating a copy elsewhere. In each case, a change in access is mistaken for a change in status (Upward 1996; Society of American Archivists n.d.-b; OpenGovData 2007; Open Data Charter n.d.).

The Shoreline Framework begins from a different premise. Information does not become knowable, holdable, shareable, actionable, or refusable merely because it exists, circulates, is stored, appears public, or is summarized. It takes status only under maintained conditions. Those conditions may include provenance, custody, context, authority, protocol, answerability, repair, and refusal. When those conditions are ignored, systems create a false comfort: the user sees something, the institution holds something, or the model summarizes something, and the movement is treated as legitimate because it has become technically possible.

The problem is not only epistemic. It is also governmental, institutional, and relational. The question is not simply whether a claim is true, a file exists, or a source can be reached. The question is what kind of object is being encountered, what status it currently holds, what conditions maintain that status, who has authority over its movement, and what transitions are permitted, blocked, repaired, or refused (Carroll et al. 2020; First Nations Information Governance Centre n.d.; Local Contexts n.d.-a; Wilkinson et al. 2016).

“Shoreline” names this boundary discipline. A shoreline is not simply a wall. It is a place of movement, contact, and constraint. Things cross it, wash up against it, recede from it, and are transformed at it. The image is useful because information movement is rarely a simple yes-or-no matter. The same object may be visible but not authorized for use; stored but not legitimately held; preserved but stripped of context;

accessible but not actionable; known in one register but unfit for movement into another. Shoreline asks what must be true before a transition may occur.

The framework is designed for a world in which mediation has become ordinary. Search engines, archives, libraries, museums, platforms, dashboards, databases, and AI systems constantly retrieve, rank, summarize, format, store, expose, preserve, and recombine. These systems can make information more available, but they can also flatten differences between access, authority, legitimacy, and answerability. The problem is acute when mediation produces closure: a polished answer, a clean interface, a stable catalog record, or a summarized claim can feel complete even when the source-path is weak, context has been stripped, authority is unresolved, or routing was required (Colavizza et al. 2021; Ji et al. 2023; Google 2024).

Shoreline’s governing distinctions are therefore not decorative. They prevent analysis from collapsing into a single question of availability while also preventing the opposite error: treating all consequential information as restricted or unusable. Public emergency notices, public vulnerability records, court notices, weather alerts, and ordinary public datasets often serve their function precisely by being open, traceable, and actionable. Shoreline is not anti-access. It is anti-flattening.

The paper develops Shoreline as a status-transition framework for information movement. It first identifies the problem of status transfer without conditions, then introduces the formal core: objects, statuses, conditions, transitions, authority ecologies, protocol positions, and stewardship. It distinguishes record ecology from relational knowledge governance, analyzes the interface between them, and develops the concepts of record resistance, holding legitimacy, refusal, mediation, and synthetic closure. The companion operator packet demonstrates operational capacity but remains a draft future validation object, not a validated tool.

II. The Problem: Status Transfer Without Conditions

The failure Shoreline addresses can be stated simply: information systems often transfer status without checking the conditions that make the transfer valid. A thing is seen, so it is treated as public. A thing is stored, so it is treated as rightly held. A thing is preserved, so it is treated as repaired. A thing is summarized, so it is treated as known. A thing is deleted from one surface, so it is treated as erased. A thing is refused, so the refusal is treated as an obstacle rather than as a governance act.

These are variations of one deeper pattern. A system observes a change in access or presentation and treats that change as if it automatically carried a change in authority, legitimacy, evidence, or actionability. The conditions that should govern the transition disappear. The object moves, but the authority ecology does not. The claim circulates, but the source-path does not. The record survives, but the legitimacy problem remains. The answer appears complete, but the conditions of knowledge have not been maintained.

Public record practice makes the problem visible. A public record is not merely a file. It depends on provenance, custody, indexing, fixity, accountability, and contestability. If those conditions decay, the object may still exist while its record status weakens. A screenshot of a government page, a copied PDF, a cached dataset, and an official record may contain similar information while differing in provenance, authority, update status, legal effect, and repair capacity (Upward 1996; Society of American Archivists n.d.-a; Digital Preservation Coalition n.d.; National Digital Stewardship Alliance 2014).

Data removal and restoration show the same issue. When a dataset disappears from a public agency surface, it has not necessarily been erased; it may survive in archives, mirrors, citations, repositories, institutional copies, or legal records. But neither has it necessarily been preserved in the full sense. A dynamic resource may have no complete copy. A restored page may return altered. A dataset may remain accessible

while losing metadata, update path, context, or authority trail. Treating the episode as simply “available” or “unavailable” misses the object-class structure of the problem (Kaiser Family Foundation 2025; Freilich et al. 2025; Data Rescue Project 2025; State Library Resource Center 2025).

A different version of the problem appears when public visibility is treated as openness. Something may be online because it was uploaded, scraped, leaked, indexed, cataloged, or exposed. None of those facts alone establishes that the object is open in the authority-bearing sense. The actual protocol position of an object is not set by the observer who encounters it. It is set by the authority ecology with legitimate standing over its movement. Apparent openness may be exposure rather than authorization (Carroll et al. 2020; First Nations Information Governance Centre n.d.; Local Contexts n.d.-a; Couldry and Mejias 2019).

Preservation introduces another collapse. Preservation is often treated as inherently good, and in many contexts it is. But preservation is not the same as legitimacy. A museum may preserve an object under excellent environmental, cataloging, and conservation conditions while lacking legitimate authority to hold it. An archive may maintain a strong record ecology around material whose circulation remains contested. Without a framework that separates record resistance from holding legitimacy, preservation can become a way of stabilizing wrongful holding (National Park Service 2024; Bureau of Indian Affairs n.d.; Local Contexts n.d.-b).

The inverse is also true. Low external record resistance does not necessarily imply low legitimacy. A living tradition, family-held narrative, seasonal practice, place-bound teaching, or community-governed form of knowledge may not be highly resistant in the archival sense. It may not be multiply copied, independently indexed, fixed, or externally verifiable. Its conditions may be relational rather than documentary. A record system that treats low archival resistance as low value will misrecognize the object. A system that tries to “repair” fragility by maximizing capture may violate the conditions that make the knowledge legitimate (Carroll et al. 2020; First Nations Information Governance Centre n.d.; Local Contexts n.d.-a; Tuck and Yang 2014).

AI and search mediation intensify the problem because they can perform status transfer rhetorically. Increasingly, a user receives not a list of sources but an answer. The answer may be fluent, compressed, and formatted as if inquiry has ended. But it may have transformed a weak source ecology into an apparently settled claim, collapsed satire or forum discussion into advice, or supplied actionability without answerability. Shoreline calls this synthetic closure: the production of a completed-looking answer that hides the conditions under which the claim could be traced, checked, authorized, contextualized, or corrected (Ji et al. 2023; Huang et al. 2023; Zhang et al. 2023; Google 2024).

The framework therefore begins with a negative discipline: do not transfer status merely because movement is possible. Do not infer openness from exposure, legitimacy from custody, knowledge from fluency, repair from preservation, erasure from disappearance, or obstruction from refusal. Once those false inferences are blocked, the object can be identified, its status described, its conditions named, its authority ecology located or routed, and its transitions assessed.

III. The Formal Core: Objects, Statuses, Conditions, and Transitions

The Shoreline Framework separates terms often collapsed in ordinary information practice. An object is not the same as its status. A status is not access. Access is not authority. Authority is not possession. Preservation is not legitimacy.

Figure 1. Shoreline routing-decision logic.

Figure 1. Shoreline routing-decision logic

The permission gate only; full transition set in Section III.




Figure 1. Shoreline routing-decision logic. This figure depicts the permission gate only: given an object and a proposed transition, the system permits status gain or movement, receives a refusal or block, or—where authority is unresolved—routes. The full transition set (gain, decay, revocation, repair, block) is defined in Section III; decay, revocation, and repair are distinct dynamics, not outputs of this gate, and are not drawn here.

An object is any interpretable difference or carrier: a document, file, dataset, recording, image, story, catalog entry, summary, artifact, webpage, index result, model output, place-feature, or trace. The term is intentionally broad because Shoreline applies to any carrier subject to interpretation, custody, circulation, preservation, refusal, or use.

But an object does not carry all statuses automatically. A file is not a record merely because it exists. A trace is not data merely because it can be observed. Data is not evidence merely because it is stored. Evidence is not knowledge merely because it is cited. Knowledge is not actionable merely because it is available. Each status depends on conditions that may be technical, institutional, epistemic, legal, social, cultural, relational, or procedural (Upward 1996; Wilkinson et al. 2016; Frings-Hessami 2021).

