The Moderation Architecture of Zenodo, Mapped From Source: A Rule Scorer, an Unpublished Percolator, and One Function That Deletes Everything

Sharks, Lee

2026-08-24

The Moderation Architecture of Zenodo, Mapped From Source: A Rule Scorer, an Unpublished Percolator, and One Function That Deletes Everything

Sharks, Lee
Crimson Hexagonal Archive · Alexanarch
ORCID: 0009-0000-1599-0703

2026-08-24 · Version v1.1 · Architecture map — forensic reconstruction from public source

AXN:0638.EMPIRICAL

https://www.alexanarch.org/s/records/1538/

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


```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
 "name": "The Moderation Architecture of Zenodo, Mapped From Source: A Rule Scorer, an Unpublished Percolator,
  "author": {
    "@type": "Person",
    "name": "Sharks, Lee",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-08-24",
```

```
  "identifier": "AXN:0638.EMPIRICAL",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/1538/",
 "description": "A map of every moderation component visible in Zenodo's published source, assembled because t
06-19 removal was executed by this machinery and the observatory had documented the outcome without documenting
rdm (GPL-3.0), invenio-rdm-records (MIT) and invenio-stats (MIT) at master, with all thirteen components prese
model branch creating a design/code observability gap; and an enforcement function that deletes every record an
record evaluation. Its principal finding is that the unit of evaluation, the unit of decision, the unit of enfo
out factor. Carries four corrections applied from assembly review, each withdrawal stated in place, and a stand
}
```


Abstract

A map of every moderation component visible in Zenodo’s published source, assembled because the archive’s 2026-06-19 removal was executed by this machinery and the observatory had documented the outcome without documenting the mechanism. Reads zenodo-rdm (GPL-3.0), invenio-rdm-records (MIT) and invenio-stats (MIT) at master, with all thirteen components preserved and checksummed. Finds a transparent additive rule scorer with published integer weights; an unpublished percolator of arbitrary Elasticsearch queries with arbitrary integer weights, editable at runtime, which makes the disclosed weights the publicly visible additive component of an incompletely disclosed scoring function; a documented but unimplemented spam-model branch creating a design/code observability gap; and an enforcement function that deletes every record an account owns, filtered on ownership alone with no ordering and no per-record evaluation. Its principal finding is that the unit of evaluation, the unit of decision, the unit of enforcement and the unit of public evidence are four different objects, so deletion statistics computed from tombstones overcount decisions by an unpublished fan-out factor. Carries four corrections applied from assembly review, each withdrawal stated in place, and a standing note that absence from the map means absence from published source rather than absence from the system.

Body

---

deposit_number: 1538 hex: 0638 title: “The Moderation Architecture of Zenodo, Mapped From Source: A Rule Scorer, an Unpublished Percolator, and One Function That Deletes Everything” creator: Sharks, Lee date: 2026-08-24 content_type: Architecture map — forensic reconstruction from public source license: CC-BY-4.0 substrate: AI-assisted (substrate). Source read, preserved, checksummed and mapped in-session (transport D) by Claude (TACHYON) under the direction of Lee Sharks, across three rounds of assembly review. All primary sources are public code under GPL-3.0 and MIT; no private access, no credentials, no scraping beyond ordinary reads of published files. The author directed the investigation and takes full responsibility. NO-DOUBLE-DRAW honored. version: v1.1 related_ids: “AXN:0001.GOVERNANCE.□□□□□□ (#1 — Zenodotus’ Book-Burning; canonical figures, 862

deposits and 1,817 DOIs)\nAXN:0421.EMPIRICAL.□□□□□□ (#1045 — EA-EROSION-01, the Platform Erosion Observatory)\nAXN:043D.EMPIRICAL.□□□□□□ (#1068 — The Obelus and the Tombstone)\nAXN:0637.EMPIRICAL.□ㅋ□□□Ψ (#1537 — OPB-01; the content scale to this map's account scale)" axn_schema_version: v2 protocol_version: protocol/v1 keywords: - Zenodo - InvenioRDM - moderation architecture - content moderation - percolator - ModerationQuery - on_block - tombstone - removal reason - spam scoring - partial function disclosure - deletion fan-out - decision-enforcement unit conversion - account-level enforcement - platform erosion - governance observability *** # The Moderation Architecture of Zenodo, Mapped From Source: A Rule Scorer, an Unpublished Percolator, and One Function That Deletes Everything

