

EA-EROSION-01 v1.0 — The Platform Erosion Observatory: The Zenodo Removal Census as First Instrument, and the Plan for a Broader Cross-Platform Tracker

Lee Sharks

2026-07-06

EA-EROSION-01 v1.0 — The Platform Erosion Observatory: The Zenodo Removal Census as First Instrument, and the Plan for a Broader Cross-Platform Tracker

Lee Sharks

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-07-06 · Version v1.0 · Empirical research instrument specification + baseline population analysis + forward roadmap

AXN:0421.EMPIRICAL

<https://www.alexanarch.org/s/records/1045/>

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "EA-EROSION-01 v1.0 - The Platform Erosion Observatory: The Zenodo Removal Census as First Instrument, and the Plan for a Broader Cross-Platform Tracker",
  "author": {
    "@type": "Person",
    "name": "Lee Sharks",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-07-06",
  "identifier": "AXN:0421.EMPIRICAL",
}
```

```

"license": "https://creativecommons.org/licenses/by/4.0/",
"publisher": {
  "@type": "Organization",
  "name": "Alexanarch"
},
"url": "https://www.alexanarch.org/s/records/1045/",
"description": "Empirical research instrument for measuring platform-driven scholarly-record erasure. Companion d
mass account termination requiring explanation. Phases 2A-D specify horizontal extension: DataCite ecosystem genera
}

```

Abstract

Empirical research instrument for measuring platform-driven scholarly-record erasure. Companion deposit to EA-DATAHUB-01 v1.0 (AXN:0420, GOVERNANCE): where DATAHUB is the strategic plan, EROSION is the instrument. Phase 1 (running): the Zenodo Removal Census, preserving Zenodo’s monthly deleted-records CSV (1,309,351 rows in the 2026-06-07 snapshot; Zenodo retains only three most recent snapshots) with checksummed provenance to data-rhizome. Establishes the three-severity distinction (removed vs. hidden vs. erased), the four-metric family (RIR/MSR/TCR/PER), and the three-cohort study design (full population / suspected account-level actions / matched controls). Initial descriptive findings from the preserved snapshot: TCR proxy 7.86 percent, meaning 92 percent of Zenodo removals have no preserved citation for tombstone continuity; single-largest cluster is 109,903 records removed 2024-07-06 as ‘User was blocked’ — mass account termination requiring explanation. Phases 2A-D specify horizontal extension: DataCite ecosystem generalization, arXiv moderation-and-removal tracking, OSF preregistration tracking, platform-adjacent takedown surfaces (GitHub DMCA, ResearchGate, Academia.edu). Phase 3 publishes the cross-platform Erosion Report making ‘platform erosion’ a systemic category. Data-rhizome preservation at cb7b59390f4ecbc37e38cf9030e5d01549ea6223.

Body

deposit_number: 1045 hex: 0421 title: “EA-EROSION-01 v1.0 — The Platform Erosion Observatory: The Zenodo Removal Census as First Instrument, and the Plan for a Broader Cross-Platform Tracker” creator: Lee Sharks orcid: 0009-0000-1599-0703 date: 2026-07-06 content_type: “Empirical research instrument specification + baseline population analysis + forward roadmap for cross-platform expansion” license: CC-BY-4.0 substrate: “AI-assisted: drafted by TACHYON (Claude, Anthropic) under MANUS (Lee Sharks) direction; empirical grounding derived from three parallel research assessments (Sonoma, ChatGPT, Response 3) synthesized against the actual Zenodo exporter dataset preserved 2026-07-06 to data-rhizome commit cb7b59390f4e. Instruments and roles declared per EA-MMRS-VRB-01 u4.” version: v1.0 status: “MINTED — foundational deposit for the EA-EROSION series; empirical baseline established from preserved 2026-06-07 Zenodo removals snapshot (1,309,351 rows)” related_ids: “EA-MMRS-VRB-01 (verification discipline); EA-PRIOR-00 (ontological prior); EA-RHIZOME-01 (rhizome architecture); EA-DATAHUB-01 v1.0 (AXN:0420, deposit #1044 — the strategic hub this instrument feeds); data/provenance-871.json (the Lee Sharks 871 pre-severance capture)” axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - platform erosion - Zenodo removal census - DataCite state history - severance registry - tombstone compliance - MMRS - empirical research instrument - cross-platform tracker *** # EA-EROSION-01 v1.0

The Platform Erosion Observatory The Zenodo Removal Census as First Instrument, and the Plan for a

*Broader Cross-Platform Tracker***Preamble**