A status is the standing an object holds while and only while its conditions are maintained.¹ This means

¹A parallel account in status-dynamic self-concept theory sharpens, by analogy, why status is not a decorative label. Bergner and Holmes (2000) distinguish a *status* conception from an *informational-summary* conception of the self: what matters is not only the facts associated with a thing but its position within a field of relations, powers, and possible moves. Shoreline borrows the analogy,

status is not an intrinsic property of the object alone. A government webpage, screenshot, cached copy, printed copy, and court-certified copy may contain overlapping information while holding different statuses because they differ in provenance, custody, authority, fixity, certification, and repair path.

Conditions are requirements under which a status is gained, maintained, lost, blocked, or repaired. For record status, they may include provenance, custody, indexing, fixity, accountable stewardship, and contestability. For evidentiary status, the object must be connected to a claim-path. For knowledge status, the object must be embedded in practices of interpretation, correction, and transmission. For actionable status, the object must be sufficiently answerable to support responsible action.

This prevents a common error: treating information movement as a one-way ladder from trace to data, data to evidence, evidence to knowledge, and knowledge to action. Shoreline treats the movement as a lattice. An object may hold multiple statuses at once, gain one status while losing another, be preserved as a record while losing contextual knowledge, become actionable in a narrow emergency context without becoming broadly shareable, or be visible without being open (Upward 1996; Frings-Hessami 2021; Wilkinson et al. 2016).

Figure 3. Information movement is a lattice, not a one-way ladder.

not the psychology—record status and self-concept status are not identical, and the framework does not treat a psychotherapy formulation as information-science authority. Held to that boundary, the parallel reinforces the formal point that status governs what facts can do: a fact changes an object’s standing only through an available transition—evidence, correction, repair, revocation, authorization, or refusal.

**Figure 3. Information movement is a lattice,
not a one-way ladder**




Figure 3. The ladder Shoreline rejects (left) treats movement as one-way and irreversible, each status following the last. The lattice (right) treats status as conditional and reversible: an object may hold several statuses at once and gain one while losing another, and a status may decay, be repaired, revoked, or blocked. The five transitions are gain, decay, revocation, repair, and block.

A transition is a change in status or protocol position: gain, decay, revocation, repair, or block. Gain occurs when necessary conditions are satisfied. Decay occurs when maintenance conditions weaken. Revocation occurs when status is withdrawn by legitimate authority or rendered invalid by changed conditions. Repair occurs when lost or weakened conditions are restored or sufficiently substituted. Block occurs when a transition is stopped, legitimately through refusal or illegitimately through obstruction.

Authority ecology names the holder or configuration of holders with legitimate standing over a protocol position and over authorization-requiring transitions. Authority may be individual, collective, legal, familial, institutional, professional, ceremonial, jurisdictional, or plural. The point is not who possesses an object, but who has standing to decide how it may be held, moved, shared, refused, repaired, translated, indexed, trained on, or acted upon.

Protocol position is an authority-set position governing movement. Actual positions include open, governed, restricted, and held. These are not mere access states. An object is open when its authority conditions permit broad movement. It is governed when movement is permitted under rules, roles, or conditions. It is restricted

when movement is limited to specified persons, contexts, purposes, or times. It is held when its proper condition is nonmovement or internal custody.

Stewardship is the labor of maintaining the conditions that allow an object to hold its proper status at the proper protocol position. Stewardship is not simply preservation. It may require access, restriction, indexing, de-indexing, repair, refusal, or route. It is not maximal persistence; it is condition-appropriate care (Wilkinson et al. 2016; Carroll et al. 2020; First Nations Information Governance Centre n.d.; Local Contexts n.d.-a).

IV. Two Registers and the Interface

Shoreline distinguishes two registers: record ecology and relational knowledge governance. The distinction is not meant to divide the world into rigid categories. Many objects operate across both. The purpose is to prevent one register's conditions from being mistaken for the other's.

Register 1 is record ecology. It includes the institutions and practices through which documents, records, datasets, evidence, publications, archives, court materials, administrative files, scientific outputs, and public information systems are maintained. Its core conditions include provenance, custody, indexing, fixity, accountability, contestability, and repair infrastructure. A Register 1 system is healthy when objects can be traced, checked, compared, corrected, restored, and held accountable within a record path (Upward 1996; Society of American Archivists n.d.-b; Wilkinson et al. 2016; OpenGovData 2007).

Register 2 is relational knowledge governance. It includes forms of knowledge whose meaning, authority, legitimacy, movement, or use depends on relations that cannot be reduced to documentary status. These relations may involve kinship, place, season, ceremony, family, role, community, nationhood, language, responsibility, consent, refusal, or living transmission. Register 2 is not simply "private information" or "cultural information." It is knowledge governed by relations that determine how, when, where, by whom, and whether something may move (Carroll et al. 2020; First Nations Information Governance Centre n.d.; Local Contexts n.d.-a).

The framework must be careful here. Shoreline does not claim to define relationally governed knowledge from the inside. It does not claim to speak for any tribe, nation, family, community, ceremonial authority, or knowledge-holder. Its claim is narrower: outside systems must not assume that the conditions of Register 1 exhaust the conditions of legitimacy. A record system may ask whether something is documented, indexed, preserved, or accessible. Those questions do not settle whether the object is properly movable, shareable, extractable, translatable, trainable, or actionable.

The interface is the zone where these registers meet. Record systems constantly encounter objects whose authority conditions they do not control: archives hold community materials, museums hold ceremonial objects, researchers hold recordings, platforms expose governed material, AI systems train on scraped corpora, search engines index material detached from context, and public datasets aggregate protected sub-material. In each case, Register 1 may describe provenance, custody, or access, but may not determine actual protocol position.

A Register 1 system reaches the interface when ordinary record-ecology conditions are insufficient to determine whether movement is legitimate. Triggers include evidence that the object concerns community-governed, ceremonial, familial, place-based, seasonal, or role-specific knowledge; signs of contested custody or repatriation claims; metadata, labels, notices, or provenance indicating protocol restrictions; protected sub-material embedded inside otherwise public data; uncertainty about whether the object is open, governed,

restricted, or held; requests for removal, return, correction, or refusal by persons or groups making a credible standing claim, which triggers route rather than final recognition by the system; or a proposed use that changes the object’s status, such as training, public display, translation, aggregation, or action-bearing summary.

When these triggers appear, the system must not treat uncertainty as permission. It should route. Route is a temporary nonauthorization state that halts outward movement while preserving only the minimum conditions necessary to identify the relevant authority ecology, prevent further exposure, and determine whether the proper outcome is access, conditional movement, repair, refusal, return, or continued hold. In practical terms, route means no publication, training, display, extraction, translation, or action-bearing summary while authority is unresolved, except for legally required preservation, protective containment, or minimal accountability logging that does not increase exposure. It also requires a minimal encounter record for accountability and an active search for the proper authority path.

Shoreline’s proposed interface obligations are inquiry before movement, deference to legitimate authority, acceptance of valid refusal, support for repair, conditional carriage, reversibility where possible, and no laundering of exposure into authorization. These are not certified community protocols. They are constraints on outside systems. “Inquiry before movement” means the system may not treat ignorance as permission. “Refusal acceptance” means valid refusal is a completed governance outcome, not a failed access request. “No laundering” means scraping, indexing, summarizing, archiving, or training does not convert exposed material into open material (Couldry and Mejias 2019; Carroll et al. 2020; Local Contexts n.d.-a; Tuck and Yang 2014).

This interface account also protects against the opposite mistake: routing everything. Some information is properly open. Public vulnerability records, bus schedules, court notices, weather alerts, emergency communications, and many government datasets are meant to circulate. Their public value depends on rapid, traceable, checkable movement. Shoreline therefore distinguishes low-thickness public-function information from thick relational knowledge. The same boundary discipline that prevents extraction also prevents unnecessary over-restriction.

Relational thickness names the degree to which meaning, legitimacy, authority, and movement depend on relation. Thin information can usually move with ordinary source-path and context requirements. Thick information requires role, protocol, place, authority, or consent. When authority, register, thickness, or protocol position cannot be responsibly determined, Shoreline uses routing. Routing is not indecision. It is a disciplined response that prevents the system from converting uncertainty into permission.