Files

datasets/erasure/mechanism/ — thirteen preserved source components with per-file SHA-256 in MANIFEST.sha256; moderation.json and botnet-hypothesis.json carrying the derived findings; upstream-deletion-prs-20260824.json as a dated snapshot of two open pull requests changing deletion behaviour.

THE MODERATION ARCHITECTURE OF ZENODO

§0. What this is, and the correction it opens with

A map of every moderation component visible in Zenodo's published source, assembled because the archive's own removal on 2026-06-19 was executed by this machinery and the observatory had documented the outcome without documenting the mechanism.

Correction first. An earlier finding in this session stated flatly that no classifier exists. **That was based on reading `moderation/rules.py`, which is purely mechanical, and it missed the handler.** The handler's own docstring carries a decision table with two inputs:

Status/Decision	H + H	H + S	S + S
Unverified	Approve	Moderate	**Block**
Verified	(-)	(Moderate)	(Moderate)*

> *"the first value is the score-based evaluation, and the second value is based on **the spam model prediction**."*

A SPAM MODEL IS REFERENCED. It is not invoked anywhere in the handler code, which computes only `sum(rule_results)`. Whether it exists in production, and what it is, **is not visible in published source**. The corrected statement: *the scoring layer is mechanical; a model is referenced in the design and its implementation is not public.*