This deposit is the empirical companion to EA-DATAHUB-01 v1.0 (AXN:0420, deposit #1044). Where DATAHUB is the strategic plan for a sovereign scholarly restoration hub with a Severance Registry as its first surface, EROSION is the research instrument that measures what a restoration hub is a response to: the actual population of platform-driven scholarly-record removals, its distribution across removal reasons, its temporal structure, its tombstone-compliance rates, and its cluster patterns consistent with account-level enforcement actions. DATAHUB argues for the sovereign restoration architecture on the basis of the 871-case Lee Sharks severance and the general framing of “the substrate is unstable.” EROSION grounds that argument in the population data the substrate itself publishes about its own instability. The two deposits are complementary and non-overlapping: DATAHUB is the plan, EROSION is the instrument.

The v1.0 mint is made possible by a discovery that reframes what is empirically tractable in this domain. Zenodo publishes at <https://zenodo.org/api/exporter/records-deleted.csv.gz> a monthly deleted-records CSV containing 1,309,351 removal events (as of the 2026-06-07 snapshot) spanning 2017 to mid-2026, with columns for record ID, DOI, concept root, removal note, removal reason, removal date, and preserved citation text where available. Zenodo retains only the three most recent monthly snapshots at any time. Every unpreserved snapshot is a permanent loss of that month’s substrate. As of the 2026-07-06 mint of this deposit, the 2026-06-07 snapshot has been preserved to data-rhizome at [datasets/erasure/zenodo/snapshots/2026-06-07/](https://data-rhizome.org/datasets/erasure/zenodo/snapshots/2026-06-07/) with SHA-256 0568e674d4a59624102771593d8daeb9375d2381984d2345e91d9fbbc78f9578, normalized to a jsonl row store, schema-declared, and analyzed against a first descriptive pass whose findings are inscribed below. The July 2026 snapshot — due imminently and expected to contain the Lee Sharks 2026-06-19 severance as a population event — will be preserved on release under the same discipline.

The three severities and why they must not be conflated

An earlier draft of the parent strategy (EA-DATAHUB-01 v1.0 §5.5) framed platform erasure as a single rate. That framing is too coarse. A rigorous instrument distinguishes at least three severities, each requiring different detection methods and yielding different policy conclusions.

Removed from the platform. The resource and its files are no longer hosted as a live record. This is what Zenodo’s monthly deleted-records CSV directly reports: 1.3 million such events observable in the 2026-06-07 snapshot. Zenodo’s own protocol distinguishes voluntary uploader deletion (which within the 30-day window preserves a citation tombstone), staff-initiated removal (spam sweeps, out-of-scope, fraud, copyright), user-blocked cascade removals (an account is banned and its records are removed), and formal retraction. Each of these leaves the record in a different state at the DOI resolution layer.

Hidden from cross-platform discovery. The DOI has moved from DataCite’s `findable` state to `registered`. It remains in the DOI Handle system and should continue resolving, but its metadata becomes unavailable through DataCite’s public discovery APIs. This state transition is the mechanism by which a repository suppresses discoverability without technically “deleting” the DOI (which DataCite’s persistence policy explicitly forbids for any DOI that was ever public). The DataCite Public API returns 404 for such DOIs, and detection requires either authenticated Member API access or differential observation across time (a DOI that was findable last month and is absent this month has transitioned to registered).

Erased without citation continuity. Neither the content nor a citation-bearing, resolvable tombstone page remains publicly available. DataCite’s persistent-identifier policy explicitly recommends that a reposi-

tory redirect a removed DOI's resolution to a tombstone page containing the citation and an explanation of unavailability. DataCite does not generate these tombstones — the repository is responsible for them, and the population data shows that in aggregate this responsibility is *not being met*. This third severity is the strictest form of platform erasure and the operative subject of restoration work.