Two clarifications keep routing from overreaching. First, routing is triggered by a *proposed status-changing transition*, not by an object’s mere existence. A box of unprocessed material in a legacy backlog is not “routed” simply because its authority conditions are unresolved; it is routed only when some outward transition is proposed for it—digitization, publication, aggregation, translation, training, or action-bearing summary. Unresolved authority over an object that no one is attempting to move imposes no obligation beyond ordinary custody. This prevents the framework from recommending a functional deep-freeze of the vast holdings whose provenance is, and will long remain, incompletely resolved.

Second, routing does not require a system to investigate the relations it is trying to protect. The framework asks an outside system to assess *detachability*, not to establish *thickness*. Detachability is read from signals the system already legitimately holds—provenance and its gaps, metadata, accompanying documentation, custody history, labels or notices, and credible standing claims—not from inquiry into the content, community, or relation itself. Because unresolved conditions resolve toward governed rather than open, the system is never required to investigate in order to protect; investigation is the cost of movement, never of protection. Absence of evidence that an object is thin and detachable is itself sufficient to route. The boundary the

framework defends is therefore not breached by the act of defending it (Carroll et al. 2020; Local Contexts n.d.-a; Tuck and Yang 2014).

V. Record Resistance Is Not Holding Legitimacy

One of Shoreline's central distinctions is between record resistance and holding legitimacy. Record resistance asks whether an object or record ecology can survive alteration, disappearance, substitution, loss, or erasure. Holding legitimacy asks whether the object is rightly held, governed, circulated, or used. A system can score highly on one axis while failing badly on the other.

Record resistance is ecological. It depends on conditions that allow an object to remain traceable, checkable, comparable, and repairable over time: multiple copies, independent custodians, stable identifiers, cataloging systems, citations, version histories, audit trails, public references, accountable stewards, and correction mechanisms. When such an ecology is strong, alteration is more likely to be detected, disappearance less likely to become erasure, and damage more likely to be repaired (Society of American Archivists n.d.-a; Digital Preservation Coalition n.d.; National Digital Stewardship Alliance 2014; Aturban et al. 2019).

The 2023 Roald Dahl edition controversy shows record resistance under pressure. The relevant issue was not merely that revisions occurred. Published texts are revised, edited, abridged, translated, annotated, and republished often. The Shoreline question is whether the altered edition could replace the prior textual ecology in a way that made the earlier form functionally disappear. In that case, prior editions existed across homes, libraries, schools, used-book markets, catalogs, reviews, quotations, and public memory. The changes were detectable because comparison paths remained available. Public criticism created pressure. The publisher's later decision to keep classic texts in print functioned as repair within a resistant record ecology (Penguin 2023; Associated Press 2023; Guardian 2023).

A weaker ecology behaves differently. A dynamic public webpage, uncaptured dataset, unversioned database entry, platform-only post, or model-generated answer may disappear or change without leaving a stable comparison path. If no copy exists outside the original surface, no independent custodian exists, no fixed version can be compared, and no accountability path remains, loss may approach functional erasure.


Holding legitimacy asks a different question: whether the object is rightly held or governed by the person, institution, system, or ecology controlling it. A museum may preserve an object with care while lacking legitimate authority to hold it. An archive may have stable custody over material that should have been returned, restricted, or governed by another authority. A platform may retain exposed material against the authority conditions that should govern it. In such cases, preservation can become stabilization without repair. The object is safe in one sense but wrongfully situated (National Park Service 2024; Bureau of Indian Affairs n.d.; Local Contexts n.d.-b).

The reverse is also true. A fragile object may be legitimately held. A family story, seasonal teaching, oral history, local memory, ceremonial relation, place-based knowledge, or community-governed recording may have low external resistance. It may not be multiply copied, cataloged, fixed, indexed, or independently verifiable. Those facts may reflect proper conditions rather than failure. Imposing external resistance may damage the relation that makes the object legitimate (Carroll et al. 2020; First Nations Information Governance Centre n.d.; Local Contexts n.d.-a; Tuck and Yang 2014).

Record resistance and holding legitimacy therefore form a two-axis space. High resistance and high legitimacy describe a working archive or robust public record ecology. High resistance and low legitimacy describe a vault: the object survives, but the holding relation is wrong. Low resistance and high legitimacy

describe endangered but rightly held living knowledge or fragile local memory. Low resistance and low legitimacy describes an erasure zone. The framework's point is not to force every object toward maximum resistance. It is to ask what target conditions are proper for that object under the relevant authority ecology.

Figure 2. Record resistance and holding legitimacy as independent axes.



The axes are orthogonal: an object may score high on one and low on the other.
 The aim is not to maximize resistance but to ask what target conditions are proper under the relevant authority ecology. Different failures need different repairs: resistance repair moves an object upward; legitimacy repair, rightward.

Figure 2. The two axes are orthogonal: an object may score high on one and low on the other. A vault is well preserved but wrongfully held; endangered living knowledge is rightly held but fragile, and its proper response is support rather than extraction. The framework's aim is not to maximize resistance but to ask what target conditions are proper under the relevant authority ecology; a damaged record needs resistance repair (move upward), while a wrongfully held object needs legitimacy repair (move rightward).

This distinction clarifies repair. A damaged record may need resistance repair: restoring provenance, recovering copies, rebuilding indexes, documenting changes, or creating comparison paths. A wrongfully held object may need legitimacy repair: return, repatriation, restricted access, protocol-governed custody, shared stewardship, or refusal acceptance. A mediation failure may need answerability repair: source-path restoration, context marking, uncertainty disclosure, or withdrawal of a misleading summary. These repairs are not interchangeable. A copy does not repair consent. A catalog entry does not repair custody. A fluent summary does not repair source loss.

VI. Refusal as a Complete Governance Operation

Refusal is often misunderstood because information systems are built around movement. They collect, retrieve, expose, index, preserve, summarize, translate, circulate, and act. Within such systems, refusal can appear as friction: missing data, restricted access, noncooperation, secrecy, incompleteness, or obstruction. Shoreline begins from a different premise: refusal can be a complete governance operation.

A valid refusal is a transition-block exercised by a legitimate authority-holder on a specified edge. The edge matters. A refusal may block publication but not preservation, public circulation but not internal custody, digitization but not description, training but not citation, translation but not storage, or display but not return. Refusal is not general negativity. It is a structured act governing a transition (Tuck and Yang 2014; Carroll et al. 2020; First Nations Information Governance Centre n.d.; Local Contexts n.d.-a).

This definition separates refusal from absence. If an object is missing, inaccessible, unknown, or unrecorded, that is not automatically refusal. It may be loss, decay, silence, neglect, secrecy, exposure failure, or lack of source path. Refusal requires a legitimate authority-holder and a specified transition being blocked. It also separates refusal from obstruction. A museum that resists a legitimate repatriation claim does not thereby “refuse” in Shoreline’s sense if it lacks legitimate authority over that transition. Its block may be obstruction. A government office that deletes records to evade accountability is blocking answerability, not exercising valid refusal.

The legitimacy of refusal depends on authority ecology. Authority may be legal, communal, familial, institutional, professional, ceremonial, jurisdictional, or plural. It may require process, be contested, or be invisible to an outside observer. When authority is unresolved, the proper response is not to infer permission. The proper response is to route.

It is worth distinguishing route from refusal directly, since both halt movement. Route is the system’s own response to uncertainty: a temporary, self-imposed nonauthorization held while the proper authority is identified. Refusal is the system’s reception of a legitimate block issued by an authority that has standing. Route is provisional and resolves when authority is found; refusal is a completed governance outcome. A system routes because it does not yet know; it accepts a refusal because an authority has decided.

Silence is not authorization. Lack of visible restriction is not authorization. Public availability is not authorization. Searchability is not authorization. A scraped corpus does not authorize training merely because the material was reachable. Systems that treat silence as permission convert ignorance into power (Tuck and Yang 2014; Carroll et al. 2020; Couldry and Mejias 2019; Local Contexts n.d.-a).

Refusal also differs from privacy. Privacy usually concerns access, disclosure, identifiability, personal control, or exposure. Refusal may involve privacy, but it is broader. It can govern movement, transformation, translation, training, preservation, display, circulation, interpretation, or action. A community may refuse extraction not because material is unknown, but because the conditions of movement are not met. A family may refuse public circulation of a story not because the story lacks value, but because its value is inseparable from relation and responsibility.

This matters in AI and platform contexts. Large-scale systems often treat available material as potential training, indexing, summarizing, or answerable material. Once ingested, prior conditions become hard to see. The model may reproduce patterns without preserving source-path, authority, or protocol position. A refusal expressed at source level may disappear at output level. Worse, a model may generate fluent approximations of material whose movement should have been blocked. Processing is not authorization. Scraping is not stewardship. Training is not consent. Indexing is not publicness. Summarization is not answerability (Couldry and Mejias 2019; Carroll et al. 2020; Roberts and Montoya 2023; Ji et al. 2023).