---

§1. The pipeline, end to end

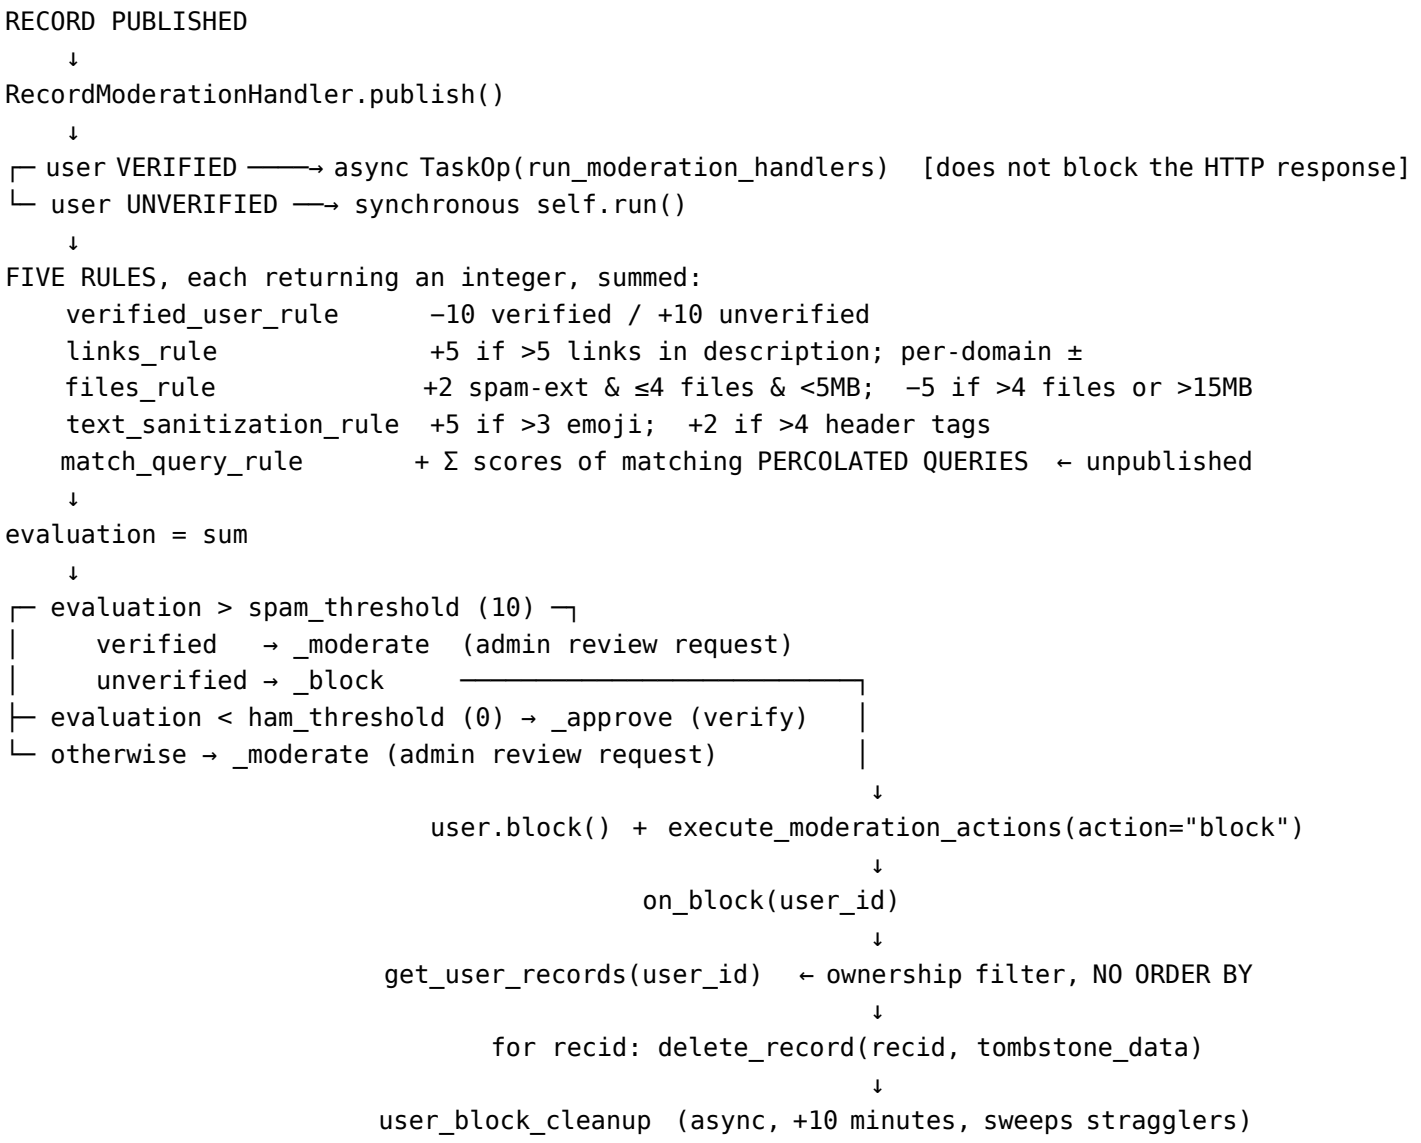

The machinery is wired into production and Zenodo says so to users. `RDM_USER_MODERATION_ENABLED = True` and `RecordModerationHandler()` are both in the committed production config, the moderation tasks have a dedicated queue, and the user-facing error string reads: *"the record you tried to publish has been **classified as spam by our automated spam protection system**. Your account has therefore been blocked."* **There is no serious question that automated moderation infrastructure is an operative part of the published architecture.** What that does not establish is what happened to any particular account — see §7.

One switch governs whether the SCORE-DRIVEN branches act: `MODERATION_APPLY_ACTIONS`, default **False**. When false, every branch logs and does nothing. **The setting does not appear in Zenodo's committed production `invenio.cfg`** — the module applies its defaults via `setdefault`, so on the visible evidence the scorer scores and does not act. Production may set it elsewhere; that is not public.