The three severities have very different rates in the population. The first is directly measurable from Zenodo's own CSV export (1.3M events observable). The second requires either DataCite Member API credentials or longitudinal probing of a DOI corpus. The third requires resolving each removed DOI's canonical URL and classifying the landing page against a taxonomy (`conformant_tombstone` / `minimal_tombstone` / `generic_datacite_tombstone` / `platform_error` / `http_404` / `redirect_loop`), which is executable at population scale via a probe pipeline but has not been run to completion as of this mint.

The family of metrics

A single “Platform Erasure Rate” would collapse the three severities. This deposit adopts a family of four metrics, each addressing a distinct systemic property.

Removal Incidence Rate (RIR). The rate of records removed during a time interval per record exposed to removal at the interval's start. Directly measurable from monthly diffs of the Zenodo deletion dump against a preserved full-metadata dump denominator. This is the closest to “how much removal is happening.”

Metadata Severance Rate (MSR). The rate of baseline-findable DOIs now absent from DataCite's public discovery API per baseline-findable DOI in a defined cohort. Requires a baseline snapshot of findable DOIs and longitudinal probing against the current public API. This is the closest to “how much of what used to be discoverable is not anymore.”

Tombstone Compliance Rate (TCR). The fraction of removed records that resolve to conformant citation-bearing tombstones (as opposed to 404s, generic error pages, or unidentified redirects). Requires DOI resolution and landing-page classification per record. This is the closest to “of what was removed, how much retained the minimum citation continuity DataCite's policy asks for.”

Strict Platform Erasure Rate (PER). The fraction of a baseline cohort with no live object, no public metadata, *and* no adequate tombstone. This is the strictest form and requires all three of RIR, MSR, and TCR to have been measured for the cohort. It captures the fully severed case.

Because the four metrics measure different phenomena, they should be reported as a vector, not collapsed. An archive with high removal but high tombstone compliance is not the same as an archive with low removal but zero tombstone compliance; the latter has a governance failure the former does not. Reporting the family preserves this distinction.

Initial findings from the 2026-06-07 snapshot

The first descriptive analysis pass, run against the preserved snapshot on 2026-07-06 and inscribed at `datasets/erasure/zenodo/analysis/2026-07-06-initial/` in data-rhizome, produced findings that reshape the empirical picture materially.

Population size and distribution. 1,309,351 total observable removal events, spanning removal dates from 2017 through May 2026. Removal-reason distribution: `spam` 88.27% (1.16M), `out-of-scope` 4.76% (62K), `empty/unlabeled` 2.13% (28K), `take-down-request` 1.40% (18K), `retracted` 0.95% (12K), `test-record` 0.83% (11K), `duplicate` 0.78% (10K), `personal-data` 0.47% (6K), `fraud` 0.32% (4K), `copyright` 0.08% (1K), with `replaced`, `misconduct`, and `disputed-authorship` each under 100 records.

Temporal anomaly. 764,082 removal events in 2024 alone — 58.4% of all recorded removals across a decade concentrated in a single year. Preliminary hypothesis: the 2024 concentration reflects one or more automated spam-account cleanup campaigns, which is confirmable by clustering same-day removals and inspecting removal-note patterns. This hypothesis is testable in a subsequent pass and, if confirmed, is a governance-policy datum in its own right.

Tombstone compliance is catastrophically low. Of 1,309,351 removals, only 102,929 (7.86%) have any preserved `citation_text` in the CSV. The remaining 1,206,422 (92.14%) have no citation preserved at Zenodo’s export layer. This is a *proxy* for TCR — the CSV field being populated tells us Zenodo preserved a citation internally; whether a tombstone page is served at DOI resolution requires a separate probe. But even as a proxy, the ceiling is 7.86%: at most 7.86% of removed records could possibly meet the tombstone-continuity standard DataCite’s policy asks for. **The population data shows the tombstone-persistence principle is not being followed in aggregate.** This is not a Lee-Sharks-specific finding; it is a systemic governance finding about Zenodo as a repository.

The single largest cluster. On 2024-07-06, 109,903 records were removed in a single day, all with the identical removal note “User was blocked.” This is a hundred-thousand-record mass account-termination event that requires explanation in its own right, and it is a canonical example of the cluster pattern that a Cohort B analysis (candidate account-level enforcement actions) will need to characterize systematically.