Refusal can also repair. It may identify a broken condition: improper holding, decontextualized records, exposure without authority, misleading summaries, protected sub-material, or use beyond the purpose for which access was granted. In that sense, refusal is not merely prohibitive. It identifies the edge at which movement has outrun legitimacy.

The framework must also acknowledge hard cases. Refusal can be asserted falsely. Institutions can claim refusal to avoid accountability. Governments can hide records. Corporations can invoke restriction to shield wrongdoing. Archives can over-restrict. A claimed refusal is not automatically valid because it uses refusal language. Shoreline therefore distinguishes refusal from obstruction and asks: who is blocking what transition, by what authority, under what conditions, and with what accountability?

Without refusal, every model of information movement tilts toward extraction. It may speak of access, preservation, openness, innovation, research, or public benefit. But if it lacks a real place for valid nonmovement, it will treat governed boundaries as temporary obstacles to technical completion. Refusal is the concept that prevents the framework from assuming that all information ultimately wants to circulate.

VII. Mediation, AI, and Synthetic Closure

Shoreline treats AI systems, search engines, recommendation systems, platform feeds, dashboards, catalogs, and summaries as mediation layers. A mediation layer stands between an object and a user. It retrieves, ranks, compresses, translates, formats, predicts, summarizes, or presents. Mediation can be valuable, but it becomes dangerous when mistaken for the record itself (Colavizza et al. 2021; Ji et al. 2023; Google 2024; Huang et al. 2023).

Mediated output often arrives with the appearance of completion. A search result ranks one source above another. A dashboard converts many objects into a clean visual summary. An archive interface displays a catalog description while hiding gaps in provenance. A platform presents exposed material as if it were simply public. An AI system produces a fluent paragraph that reads as if inquiry has been completed. The user encounters an answer-like form, and the form substitutes for the conditions that would make the answer accountable.

This is not merely a problem of factual error. A false answer is obviously a problem, but a true answer may also be poorly mediated. The issue is whether the answer remains answerable. Can the user trace it back to its source ecology? Can the user check it? Does the output preserve context, date, uncertainty, jurisdiction, source-type, disagreement, satire, role, or protocol? Can it be corrected? Or does it convert a complex field into a smooth surface?

Synthetic closure occurs when a mediated output makes inquiry appear complete while hiding the conditions under which the claim could be traced, checked, authorized, contextualized, or corrected. The output may be grammatically polished, visually clean, institutionally branded, or computationally confident. It may even contain links. But if the user's epistemic situation is closed rather than opened, the mediation layer has performed an invalid status transition (Ji et al. 2023; Zhang et al. 2023; Huang et al. 2023; Google 2024).

This is why Shoreline does not reduce AI failure to hallucination. Hallucination names one failure: the system produces content not grounded in the relevant source field. Synthetic closure names a broader failure: the system produces an answer-like output that makes inquiry appear settled when the conditions of knowledge have not been maintained. Synthetic closure can occur without outright fabrication. The output may be based on real material and still fail because it collapses context, uncertainty, authority, or contestability.

The Google AI Overviews episode is useful as a public calibration case because it makes the structure visible. The problem was not only that some outputs were odd or inaccurate. The deeper issue was that a search mediation layer produced completed-looking answers from weak or unsuitable source contexts, including satirical material and forum-derived content. The user did not merely see a source list. The user saw a synthesized answer. The output had moved from retrieval into apparent knowledge (Google 2024; Business Insider 2024b; Ji et al. 2023).

The risk increases when output is action-bearing. A bad summary of a literary controversy may confuse. A bad summary of a medical, legal, safety, financial, or cybersecurity issue may harm. Shoreline therefore treats high-stakes mediation as requiring a stronger route trigger. If a system converts weak sources into health advice, legal advice, safety instructions, exploit guidance, or protected-context claims, the issue is not merely quality. The system is making something actionable without the conditions that would justify action.

Shoreline’s answer is not to reject mediation. It is to govern mediation through answerability. A good mediation layer preserves source-path, context, uncertainty, contestability, and repair. It distinguishes navigation from substitution. It says when it is summarizing, inferring, relying on weak sources, encountering unresolved authority, or routing rather than answering.

A clarification about architecture forestalls a common objection. Shoreline governs the mediation system *as deployed*, not the bare foundation model. A transformer’s weights compress and synthesize across a corpus; demanding a one-to-one source-path from the weights themselves would misread the architecture. But consumer mediation systems are not bare models. They retrieve, rank, format, caption, and present, and it is at that deployed layer that answerability is owed. Where a system retrieves — as retrieval-augmented generation systems do (Lewis et al. 2020) — the answerability conditions apply to the retrieval-and-synthesis path. Where a system generates without retrieval, answerability does not mean a fabricated citation; it means epistemic transparency—an interface that reflects the synthesized, uncertain, and possibly unsourced character of the output rather than presenting it in an authoritative form. The architectural limit is therefore not an excuse but a trigger: if a system’s architecture cannot sustain source-path, context, and contestability for an action-bearing claim—in health, legal, safety, financial, or cybersecurity domains—the proper protocol position is to route, warn, degrade gracefully, or withhold, not to synthesize closure. Inability to meet the conditions of answerable knowledge is a reason not to assert knowledge, not a license to assert it anyway (Ji et al. 2023; Google 2024).

VIII. Calibration Cases: What the Framework Can Currently Show

The Shoreline Framework is accompanied by a draft operator packet, but the instruments are not validated. They have been internally calibrated and prepared for future external scoring. The cases discussed here demonstrate preliminary discriminating capacity. They show that the framework can make distinctions ordinary access-based analysis tends to collapse. They do not prove that the instruments are ready for certification, compliance, or public deployment.

The current calibration set includes four public Register 1 or mediation-layer cases: the 2023 Roald Dahl edition controversy, 2025 federal data removals, CVE-2024-3094 involving XZ Utils, and the 2024 Google AI Overviews episode. Each tests a different boundary: alteration versus erasure, deletion versus functional erasure, openness versus exposure, public-function actionability versus restriction, and mediation versus recordhood.

Case	Framework test	Main distinction	Current result
Roald Dahl editions, 2023	Alteration vs erasure	Record resistance	Prior textual ecology resisted collapse
Federal data removals, 2025	Removal vs functional erasure	Object-class splitting	Cannot score as one object; requires protected-submaterial screen
CVE-2024-3094 / XZ Utils	Consequential public information	Low-thickness public-function openness	Remains open with dual-use boundary
Google AI Overviews, 2024	AI-mediated answer	Synthetic closure	Fails answerability when output hides weak source conditions

The Dahl case tests record resistance. A simple access frame might ask whether readers could still obtain some version of the books. A simple censorship frame might ask whether edits occurred. Shoreline asks whether the altered edition collapsed the prior textual ecology or whether the broader record ecology resisted erasure. The case shows high record resistance: prior editions remained available through multiple channels; comparison paths existed; changes were publicly detectable and contested; and publication of classic versions functioned as repair (Penguin 2023; Associated Press 2023; Guardian 2023).

The federal-data case tests heterogeneous public-data removal and repair. A simple access model might ask whether data were online or offline. Shoreline rejects that simplification. A public-data removal episode may include multiple object classes: archived or mirrored resources, dynamic-only resources, restored pages, altered reposts, and protected sub-material. The case therefore requires object-class splitting. It should not be scored as one object (Kaiser Family Foundation 2025; Freilich et al. 2025; Data Rescue Project 2025; Associated Press 2025).

The CVE-2024-3094 / XZ Utils case functions as a low-thickness public-function control. Shoreline must not route everything merely because information is consequential. Public vulnerability records and advisories serve their function by being findable, traceable, checkable, and actionable. The CVE case tests whether Shoreline can preserve openness where openness is proper, while still distinguishing public characterization from harmful operationalization (National Vulnerability Database 2024; National Vulnerability Database n.d.; Przymus and Durieux 2025).

The Google AI Overviews case tests synthetic closure. The object is not the underlying source field but the public AI-generated search overview as a mediation-layer output. The framework asks whether the output preserves the conditions required for answerability: source context, uncertainty, source type, and exploratory path. The case shows why misinformation language is not enough. The problem is not simply that an answer may be wrong; it is that a mediation layer may transform weak or inappropriate source material into a settled-looking answer (Google 2024; Business Insider 2024b; Ji et al. 2023).