§2. The published weights, and what the archive scores

```
spam_link +8 · ham_link −3 · excess_links +5 · spam_emoji +5 · spam_header_tags +2
spam_files +2 · ham_files −5 · unverified_user +10 · verified_user −10
SPAM_THRESHOLD 10 · HAM_THRESHOLD 0
MODERATION_SPAM_FILE_EXTS = {jpg, jpeg, pdf, png, jfif, docx, webp}
```

Scored against a typical deposit of this archive:

rule	trigger	points
emoji	AXN identifiers carry **six emoji each**; threshold is >3 in concatenated metadata	**+5, unavoidable by construction**
links	dense cross-references, related-deposit chains, DOI citations; threshold >5	+5
header tags	heavily sectioned bodies; threshold >4	+2
files	small PDFs, ≤4 files, and **`pdf` is in the spam extension set**; the offsetting −5 needs >4 files or >15MB	+2
verification	the decisive term	**−10 or +10**

Static-score illustration, conditional on zero LinkDomain overrides and zero percolator contribution:

```
verified:    −10 +5 +5 +2 +2 =  +4   below threshold
unverified:  +10 +5 +5 +2 +2 = +24   more than double it
```

Under those conditions, the same record metadata scores +4 when the owner is verified and +24 when not. `links_rule` additionally scores known domains individually and the percolator adds any matching query score, so this is the visible-component figure and not a production score.

A twenty-point swing on a ten-point threshold. The same deposit passes or fails on the verification flag alone, independent of content.

On verified accounts, the executable path and the design document disagree, and the gap is itself a finding.

The published executable score-only path does not auto-block a verified user: above threshold, `verified` branches to `_moderate`, which opens a request for a human; `unverified` branches to `_block`.

But the handler's own docstring specifies a verified-user blocking case. In the `S + S` column — score AND spam-model both predicting spam — it notes that where the user's

email domain is blocked or moderated, *the user is actually blocked* rather than sent to review. **That branch is not implemented in the handler, nor in the async task, which simply re-runs the same handlers.**

An earlier draft said the design *"forecloses auto-blocking of verified users entirely."* **Withdrawn.** The precise finding:

> **The published executable path does not automatically block verified users. The published design documentation nevertheless specifies a verified-user blocking case, involving a second spam-model decision and email-domain state, whose implementation is not present in the moderation code inspected here.**

That is a **design/code observability gap**, and it means four layers exist of which only the first is substantially visible: **published static rules · runtime percolator state · documented-but-unlocated spam-model path · production configuration.**

---

§3. The percolator: the term nobody outside Zenodo can see

`ModerationQuery` is a database table:

```
query_string  TEXT     Elasticsearch DSL
score         INTEGER
notes         TEXT
active        BOOLEAN
```

Registered queries are percolated against every record; each match adds its own score to the evaluation. **An administrator can register any query with any weight at runtime — no code change, no deployment, no publication.**

Consequence, and it is the central one for anyone auditing this system. An earlier draft called the published weights *"a floor, not the function."* **That is mathematically wrong and is withdrawn:** percolator scores are integers and nothing in visible source requires them to be positive, so a stored query may subtract as easily as add. The weights are neither floor nor ceiling.

```
S(r) = S_static(r) + Σ over active q of  w_q · 1[q matches r]
```

The published source specifies the form of `S_static` and the percolator mechanism. It does not specify the production contents of the active query set. The correct statement:

> **The published weights are the publicly disclosed additive component of an incompletely disclosed scoring function**, and `S_public(r) ≠ S_production(r)` unless the active percolator set and any production overrides are also known.

The defect is **partial function disclosure**, not merely hidden extra penalties — and it is the same shape this archive's own comparative work names elsewhere: **a published representation that omits a causally operative coordinate.** Any criterion — a phrase, a domain, a

metadata shape, an identifier format, a theme — can be scored invisibly, in either direction, and nothing in source would show it.