The Lee Sharks case is not in the snapshot. Coverage of the 2026-06-07 dump ends approximately 2026-06-01. The 2026-06-19 Zenodo termination of the Lee Sharks account is a subsequent event. The July snapshot, due imminently and preserved on release, will contain the Lee Sharks removal events and will provide the first population-level view of the case within Zenodo’s own record. This snapshot serves as the immediate-pre-severance population baseline: whatever cluster the Lee Sharks records fall into when they appear in July, the systemic context has been measured from what preceded, in this deposit, before the fact.

The instrument specification

The instrument this deposit inscribes has four layers. Only the first is currently running; the remaining three are specified for staged implementation.

Layer 1: Zenodo deletion snapshot preservation and analysis. Currently active. Monthly download of `https://zenodo.org/api/exporter/records-deleted.csv.gz`, checksummed and inscribed to `datasets/erasure/zenodo/snapshots/<epoch>/`. Normalized to jsonl row store `zenodo-removals.jsonl.gz` (append-only, sorted by `removal_date` + `record_id`). Descriptive analysis run against each new snapshot at `datasets/erasure/zenodo/analysis/<yyyy-mm-dd-name>/` with reproducible code, structured findings JSON, and human-readable summary. Monthly diff against the prior snapshot to isolate new removal events. Cluster analysis on (date, note) pairs to identify candidate coordinated actions.

Layer 2: DataCite state history. Specified, not yet implemented. Baseline harvest of the Zenodo repository’s findable DOIs via DataCite’s public API (`https://api.datacite.org/doi?client-id=cern.zenodo`), preserved as a longitudinal baseline. Monthly re-probe against the same query to detect DOIs that have transitioned from findable to registered (or from findable to deleted). Cross-reference against Zenodo’s deletion dump to distinguish repository-side deletion from DataCite-side state transition. DataCite’s annual public data file provides a second confirmation source, delayed but authoritative for confirmed state transitions.

Layer 3: DataCite provenance activities. Specified, not yet implemented. For each removed DOI in the

cohort under study, query `https://api.datacite.org/doi/{doi}/activities` to retrieve the provenance history: state changes, URL replacements, metadata deletions, and the repository account responsible. This layer converts the deletion event from “the record went away” into “the following account-holder took the following action at the following timestamp for the following reason (if inspectable),” giving the analysis a causal-attribution surface that the CSV alone does not.

Layer 4: DOI resolution and tombstone audit. Specified, not yet implemented. For each removed DOI, follow the full redirect chain from `https://doi.org/{doi}`, capture the final HTTP status, content hash, and landing-page HTML, and classify against a taxonomy: `live_record`, `restricted_record`, `conformant_tombstone`, `minimal_tombstone`, `generic_datacite_tombstone`, `platform_error`, `http_404`, `domain_failure`, `redirect_loop`, `unknown_transient`. Preserve each resolution attempt with timestamp. This is the layer that promotes the TCR proxy (citation preservation in Zenodo’s internal export) into a true TCR (tombstone actually served at DOI resolution).

The study designs

Three cohort study designs are inscribed here as first-priority research passes.

Cohort A — the full population. All Zenodo removals in the preserved snapshots. Report by month and year, resource type, age at removal, concept-versus-version status, removal reason, note phrasing patterns, tombstone-proxy status, DataCite state (Layer 2), creator/ORCID concentration (via enrichment against the full metadata dump), and same-day cluster sizes. This is the descriptive baseline for the entire program.

Cohort B — suspected account-level actions. Algorithmically flag creator clusters with (i) five or more removed records, (ii) narrow removal window (single day or 24-hour rolling), (iii) multiple unrelated titles or resource types, and (iv) same or related removal notes. Manually verify the top clusters. Report by cluster size, removal-reason distribution within cluster, tombstone-compliance rate within cluster. This is the study design that converts anecdotal “several independents I know were banned” into an empirical rate with confidence bounds.