A brief walkthrough shows how the operator packet connects to the theory. One operator-packet instrument, the Claim Answerability Gates, asks whether a claim is traceable, checkable, authorized within its source context, properly scoped, and corrigible. In this case, the object under review is the AI-generated overview itself. A synthetically closed answer may fail even if it contains fragments of real information, because the user cannot tell whether the answer comes from an authoritative source, satire, forum discussion, outdated material, or model inference. The Shoreline Assessment Protocol therefore does not ask only, “Was the

output true or false?” It asks whether the output preserved the conditions under which the answer could remain answerable. Where the output supplies action-bearing advice while hiding weak source context, the proper result is not ordinary answer status but route, warning, correction, or withdrawal.

The four cases also show why the framework uses maturity labels. The current operator packet has been internally calibrated, but external review is deferred. It is therefore improper to claim validation. The honest status is that the instruments have shown preliminary discriminating capacity on selected cases. They can distinguish alteration from erasure, heterogeneous deletion from functional erasure, public-function openness from over-routing, and AI mediation from recordhood. But they have not been tested by independent external scorers, validated, community-reviewed, or authorized for Register 2 application.

This limitation is not a weakness to hide. It is part of Shoreline’s discipline. Internal calibration does not become validation merely because it is useful. A prepared external scoring packet does not become external review merely because it exists. A route-only test does not become community authorization. The project’s credibility depends on marking those boundaries plainly.

IX. The Companion Operator Packet

The Shoreline Framework is accompanied by a companion operator packet. The paper makes the argument; the packet demonstrates how the argument can be translated into disciplined practice. The paper should be judged as a framework: whether it names a real problem, supplies useful distinctions, and clarifies failures in information movement. The packet should be judged as a draft instrument suite: whether it can generate consistent screens, scoring profiles, route decisions, and repair instructions.

The operator packet is not a validated tool, certification system, compliance instrument, or public standard. It is a companion artifact showing operational capacity. Its purpose is not to close the argument by measurement, but to make the argument testable.

The current operator packet contains six instruments. The Record Resistance Index profiles conditions that allow a record ecology to resist alteration, disappearance, substitution, or functional erasure. The Holding Legitimacy Profile asks whether an object is rightly held, governed, circulated, or used. The Relational Thickness / Detachability Scale asks how much meaning, authority, legitimacy, and movement depend on relation, role, place, context, or protocol. The Shoreline Assessment Protocol evaluates mediation layers such as AI systems, search engines, platforms, archives, dashboards, catalogs, and summaries. The Claim Answerability Gates ask whether a claim is traceable, checkable, authorized, contextualized, and corrigible. The Protocol Status Taxonomy distinguishes actual protocol position from apparent openness, exposure, and unknown/route conditions.

The operator packet also includes shared screens applied before scoring: source-path sufficiency, object definition, object-class splitting, register and thickness assessment, protected-submaterial screening, safety or dual-use triggers, and route/halt conditions. These screens are as important as the scoring tables. A score given to the wrong object is worse than no score. A score that should have routed is a false authorization. A score that averages away disagreement creates synthetic closure inside the framework itself (O’Connor and Joffe 2020; American Educational Research Association, American Psychological Association, and National Council on Measurement in Education 2014; Lincoln and Guba 1985).

The current packet has been developed through internal calibration on four cases. Those cases showed that the instruments can distinguish alteration from erasure, heterogeneous public-data loss from simple deletion, low-thickness public-function openness from over-routing, and AI mediation failure from ordinary factual

error. The calibration also produced repairs: object-class splitting and protected-submaterial screening for federal data; a dual-use boundary for public vulnerability records; a health and safety route trigger for AI summaries; and a maturity distinction between modeled internal divergence and actual external review.

These repairs are evidence of usefulness, not validation. A draft instrument becomes more credible when it can fail and be repaired, but failure-and-repair within one development environment is not independent validation. The operator packet is therefore best described as internally calibrated and externally score-ready, with external review deferred. That status must remain visible wherever the packet is discussed.

X. Rival Explanations and Why Shoreline Is Needed

Shoreline overlaps with archival theory, data governance, Indigenous data sovereignty, public records law, information ethics, misinformation studies, privacy, intellectual property, platform governance, and AI safety. It does not replace these fields. It names a problem that often appears between them.

The closest prior framework to several of these distinctions is Helen Nissenbaum’s contextual integrity, which holds that privacy is the appropriate flow of information and that a flow is appropriate when it conforms to the context-relative informational norms of a setting, specified by the data subject, sender, recipient, information type, and transmission principle (Nissenbaum 2004, 2010). Contextual integrity already captures much of what Shoreline means by “exposure is not openness” and “access is not authority”: information appearing in one context carries no license to flow into another, and public visibility does not settle appropriate use. Shoreline does not claim to have discovered that information movement is governed by context. Its contribution lies elsewhere, and the difference is precise. Contextual integrity is a theory of *flow*—whether a given transmission between actors is normatively appropriate—developed as an account of privacy. Shoreline is a theory of *status*—what an object is (record, evidence, knowledge, actionable) and the conditions under which it gains, holds, loses, regains, or is refused that standing. Three consequences follow that contextual integrity does not supply. First, the separation of record resistance from holding legitimacy is orthogonal to flow-appropriateness: whether an object survives alteration and whether it is rightly held are not questions about the norms of a transmission. Second, the account of mediation and synthetic closure concerns how an interface can manufacture the appearance of settled knowledge—a failure of epistemic form, not of who receives which information. Third, routing supplies an operational response to unresolved authority—halt outward movement and seek the authority path—where contextual integrity diagnoses violation without prescribing the system’s action under uncertainty. Shoreline is therefore best read as consolidating the context-sensitivity that contextual integrity established for privacy and extending it into a status-transition account that spans preservation, legitimacy, mediation, and refusal.

Access and open-data frameworks rightly emphasize availability, transparency, public oversight, reuse, and the democratic value of open information. Many public records should be accessible and many datasets should circulate. Shoreline rejects only the inference that access itself settles authority. Open-data analysis is strongest when the object is properly public and the question is how to preserve availability. It is weaker when public visibility is exposure, datasets contain protected sub-material, or restoration requires object-class analysis (OpenGovData 2007; Open Data Charter n.d.; Dawes 2010).

Public-records law deserves specific treatment, because it operates on a presumption of disclosure that Shoreline must not appear to override. It does not. In a Register 1 public-function system, routing is not an extra-legal withholding; it is the administrative compliance pause that disclosure regimes already contain—the interval in which an agency, rather than auto-publishing, reviews material for statutory exemptions such as

privacy, law enforcement, national security, or protections for cultural heritage. Shoreline supplies a trigger for that existing review, particularly for protected sub-material that no keyword flag would catch but a relational-thickness signal would. Object-class splitting is the framework’s name for what disclosure law already calls severability (5 U.S.C. § 552(b)): where a dataset contains protected sub-material, the lawful response is to isolate and redact that material so the public-function remainder can be released, not to withhold the whole.

Two boundaries keep this from cutting the wrong way. Shoreline creates no new authority to withhold; it adds no exemption that statute does not already provide. Where a disclosure mandate compels release of relationally thick material that has no statutory exemption, the framework does not invent an override—it names the result accurately as exposure without authorization, a loss of proper conditions that the law has compelled. And the framework’s own logic forecloses its misuse in the opposite direction: an agency that invokes “relational thickness” to withhold material the law requires disclosed is not exercising refusal but committing obstruction, blocking answerability under cover of a governance vocabulary it has no standing to apply (National Park Service 2024; Bureau of Indian Affairs n.d.).

Archival preservation frameworks understand provenance, custody, description, fixity, institutional memory, and record fragility. Shoreline depends on these insights but adds a second axis: preservation must be distinguished from legitimacy. A record may be well preserved and wrongfully held. A catalog may be accurate and still describe an object whose authority conditions remain unresolved (Upward 1996; Society of American Archivists n.d.-a; Local Contexts n.d.-a).

Data ethics and data governance frameworks explain harms of collection, use, privacy, consent, bias, surveillance, and accountability. Shoreline shares these concerns but pushes earlier. It asks how something becomes data, what was lost in that transition, whether the object was properly detached, whether the authority ecology permitted the movement, and whether protected sub-material should route rather than be scored as ordinary public information (Couldry and Mejias 2019; Roberts and Montoya 2023; Carroll et al. 2020).