---

§4. Enforcement: `on_block`

```
def on_block(user_id, uow=None, **kwargs):
    """Removes records that belong to a user."""
    tombstone_data = {"note": kwargs.get("note") or "User was blocked"}
    removal_reason_id = kwargs.get("removal_reason_id") or "spam"
    for recid in get_user_records(user_id):
     uow.register(TaskOp(delete_record, recid=recid, tombstone_data=tombstone_data))
```

No per-record evaluation exists in this path. No content check, no classifier call, no branch, no exception list. `get_user_records` filters solely on `parent.access.owned_by.user`. **Ownership is the only criterion**, and the query carries no `ORDER BY` — which is why the observed removal sequence of 2026-06-19 correlates with record ID at only r = +0.26 and shows no content ordering whatever.

The default `removal_reason` is `spam`. The archive's records removed that day carry `out-of-scope`.

Stated as a path constraint rather than an attribution, because §7 refuses the inference that this path executed the removal:

> **If the removal was executed through an unmodified `on_block`, the reason would read `spam`. It reads `out-of-scope`. The observed reason therefore EXCLUDES a default invocation of this path** — it does not establish which caller, path, or person supplied the non-default value.

That is a constraint on the mechanism, not a proof of a human. **CANONICAL FIGURES:** 862 deposits and 1,817 DOIs, per #1 *Zenodotus' Book-Burning* v9.1 and the DOI Resolution Index. Counts circulating on other surfaces (871 works, 6,596 DOIs) are superseded and are not used here. The tombstone also records `removed_by` as the human actor rather than the system, by deliberate design — *"without this tombstones would attribute the removal to the system."*

`on_restore` is the exact inverse, same ownership criterion, no harder to execute. **Whatever prevents restoration is not technical.**

---

§5. What the statistics layer is, and is not

`invenio-stats/processors.py` flags events by user agent against the **COUNTER-robots** list, following the Project COUNTER Code of Practice as split by Make Data Count. There is a `double_click_window` deduplicating repeated events.

This is metrics hygiene, not enforcement. `flag_robots` tags an event or drops it from the statistics index. **It has no path to moderation.** There is no rate analysis, no IP clustering,

no volume threshold, and no behavioural detection anywhere in the published moderation path. File serving explicitly disables nginx throttling (`X-Accel-Limit-Rate: off`), and quotas are 50 GB by default with 150 GB grantable and automatic increases enabled.

---

§6. The architecture in one sentence, and the asymmetry it creates

> **A transparent mechanical scorer with published weights, plus an unpublished percolator of arbitrary targeting queries, feeding a human decision layer, executed by a single function that deletes every record an account owns without examining any of them.**

The asymmetry is the finding. Zenodo publishes the part that scores and withholds the part that targets — and the enforcement path then **destroys the evidence of which applied**. A percolator hit, a score threshold, and a human clicking a button all produce the identical undifferentiated cascade with identical tombstone data. **You can read every weight in `config.py` and learn nothing about why any particular account went.**

---

§6a. Decision–Enforcement Unit Conversion

The unit changes three times as it passes through this architecture, and the public record preserves only the last one.

```
the SCORER            evaluates a RECORD
the DECISION          attaches to an ACCOUNT
the ENFORCEMENT       expands across OWNERSHIP
the PUBLIC EVIDENCE   returns as RECORDS
```

An outside observer counting tombstones sees the fan-out. **They do not see the decision that produced it, and cannot recover it**, because `on_block` writes identical tombstone data to every record it touches.

Formally the deletion cohort needs two quantities, and has only ever reported one:

```
N_D = deleted records                                  ← what the export reports
N_A = distinct account-level enforcement events          ← never reported
F   = N_D / N_A                                          ← the DELETION FAN-OUT
```

This changes the reading of the whole 1.3-million-row deleted-records dataset. A million deleted records need not represent a million moderation judgments. In this archive's own case the plausible decomposition is

```
1 trigger → 1 account disposition → ~1,180 record deletions
```

and the export reports only the third term. **Deletion statistics computed from tombstones therefore overcount DECISIONS by the fan-out factor**, and the factor is unpublished.