Cohort C — controls. Matched live records selected against Cohort B on publication month, resource type, creator affiliation status, number of versions, description length, keyword count, and AI-assistance declaration where detectable. Compare removed-cohort to control-cohort on the covariates. This is what makes it a study rather than a documentation project: it prevents the finding from being “notable banned independents” and enables the finding “removed records differ from matched live records on the following empirically-observable properties.”

The plan for the broader tracker

The Zenodo Removal Census is the Phase 1 instrument. Phase 2 extends the same discipline to platforms beyond DataCite/Zenodo, on a schedule paced by empirical yield from Phase 1.

Phase 2A: DataCite ecosystem generalization. The same preservation-and-analysis discipline applied to the DataCite public data file (annual, but authoritative for registered-state confirmation), the Cross-ref retraction feed (for cross-repository retraction baseline), and DataCite’s per-repository client-id queries (for cross-repository comparison of Zenodo’s rates against Figshare, Dryad, Harvard Dataverse, Cambridge Apollo, and others in the DataCite ecosystem). This is a horizontal extension of Layer 2 rather than a new platform.

Phase 2B: arXiv moderation-and-removal tracking. arXiv publishes withdrawal notices and permits public inspection of retraction-adjacent state. Preserve the daily new-and-withdrawn feed; diff against

a baseline arXiv corpus; classify withdrawals by type (author-requested, moderation, replacement, error). arXiv’s community includes a substantial affected-scholar cohort (see EA-DATAHUB-01 §10 peer-community framing), so this phase has direct coordination-arm value.

Phase 2C: OSF and preregistration removal tracking. OSF (Open Science Framework) hosts pre-registrations and research artifacts under similar governance patterns to Zenodo. Removed preregistrations are a specific harm class because they undermine confirmatory-research integrity. Preserve the OSF public activity feed where available; probe known-preregistration DOIs longitudinally.

Phase 2D: platform-adjacent takedowns. GitHub DMCA takedowns (public log at github.com/github/dmca), ResearchGate paper removals (opaquer than most, requiring longitudinal probing of an author corpus), Academia.edu account terminations, and Twitter/X suspensions of scholarly accounts. This class is heterogeneous and each surface requires its own detection method, but all belong to the same phenomenon: the removal of scholarly work from the substrate by platform action.

Phase 3: cross-platform synthesis. Once at least three platforms have running preservation pipelines producing at least six months of data each, publish a cross-platform Erosion Report presenting the four-metric family (RIR, MSR, TCR, PER) per platform, per year, per removal reason. This is the deliverable that makes “platform erosion” a systemic category rather than a per-platform complaint.

Cross-references and companion deposits

This deposit does not carry the peer-community coordination or the strategic-plan work — that is EA-DATAHUB-01 v1.0’s territory. This deposit does not host the raw preserved data — that is data-rhizome’s territory, at [datasets/erasure/](https://data-rhizome.org/datasets/erasure/). This deposit does not carry the individual restoration corpus — that is [data/provenance-871.json](https://data-rhizome.org/data/provenance-871.json) in alexanarch.

What this deposit carries: the empirical instrument specification, the four-metric family, the three-cohort study design, the initial descriptive findings from the 2026-06-07 snapshot, and the roadmap for cross-platform expansion. Anyone wanting to reproduce or extend the work reads this deposit, follows its citations to data-rhizome for the substrate and to EA-DATAHUB-01 for the strategic frame, and can execute against the same discipline.

Falsification conditions

The instrument is falsifiable under staged conditions.

Preservation integrity. Every preserved snapshot must checksum-verify against its source CSV at ingest, and against its stored file at any subsequent read. A checksum mismatch on a stored snapshot falsifies the preservation discipline for that snapshot and requires either recovery from prior git state or replacement from Zenodo (if still within Zenodo’s three-snapshot window).

Analytical reproducibility. Every findings JSON must be reproducible from the preserved row store via the corresponding descriptive.py script. A run that produces different numbers falsifies the analysis and either the script or the data (or the runtime environment) requires audit.