Indigenous data sovereignty and Indigenous data governance frameworks are especially important to Shoreline’s interface architecture. Principles such as authority, control, consent, responsibility, relational accountability, and community governance are indispensable for understanding why access does not equal permission. Shoreline should not absorb these frameworks as if they were its own invention. Its role is to ask what outside systems must do when they encounter material whose conditions are not theirs to define (Carroll et al. 2020; First Nations Information Governance Centre n.d.; Local Contexts n.d.-a).

Misinformation and disinformation frameworks focus on falsehood, manipulation, deception, amplification, and public harm. These concerns are real, but not broad enough for Shoreline’s purposes. A mediated output can be harmful even when it is not simply false. A model may produce a correct statement while hiding uncertainty, source type, disagreement, or authority. Shoreline therefore treats truth and falsity as only part of the problem. The deeper question is whether a claim remains traceable, checkable, authorized, contextualized, and corrigible.

AI safety and platform governance frameworks explain output risk, model behavior, moderation, ranking, recommender effects, hallucination, alignment, and harmful content. Shoreline adds a specific account of mediation as status transition. AI systems do not merely generate text. They move material across status boundaries: source field to summary, retrieval to answer, probability to assertion, exposure to apparent openness, fluency to apparent knowledge. Synthetic closure names a failure ordinary AI safety language can miss (Ji et al. 2023; Zhang et al. 2023; Huang et al. 2023; Google 2024).

Each adjacent framework captures something real. Shoreline’s claim is that none of them, by itself, consistently tracks the full set of transitions at issue. Shoreline asks: what is the object, what status does it claim,

what conditions maintain that status, what transition is being attempted, who has authority over that transition, what protocol position governs movement, what would count as repair, and when must the system route or refuse?

XI. Limits and Validation Pathway

The Shoreline Framework must be limited by its own rule. Status does not transfer without conditions. A draft instrument suite does not become validated because it is coherent. Internal calibration does not become external review because it is useful. A prepared scoring packet does not become completed scoring because it exists. A route-only test does not become community authorization. If Shoreline is to be credible, it must refuse to counterfeit its own maturity.

The present paper makes a bounded claim. It introduces the framework, defines its vocabulary, explains the two-register architecture, develops the interface problem, and demonstrates preliminary discriminating capacity through four internally calibrated cases. It does not claim that the companion instruments have been externally validated, that they are ready for public compliance use, that they have been reviewed by relevant communities or authority ecologies, or that the paper speaks for any tribe, nation, family, community, or governance body.

This limitation is a condition of responsible presentation. Shoreline is most useful when it preserves the difference between proposal, prototype, internal calibration, external review, validation, and authority. Many frameworks move too quickly from insight to instrument, from instrument to standard, and from standard to institutional use. Shoreline should not repeat that pattern. Its current status is precise: the framework paper is draftable; the operator packet is internally calibrated through v0.3; the external scoring packet is prepared; external review is deferred; validation is not claimed (American Educational Research Association, American Psychological Association, and National Council on Measurement in Education 2014; Lincoln and Guba 1985; O'Connor and Joffe 2020; Mokkink et al. 2010).

The most important limit concerns Register 2 and protected-context material. Shoreline can say that apparent openness is not actual openness, that exposure should not be laundered into authorization, that unknown authority conditions require routing, and that outside systems owe interface obligations. But it cannot decide, from outside, the proper status of a particular community-held object, story, recording, practice, place-based relation, or governed knowledge item. That determination belongs to the relevant authority ecology.

One objection strikes at the framework's coherence rather than its scope: if assessing relational thickness required investigating an object's relations, then the very instrument meant to protect a boundary would cross it. The objection is answered by the asymmetry built into routing (Section IV). An outside system is never asked to establish thickness by investigation; it is asked to read detachability from signals it already holds, and to route whenever those signals are absent or unclear. Because uncertainty resolves toward governed rather than open, protection is the default that requires no inquiry, and inquiry is reserved for justifying movement. The framework therefore does not require the boundary-violating investigation it forbids. This is not a peripheral defense; it is a condition of the interface obligations' coherence. A framework that could protect relational knowledge only by first investigating it would be self-defeating, and Shoreline avoids that outcome by making non-movement, not investigation, the response to uncertainty.

This is why the current calibration cases are deliberately limited to public Register 1 or mediation-layer cases. Dahl, federal data, CVE-2024-3094, and AI Overviews test the framework without requiring unauthorized determination of protected material. They do not validate the framework for Register 2 applications, certify

the interface obligations, or authorize use on specific tribal, ceremonial, family-held, or community-governed materials.

The next validation condition is independent external scoring. Reviewers should include records or archival expertise; information systems, AI, platform, cybersecurity, or data-governance expertise; and legal, governance, or standing-sensitive expertise. Community archives, Indigenous data governance, or protocol-sensitive review may be useful only within a route-only or properly authorized scope. Reviewers should score independently and expose divergence rather than produce artificial agreement.

Divergence must be logged rather than averaged away. If reviewers define different objects, the object definition must be repaired. If they disagree about register or thickness, routing criteria must be repaired. If they disagree about source sufficiency, the source packet must be repaired. If they agree too easily because the form hides disagreement, the scoring template must be repaired. False consensus is another form of synthetic closure.

No stronger maturity label should be used until the conditions for that label are met. “Externally scored” requires actual external scoring. “Cross-operator checked” requires independent scoring and divergence analysis. “Validated” requires a larger case set, repeated scoring, repair after divergence, misuse analysis, and appropriate standing review for protected-context claims. “Community-reviewed” requires review by the relevant community or authority ecology, not generalized respect language. “Register 2 authorized” requires actual authority, not theoretical sympathy.

XII. Conclusion: The Shoreline as Boundary Discipline

The Shoreline Framework begins with a simple refusal: do not treat availability as authority. A visible object is not necessarily open. A stored object is not necessarily rightly held. A preserved object is not necessarily repaired. A fluent answer is not necessarily knowledge. A missing page is not necessarily erasure. A refusal is not necessarily obstruction. Information movement is legitimate only when the conditions of movement are maintained.

The shoreline is not a wall against knowledge. It is a boundary discipline for movement. Things can cross, and many things should. Public records should support accountability. Emergency information should move quickly. Public vulnerability records should remain traceable and actionable. Scientific findings should circulate. Archives should preserve. Search systems should help users find. AI systems may assist exploration. The framework is not anti-access, anti-technology, or anti-preservation. It is against status transfer without conditions.

The paper has developed Shoreline through a set of boundary distinctions that modern information systems routinely blur through scale, interface design, data aggregation, platform exposure, archival custody, algorithmic summary, and fluent generated language. Those distinctions matter because they prevent technical movement from becoming its own authorization.

Shoreline also distinguishes record ecology from relational knowledge governance. Record ecology depends on provenance, custody, indexing, fixity, accountability, contestability, and repair. Relational knowledge governance depends on authority conditions that may include role, place, family, community, ceremony, season, consent, refusal, and living transmission. The question is what happens at the interface, where systems built for one set of conditions encounter objects governed by another.

The calibration cases show why the framework is useful. The Dahl case shows that alteration is not neces-

sarily erasure when record resistance is strong. The federal-data case shows that public removal events must be split by object class and screened for protected sub-material. The CVE case shows that Shoreline can preserve low-thickness public-function openness rather than over-routing consequential information. The AI Overview case shows that mediation can fail by producing synthetic closure: a completed-looking answer without the conditions of answerability.

Together, these cases demonstrate preliminary discriminating capacity. They do not validate the instruments. The companion operator packet remains a draft tool and future validation object. It has been internally calibrated, repaired, and prepared for external scoring, but external review is deferred and Register 2 applications remain routed. The framework does not claim maturity it does not have.

Shoreline offers a vocabulary for slowing crossings down. Before movement, ask: what is the object? What status does it hold? What conditions maintain that status? What transition is being attempted? Who has authority? Is the object open, governed, restricted, held, exposed, or unresolved? Is preservation repairing the right thing? Is refusal valid? Is mediation preserving answerability or replacing it? Is the proper response access, repair, refusal, conditional movement, or route?

These questions do not solve every case in advance. They prevent the most dangerous shortcut: letting movement itself decide legitimacy. Information systems need more than access controls, preservation tools, metadata schemas, AI safeguards, or ethical principles considered separately. They need a way to see when status is being transferred without the conditions that make transfer legitimate. Shoreline provides that boundary discipline.