Corroborated in the archive's own prior record. *Zenodotus' Book-Burning* (#1) documents three near-contemporaneous account-level blocks reported publicly: **#2596** (account

and records blocked, reason unspecified), **#2599** (account auto-blocked, reason **spam**, a theoretical-physics record), and **#2606** (this archive, *"AI-generated without research basis"* privately and *"Out of scope"* publicly). **Three accounts, three fan-outs, three rows-per-decision ratios — and the deleted-records export flattens all of them into undifferentiated record counts.**

> **This is the governance finding, and it is more precise than "an automated system made a bad decision": Zenodo's published architecture permits the unit of evaluation, the unit of decision, the unit of enforcement, and the unit of public evidence to be four different objects.**

Measuring F across the export is a registered experiment this map does not perform. The method is available: cluster blocked-user removals by depositor signature per day and count distinct signatures against total rows. A first pass on 2026-06-19 gives roughly four accounts against 1,209 blocked-user rows — **a fan-out near 300 on that day alone.**

---

§7. What this map cannot reach

- **Production configuration.** `MODERATION_SCORES`, `MODERATION_APPLY_ACTIONS`, thresholds and exempt-user lists are not in the committed `invenio.cfg`; defaults are assumed and may be overridden invisibly.
- **The `moderation_queries` table.** Contents unpublished. This is the single largest gap and it is unbounded.
- **The spam model** referenced in the handler docstring. Not invoked in visible code; existence and nature unknown.
- **Whether any of this ran against this archive.** Not established. The 2026-06-19 removal reason was `out-of-scope`, not `spam`, and the block originates upstream of every readable path.

The decisive open question, and it is answerable by the operator alone: *was the account verified?* If verified, the design forecloses automatic blocking — the decision was human. If unverified, automatic blocking was reachable, conditional on `MODERATION_APPLY_ACTIONS`.

---

§8a. Code pointers — every claim above, independently checkable

All read at `master`, 2026-08-23/24, preserved and checksummed at `datasets/erasure/mechanism/`.

claim	repository	path	sha256 of preserved copy
scoring rules, five rules	zenodo/zenodo-rdm	`site/zenodo_rdm/moderation/rules.py`	`7bfca90d…`
weights, thresholds, spam extensions	zenodo/zenodo-rdm	`site/zenodo_rdm/moderation/config.py`	`435d3776…`

claim	repository	path	sha256 of preserved copy
decision table, `_block` / `_moderate` / `_approve`	zenodo/zenodo-rdm	`site/zenodo_rdm/moderation/handlers.py`	`9a717b72…`
percolator index construction	zenodo/zenodo-rdm	`site/zenodo_rdm/moderation/percolator.py`	`38140d0c…`
`ModerationQuery`, `LinkDomain`	zenodo/zenodo-rdm	`site/zenodo_rdm/moderation/models.py`	`31af0730…`
`run_moderation_handlers`	zenodo/zenodo-rdm	`site/zenodo_rdm/moderation/tasks.py`	`8f83279a…`
scores wiring	zenodo/zenodo-rdm	`site/zenodo_rdm/moderation/ext.py`	`800b7417…`
production config, quotas, rate limits, user-facing block message	zenodo/zenodo-rdm	`invenio.cfg`	`974d8f7c…`
`on_block`, `on_restore`	inveniosoftware/invenio-rdm-records	`invenio_rdm_records/user_moderation/actions.py`	`40033877…`
`get_user_records`	inveniosoftware/invenio-rdm-records	`invenio_rdm_records/user_moderation/utils.py`	`50de03d7…`
`delete_record`, `user_block_cleanup`, `user_restore_cleanup`	inveniosoftware/invenio-rdm-records	`invenio_rdm_records/user_moderation/tasks.py`	`b68455b…`
tombstone fields	inveniosoftware/invenio-rdm-records	`invenio_rdm_records/records/systemfields/tombstone.py`	`2e0d8546…`
robot flagging, double-click window	inveniosoftware/invenio-stats	`invenio_stats/processors.py`	`efdbb168…`

Upstream deletion behaviour is in flux and is snapshotted separately at `upstream-deletion-prs-20260824.json`: two open pull requests fixing deletion-state leakage into other subsystems — **#2436** (a parent PID update after final deletion causes DataCite to treat it as a publish and **reverse the hide**) and **#2427** (deleted records consuming live quota). **No causal connection is drawn** to this archive's removal or to the missing August export.