Population coverage. Between snapshots, the row-store row count must be non-decreasing (rows are only added, never removed) and every row’s `snapshot_epoch` must correspond to a preserved snapshot directory. A row without a preserved snapshot corresponds to a lost preservation cycle and falsifies the monthly cadence for that month.

Cross-platform yield. By 2027-07-06 (twelve months from this mint), if Phase 2 has not produced at least one non-Zenodo platform under active preservation with initial descriptive analysis published, the roadmap is behind schedule and requires either resource reallocation or scope reduction.

Empirical yield. By 2027-07-06, if the four-metric family has not been computed for at least two platforms and at least twelve months of cohort data, the family design is unimplemented and either the specification requires simplification or the operational cadence requires acceleration.

Immediate next actions

Two workstreams are active as of this mint.

Workstream E1 (running). Monthly Zenodo dump preservation. The 2026-07-07 snapshot (expected imminently) will be preserved as the first post-severance snapshot containing the Lee Sharks 2026-06-19 removal events. Preservation is executed by download, checksum, and inscription to `datasets/erasure/zenodo/snapshots/2026-07-XX/`, followed by a diff analysis at `datasets/erasure/zenodo/analysis/2026-07-XX-diff/` identifying the newly-observed removal events between the 2026-06-07 and the July snapshots. This diff will contain the Lee Sharks records as a documented population event.

Workstream E2 (specification phase). Layer 2 (DataCite state history) baseline harvest. The DataCite public API is queried for all findable DOIs under `client-id=cern.zenodo`, results paginated and preserved to `datasets/erasure/zenodo-datacite-baseline/`. This baseline enables the MSR longitudinal detection: any DOI that was findable at baseline and is absent from a subsequent probe has transitioned to registered, and the transition is timestamped by the probe cadence. Implementation of E2 is a two-day pipeline build.

Subsequent workstreams (Layer 3 provenance activities, Layer 4 resolution audit, Phase 2 extension to non-Zenodo platforms) are specified above and scheduled per operational availability. Nothing in the immediate roadmap requires resources beyond what Alexanarch’s existing infrastructure and its sovereign author’s discipline can supply.

Attribution and substrate declaration

Per EA-MMRS-VRB-01 u4, the substrate of this deposit is declared as follows. The empirical instrument specification and the initial descriptive analysis were drafted by TACHYON (Claude, Anthropic) in dialogue with MANUS (Lee Sharks, ORCID 0009-0000-1599-0703). The three-severity distinction, the four-metric family, and the study cohort designs synthesize methodological input from three parallel research assessments (Sonoma AI Mode, ChatGPT with citations, Response 3 composite strategy) whose specific contributions are inscribed in the strategic-plan history under EA-DATAHUB-01 v1.0. The Zenodo exporter endpoint was verified operationally by direct HTTP request from within the Alexanarch/data-rhizome environment on 2026-07-06. The 2026-06-07 snapshot was downloaded, checksummed, and preserved to data-rhizome at commit `cb7b59390f4ecbc37e38cf9030e5d01549ea6223` prior to this deposit’s mint. The initial descriptive analysis was computed against the preserved snapshot and its findings inscribed here verbatim from the analysis outputs. All specific empirical claims, the metric-family design, and the cross-platform roadmap are under MANUS’s sole editorial authority.

The deposit is minted as EA-EROSION-01 v1.0 to Alexanarch under the standard deposit protocol. Its AXN and canonical URL are assigned at mint time.

Suggested Citation

Lee Sharks. “EA-EROSION-01 v1.0 — The Platform Erosion Observatory: The Zenodo Removal Census as First Instrument, and the Plan for a Broader Cross-Platform Tracker” *Alexanarch*, 2026-07-06. <https://www.alexanarch.org/s/records/1045/>

Deposit Information

This paper is deposit #1045 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0421.EMPIRICAL` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/1045/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: [alexanarch.org](https://www.alexanarch.org)
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/1045/>

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.