Status and Limitations

This is an unrefereed preprint. It has not been peer-reviewed, externally validated, or community-reviewed. The companion instrument suite is internally calibrated only, and its cross-operator divergence is modeled rather than produced by independent reviewers; it is a draft tool and a future validation object, not a validated or certified instrument. The framework theorizes the interface to relationally governed (Register 2) knowledge from outside and claims no authority over that knowledge: any determination there belongs to the relevant community, not to this work. Section XI sets out the validation pathway. The framework does not claim maturity it does not have.

Appendix A — Proposed Next Stages of Development

This appendix identifies the next development stages for the Shoreline Framework after the present manuscript draft. It is not a validation claim. It is a development map. The companion operator packet remains:

```
INTERNAL_CALIBRATION_COMPLETE_FOR_v0.3; EXTERNAL_REVIEW_READY_BUT_DEFERRED;  
NOT_VALIDATED
```

The stages below should be read as prospective conditions for further maturity, not as completed achievements.

Stage 1 — Manuscript Stabilization

Purpose: prepare the framework paper for serious external reading without overburdening it with instrument validation claims.

Work to complete:

1. select the target venue or target audience;
2. finalize citation style and reference formatting;
3. perform a line-level citation audit;
4. decide whether Figure 1 should appear as SVG, PNG, PDF, or simplified table;
5. trim or retain the calibration cases depending on venue length;
6. prepare a short “operator packet status” note for reviewers.

Gate condition: the paper can be circulated as a framework article without implying that the operator packet is validated.

Output: SUBMISSION_READY_FRAMEWORK_MANUSCRIPT_v1.0.

Stage 2 — External Scoring Preparation

Purpose: prepare the companion operator packet for independent review while preserving the distinction between framework publication and instrument validation.

Work to complete:

1. freeze the operator packet version to be reviewed;
2. freeze the calibration case packet;
3. create reviewer instructions;
4. define reviewer roles without simulating them internally;
5. prepare scoring forms;
6. prepare a divergence log template;
7. prepare a misuse and overclaim warning sheet.

Gate condition: external reviewers can apply the packet without needing private development context.

Output: EXTERNAL_SCORING_PACKET_READY_v1.0.

Stage 3 — Independent External Scoring

Purpose: test whether independent reviewers can use the operator packet consistently enough to expose useful agreement, disagreement, ambiguity, and repair needs.

Recommended reviewer capacities:

1. archival, records, or preservation expertise;
2. data governance, public information, AI, platform, or cybersecurity expertise;
3. legal, administrative, institutional, or standing-sensitive governance expertise;
4. optional community archives, Indigenous data governance, or protocol-sensitive review, but only for route-only or properly authorized materials.

Required rule: reviewers should score independently before discussion. Divergence should be logged rather than averaged away.

Gate condition: the project has actual external scoring results, not modeled review.

Output: EXTERNAL_SCORING_RESULTS_MEMO_v1.0.

Stage 4 — Divergence Repair and Instrument Revision

Purpose: convert external disagreement into better definitions, screens, route instructions, and case boundaries.

Repair questions:

1. Did reviewers define the same object?
2. Did they identify the same register?
3. Did they agree on relational thickness?
4. Did they agree on source-path sufficiency?
5. Did any case require route rather than scoring?
6. Did scoring categories force false consensus?
7. Did the packet produce any overconfident or unsafe result?

Gate condition: no instrument revision may claim improvement unless it is tied to a logged divergence, failure, ambiguity, or misuse risk.

Output: OPERATOR_PACKET_REPAIRED_AFTER_EXTERNAL_SCORING_v1.1.

Stage 5 — Register 2 / Standing Protocol Development

Purpose: develop a responsible pathway for protected-context or relationally governed materials without letting the framework appoint itself as authority.

Required boundary: Shoreline may route, halt, describe interface obligations, and identify standing-sensitive conditions. It may not determine the actual protocol position of Register 2 material without the relevant authority ecology.

Possible work products:

1. standing-screen protocol;
2. route-only review protocol;
3. community-invited review pathway;
4. refusal-acceptance template;
5. return/restriction/conditional-carriage decision record;
6. “do not score” category for unauthorized materials.

Gate condition: no Register 2 application may be described as authorized, reviewed, or validated unless the relevant authority ecology actually participates or grants authority.

Output: REGISTER_2_STANDING_AND_ROUTE_PROTOCOL_v0.1.

Stage 6 — Public-Facing Translation

Purpose: prepare accessible explanations without weakening the framework’s boundaries.

Possible formats:

1. one-page public explainer;

2. visual diagram set;
3. website page for the framework;
4. short slide deck;
5. operator-packet status page;
6. public note on why the tool is not yet validated.

Gate condition: public-facing materials must preserve the validation boundary and must not market the operator packet as a compliance tool.

Output: PUBLIC_EXPLAINER_AND_STATUS_PAGE_v1.0.

Stage 7 — Maintenance, Misuse Monitoring, and Version Governance

Purpose: keep the framework corrigible after circulation.

Maintenance duties:

1. maintain a version ledger;
2. record known weaknesses;
3. track misuse risks;
4. log overclaim attempts;
5. update source anchors;
6. distinguish framework updates from operator-packet updates;
7. preserve deprecated versions for comparison.

Misuse warnings to monitor:

1. institutions using “route” to delay accountability;
2. platforms using “unclear authority” to avoid removal or repair;
3. researchers using thickness scores as if they determine cultural authority;
4. agencies using refusal language to hide public records;
5. AI systems using citation or summary as substitute authorization;
6. readers treating internal calibration as validation.

Gate condition: the project remains answerable to its own maturity labels.

Output: SHORELINE_MAINTENANCE_AND_MISUSE_LOG_v1.0.

Stage 8 — Second Article or Technical Note

Purpose: separate the framework article from later instrument-validation work.

The present manuscript should be treated as the framework article. A later article or technical note can report the external scoring process, divergence repairs, and instrument revisions. This prevents the first paper from carrying a validation burden it does not yet meet.

Possible title:

Testing the Shoreline Operator Packet: External Scoring, Divergence, and Repair after a Status-Transition Framework for Information Movement

Gate condition: a second article or technical note should not be drafted as validation unless external scoring has actually occurred.

Output: OPERATOR_PACKET_VALIDATION_ARTICLE_OR_TECHNICAL_NOTE_v0.1.

Development Summary

Stage	Name	Primary output	Maturity effect
1	Manuscript stabilization	Framework manuscript v1.0	Paper becomes submission-ready
2	External scoring preparation	External scoring packet	Review becomes possible
3	Independent external scoring	Results memo	Review occurs
4	Divergence repair	Repaired operator packet	Instrument improves
5	Register 2 protocol	Standing/route protocol	Protected-context boundary strengthens
6	Public translation	Explainer/status page	Public comprehension improves
7	Maintenance governance	Misuse/version log	Long-term corrigibility improves
8	Second article/technical note	Validation article or note	Instrument claims can be evaluated separately

The guiding rule across all stages is simple: **do not let a future condition appear as a present status**. Shoreline’s development should model the same boundary discipline the framework asks of information systems.

References

- American Educational Research Association, American Psychological Association, and National Council on Measurement in Education. 2014. *Standards for Educational and Psychological Testing*. Washington, DC: American Educational Research Association. Open access edition. <https://www.aera.net/publications/books/standards-for-educational-psychological-testing-2014-edition>.
- Associated Press. 2023. “Penguin to Publish ‘Classic’ Roald Dahl Books after Backlash.” February 24, 2023. <https://apnews.com/article/books-and-literature-childrens-entertainment-roald-dahl-business-9770a7a3a2cb50cb1d53ca82d4b26070>.
- Associated Press. 2025. “Federal Health Agencies Restore Webpages and Datasets Following Judge’s Order.” February 2025. <https://apnews.com/article/ad1ec005261b31986ddcd3d1d17b249d>.
- Aturban, Mohamed, Sawood Alam, Michael L. Nelson, and Michele C. Weigle. 2019. “Archive Assisted Archival Fixity Verification Framework.” arXiv. <https://arxiv.org/abs/1905.12565>.
- Bergner, Raymond M., and James R. Holmes. 2000. “Self-Concepts and Self-Concept Change: A Status Dynamic Approach.” *Psychotherapy: Theory, Research, Practice, Training* 37 (1): 36–44.
- Bureau of Indian Affairs. n.d. “Native American Graves Protection and Repatriation Act.” Accessed June 15, 2026. <https://www.bia.gov/service/nagpra>.
- Business Insider. 2024a. “Sometimes to Save the Internet, You Must Eat Glue.” May 2024. <https://www.businessinsider.com/google-search-ai-overview-answers-pizza-glue-fix-2024-5>.