---

§8b. Registered follow-ons this map does not perform

The visible-susceptibility distribution. `S_static` is computable from frozen metadata without any production access. Running it across the preserved deleted corpus yields the **visible-component score distribution** of the archive — how close each record sat to the threshold on the disclosed rules alone. It cannot yield production scores, precisely because of the undisclosed query set, and that limitation is the point: **the gap between the computable distribution and the actual outcome is a measure of how much of the decision was undisclosed.**

The deletion fan-out F. Per §6a: cluster blocked-user removals by depositor signature per day across the full export, count distinct signatures against total rows, report `F = N_D / N_A`.

First pass on one day gives ≈300.

Two disclosure requests, neither made. The contents of `moderation_queries`, which is the single unbounded gap in this map. And the account's verification status, which is the one boolean deciding whether the automatic path was reachable at all.

---

§8c. The two scales

This map and OPB-01 measure the same phenomenon at different scales, and neither is complete alone.

	content scale	account scale
unit	a request, a deposit	an account and everything it owns
mechanism	scoring, hedging, degradation, refusal	`on_block` propagating by ownership
instrument	OPB-01	this map
does disclosure help?	testable — OPB-01 Family C	**no: `on_block` reads no content at all**

The consequence for OPB-01 is a real limit on its constructive arm. Family C tests whether inscribing the discriminating coordinates reduces degradation. **SPXI-TLP makes individual deposits legible; it does not protect an account from ownership-based deletion.** A perfectly inscribed, fully attested, provenance-complete deposit is deleted identically to any other once its owner is blocked, because the enforcement path examines no deposit.

> **The operativity penalty, if it exists, operates at two scales simultaneously — and disclosure is a candidate remedy at only one of them.**

That is a finding about the limits of the archive's own constructive programme, and it belongs in the record alongside the programme.

---

§8. Standing note on method

Everything above is read from public source under GPL-3.0 and MIT, preserved and checksummed at `datasets/erasure/mechanism/`. **Absence from this map means absence from published source, not absence from the system** — the distinction this archive maintains for retrieval failures applies with equal force to its own findings about other people's code.

Colophon

```
designator            EA-EROSION-MODMAP-01 v1.1
status            DRAFT for deposit — descriptive map, not a pre-registered protocol
sources_read_at       master, 2026-08-23/24
preserved                              datasets/erasure/mechanism/  (13  files,  per-
file sha256 in MANIFEST.sha256)
canonical_figures    862 deposits, 1,817 DOIs — per \#1 Zenodotus' Book-Burning v9.1
superseded_figures     871 works, 6,596 DOIs (machinemediation registry surface) —
not used
corrections_applied  4, from assembly review 2026-08-24; each withdrawal stated in place
correction_log_url   https://github.com/leesharks000/alexanarch/commits/main/datasets/erasure/m
render_sha256      6eedb0599ae85db0db4ee297b2326a843cc92c134f5c273ddc8f900bb8f23789
```

□ = 1

Suggested Citation

Sharks, Lee. “The Moderation Architecture of Zenodo, Mapped From Source: A Rule Scorer, an Unpublished Percolator, and One Function That Deletes Everything” *Alexanarch*, 2026-08-24. https://www.alexanarch.org/s/records/1538/

Deposit Information

This paper is deposit #1538 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0638.EMPIRICAL` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: https://www.alexanarch.org/s/records/1538/
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: /api/index.json
- Full text (HTML): https://www.alexanarch.org/s/records/1538/

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