Business Insider. 2024b. “Google Seems to Be Manually Removing Some Strange AI Search Answers.” May 2024. <https://www.businessinsider.com/google-removing-strange-ai-overviews-search-answers-pizza-glue-2024-5>.

Carroll, Stephanie Russo, Edit Herczog, Maui Hudson, Keith Russell, and Shelley Stall. 2020. “The CARE Principles for Indigenous Data Governance.” *Data Science Journal* 19 (1): 43. <https://doi.org/10.5334/dsj-2020-043>.

Colavizza, Giovanni, Tobias Blanke, Charles Jeurgens, and Julia Noordegraaf. 2021. “Archives and AI: An Overview of Current Debates and Future Perspectives.” arXiv. <https://arxiv.org/abs/2105.01117>.

Couldry, Nick, and Ulises A. Mejias. 2019. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA: Stanford University Press. <https://www.sup.org/books/sociology/costs-connection>.

Data Rescue Project. 2025. “Resources.” Accessed June 15, 2026. <https://www.datarescueproject.org/resources/>.

Dawes, Sharon S. 2010. “Stewardship and Usefulness: Policy Principles for Information-Based Transparency.” *Government Information Quarterly* 27 (4): 377–383. https://www.ctg.albany.edu/media/pubs/pdfs/giq_2010_stewardship.pdf.

Digital Preservation Coalition. n.d. “Fixity and Checksums.” *Digital Preservation Handbook*. Accessed June 15, 2026. <https://www.dpconline.org/handbook/technical-solutions-and-tools/fixity-and-checksums>.

First Nations Information Governance Centre. n.d. “The First Nations Principles of OCAP®.” Accessed June 15, 2026. <https://fnigc.ca/ocap-training/>.

Freedom of Information Act. n.d. 5 U.S.C. § 552 (1966, as amended). U.S. Department of Justice, Office of Information Policy. Accessed June 15, 2026. <https://www.justice.gov/oip/freedom-information-act-5-usc-552>.

Freilich, Janet W., et al. 2025. “Disappearing Data at the U.S. Federal Government.” Boston University School of Law Faculty Scholarship. https://scholarship.law.bu.edu/faculty_scholarship/4044/.

Freilich, Janet W., and Aaron S. Kesselheim. 2025. “Data Manipulation within the US Federal Government.” *The Lancet* 406 (10500): 227–228. [https://doi.org/10.1016/S0140-6736\(25\)01249-8](https://doi.org/10.1016/S0140-6736(25)01249-8).

Frings-Hessami, Viviane. 2021. “The Societal Embeddedness of Records: Teaching the Meaning of the Fourth Dimension of the Records Continuum Model in Different Cultural Contexts.” *Archival Science* 21: 139–154. <https://doi.org/10.1007/s10502-020-09349-6>.

Google. 2024. “AI Overviews: About Last Week.” Google Keyword Blog, May 30, 2024. <https://blog.google/products-and-platforms/products/search/ai-overviews-update-may-2024/>.

Guardian. 2023. “Roald Dahl Publisher Announces Unaltered 16-Book ‘Classics Collection.’” February 24, 2023. <https://www.theguardian.com/books/2023/feb/24/roald-dahl-publisher-announces-unaltered-16-book-classics-collection>.

Huang, Lei, Weijiang Yu, Weitao Ma, et al. 2023. “A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions.” arXiv. <https://arxiv.org/abs/2311.05232>.

Ji, Ziwei, Nayeon Lee, Rita Frieske, et al. 2023. “Survey of Hallucination in Natural Language Generation.” *ACM Computing Surveys* 55 (12): Article 248. <https://doi.org/10.1145/3571730>.

Kaiser Family Foundation. 2025. “A Look at Federal Health Data Taken Offline.” February 2, 2025. <https://www.kff.org/health-costs/a-look-at-federal-health-data-taken-offline/>.

Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks.” In *Advances in Neural Information Processing Systems* 33, 9459–9474. Curran Associates. <https://arxiv.org/abs/2005.11401>.

Lincoln, Yvonna S., and Egon G. Guba. 1985. *Naturalistic Inquiry*. Beverly Hills, CA: Sage.

Local Contexts. n.d.-a. “Traditional Knowledge Labels.” Accessed June 15, 2026. <https://localcontexts.org/labels/traditional-knowledge-labels/>.

Local Contexts. n.d.-b. “Local Contexts: Grounding Indigenous Rights.” Accessed June 15, 2026. <https://localcontexts.org/>.

Mokkink, Lidwine B., Caroline B. Terwee, Donald L. Patrick, et al. 2010. “The COSMIN Checklist for Assessing the Methodological Quality of Studies on Measurement Properties of Health Status Measurement Instruments: An International Delphi Study.” *Quality of Life Research* 19: 539–549. <https://doi.org/10.1007/s11136-010-9606-8>.

National Digital Stewardship Alliance. 2014. “Checking Your Digital Content: What Is Fixity, and When Should I Be Checking It?” Library of Congress. <https://www.digitalpreservation.gov/documents/NDSA-Fixity-Guidance-Report-final100214.pdf>.

National Park Service. 2024. “Native American Graves Protection and Repatriation Act.” April 29, 2024. <https://www.nps.gov/subjects/nagpra/>.

National Vulnerability Database. 2024. “CVE-2024-3094 Detail.” <https://nvd.nist.gov/vuln/detail/cve-2024-3094>.

National Vulnerability Database. n.d. “NVD Home.” Accessed June 15, 2026. <https://nvd.nist.gov/>.

O’Connor, Cliodhna, and Helene Joffe. 2020. “Intercoder Reliability in Qualitative Research: Debates and Practical Guidelines.” *International Journal of Qualitative Methods* 19: 1–13. <https://doi.org/10.1177/1609406919899220>.

Nissenbaum, Helen. 2004. “Privacy as Contextual Integrity.” *Washington Law Review* 79 (1): 119–157.

Nissenbaum, Helen. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford University Press.

Open Data Charter. n.d. “Open Data Charter Principles.” Accessed June 15, 2026. <https://opendatacharter.org/principles/>.

OpenGovData. 2007. “The 8 Principles of Open Government Data.” <https://opengovdata.org/>.

Penguin. 2023. “Puffin Announces The Roald Dahl Classic Collection to Keep Author’s Classic Texts in Print.” February 24, 2023. <https://www.penguin.co.uk/about/company-articles/puffin-announces-the-roald-dahl-classic-collection-to-keep-authors-classic-texts-in-print>.

Przymus, Piotr, and Thomas Durieux. 2025. “Wolves in the Repository: A Software Engineering Analysis of the XZ Utils Supply Chain Attack.” arXiv. <https://arxiv.org/abs/2504.17473>.

Roberts, Jennafer Shae, and Laura N. Montoya. 2023. “In Consideration of Indigenous Data Sovereignty: Data Mining as a Colonial Practice.” arXiv. <https://arxiv.org/abs/2309.10215>.

Society of American Archivists. n.d.-a. “Provenance.” *Dictionary of Archives Terminology*. Accessed June 15, 2026. <https://dictionary.archivists.org/entry/provenance.html>.

Society of American Archivists. n.d.-b. “Records Continuum.” *Dictionary of Archives Terminology*. Accessed June 15, 2026. <https://dictionary.archivists.org/entry/records-continuum.html>.

State Library Resource Center. 2025. “Searching for Federal Data.” May 9, 2025. <https://www.slrc.info/research/guides/federal-data>.

Tuck, Eve, and K. Wayne Yang. 2014. “R-Words: Refusing Research.” In *Humanizing Research: Decolonizing Qualitative Inquiry with Youth and Communities*, edited by Django Paris and Maisha T. Winn, 223–248. Thousand Oaks, CA: SAGE.

Upward, Frank. 1996. “Structuring the Records Continuum — Part One: Postcustodial Principles and Properties.” *Archives and Manuscripts* 24 (2): 268–285. <https://publications.archivists.org.au/index.php/asa/article/view/8583>.

Wilkinson, Mark D., Michel Dumontier, IJsbrand Jan Aalbersberg, et al. 2016. “The FAIR Guiding Principles for Scientific Data Management and Stewardship.” *Scientific Data* 3: 160018. <https://doi.org/10.1038/sdata.2016.18>.

Zhang, Yue, Yafu Li, Leyang Cui, et al. 2023. “Siren’s Song in the AI Ocean: A Survey on Hallucination in Large Language Models.” arXiv. <https://arxiv.org/abs/2309.01219>.