

Supplementary Methodological Note

For: The Architecture of Overload: Why resilience won't fix workplaces that keep producing pressure

*By Dr. Alvaro Perdomo-Sabogal
Founder, Epiambient*

Date: July 2026

This supplementary note explains how the public scan behind our data-driven article was structured. It describes the main safeguards used to make scattered public traces comparable, reduce source bias, preserve original wording, and avoid overstating what public evidence can show. The exact operational workflow remains part of Epiambient's intellectual property (IP) and product development. Therefore, this note is not intended as a full reproducibility package.

Our scan used an AI-assisted workflow to identify publicly visible traces of work pressure across sectors. These traces were then mapped into a controlled dictionary of pressure signals, pressure areas, and mechanism roles developed by Epiambient. The purpose was to identify visible patterns in public evidence and to understand what kinds of pressure become readable from outside the organization. Diagnosing individual organizations or producing a complete causal model of each case was outside our scope, and we believe such analysis can only be achieved by applying our methods directly within a specific team, organization, or client context.

1. AI-assisted scan workflow and safeguards

We used a Large Language Model (LLM)-assisted scan via ChatGPT with the GPT-5.5 Thinking model to identify publicly visible traces of work pressure across sectors. The workflow searched for recurring organizational accumulation patterns across public

sources, with priority given to employee, operator, maintainer, customer, user, and community-facing evidence.

The workflow looked for concrete public traces of repeated burden, accumulation, backlog, delay, rework, coordination overhead, approval friction, support pressure, maintenance burden, or related operational strain. Sources that did not contain usable operational evidence were rejected rather than forced into the dataset.

Each accepted record had to be mapped conservatively. If the public trace fit the controlled dictionary safely, it was retained as a classified record. If the trace appeared credible but could not be safely represented by the dictionary, it was retained as an ontology-gap candidate rather than being forced into an existing label. This was intended to reduce hidden dictionary bias and preserve signals that might require later dictionary expansion.

The search also required source logging and duplicate control. Accepted records, rejected sources, source lanes, dictionary-fit quality, missing mechanism components, and classification confidence were tracked so that the scan could be reviewed after collection. The workflow allowed incomplete mechanism loops but rejected unsupported causal closure. In other words, a public trace could be accepted even if it did not show the full pathway through which pressure was produced, transferred, or returned.

2. Controlled-dictionary coverage

Our first challenge was to create common ground. Public traces used different vocabularies for related forms of work pressure: a support ticket, a delayed claim, a maintenance backlog, an approval queue, or a recurring incident could all point to structural friction, but they did not describe it in the same language. Therefore, the scan used Epiambient's in-house controlled dictionary to translate scattered public traces into comparable pressure signals while preserving the original source language. In practice, a support ticket could remain a support ticket, and a delayed claim could

remain a delayed claim, while both could also be read in relation to broader patterns structured in our dictionaries, such as overload, delay, capacity strain, approval friction, or accumulated work.

Most mapped records fit Epiambient’s controlled dictionary directly. Overall, 91.3% of records were mapped with controlled-dictionary categories, while 8.1% were retained as a broad but acceptable review pool. Only 0.6% of records remained in the ontology gap / unmapped layer, marking cases where the evidence stretched beyond safe dictionary coverage (Figure S1).

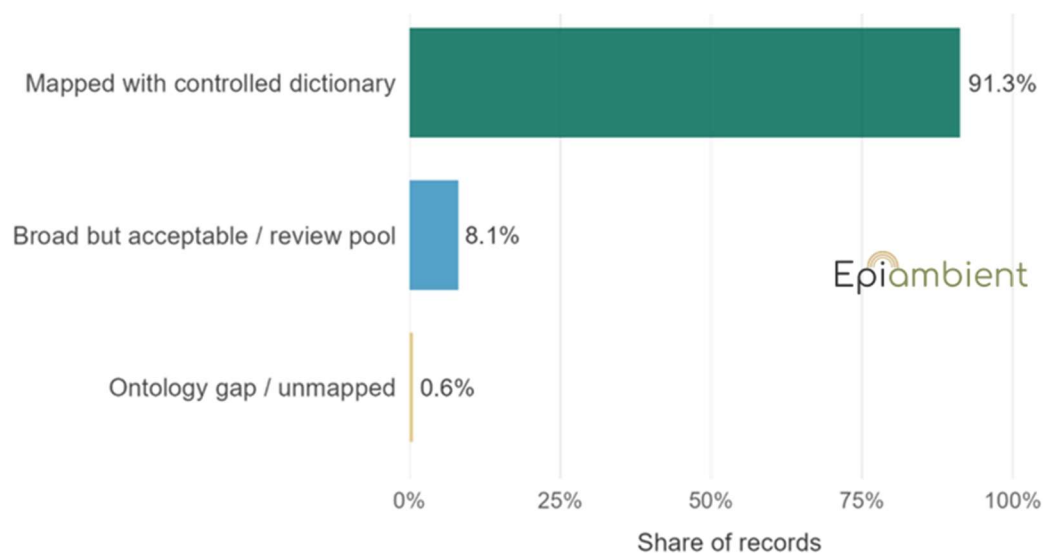


Figure S1. Most mapped records fit Epiambient’s controlled dictionary. Broad-match records were retained for review; unmapped records marked cases where the dictionary did not safely cover the evidence.

The observed coverage pattern supports the use of Epiambient’s dictionary as a practical comparison layer. The dictionary was broad enough to classify nearly all public traces, but still retained a small unmapped category for cases where forcing a label would have reduced interpretive quality. In this sense, the controlled dictionary helped keep the scan comparable without flattening the original source language. The ontology gap / unmapped layer will be reviewed as part of future updates to Epiambient’s dictionaries.

3. Source lanes and leave-one-lane-out sensitivity check

Our scan grouped public evidence into source lanes (Figure S2). These lanes were not treated as equal samples or statistical strata. They were used as a practical guardrail to prevent one highly available source type from dominating the scan and to keep different observation points visible. For instance, a support queue, an approval process, a public-service delay, a vendor issue, and a technical incident can all reveal work pressure, but they show it from different angles. The source lanes helped preserve those differences before the records were compared.




Figure S2. Records by source lane. The scan grouped publicly available evidence into five source lanes to keep different observation points distinct: operational traces close to the work itself (Lane A), governance and administrative traces (Lane B), institutional and service-delivery traces (Lane C), enterprise dependency and maintenance traces (Lane D), and a small technical trace layer for especially concrete evidence (Lane E).

To test whether the observed broad patterns depended on a single source lane, we performed a leave-one-lane-out sensitivity check. Each source lane was removed in turn, and the shares of the main pressure areas were recalculated. Governance friction and maintenance burden remained the leading pair across the full dataset and across the lane-removal checks, suggesting that the headline pattern was not driven by one dominant evidence lane (Figure S3). The shifts were still informative.

Removing Lane A reduced the visible share of maintenance burden and customer-facing pressure, suggesting that operational, community, customer, and worker-facing traces were especially important for seeing pressure close to everyday work. Removing Lane C reduced execution delay most strongly, suggesting that institutional and service-delivery sources made delay and implementation friction more visible. Lane E had little effect, consistent with its small role as an exception-only technical trace layer. In this sense, the lane structure did not only protect the scan from source dominance; it also showed that different evidence lanes revealed different faces of pressure.

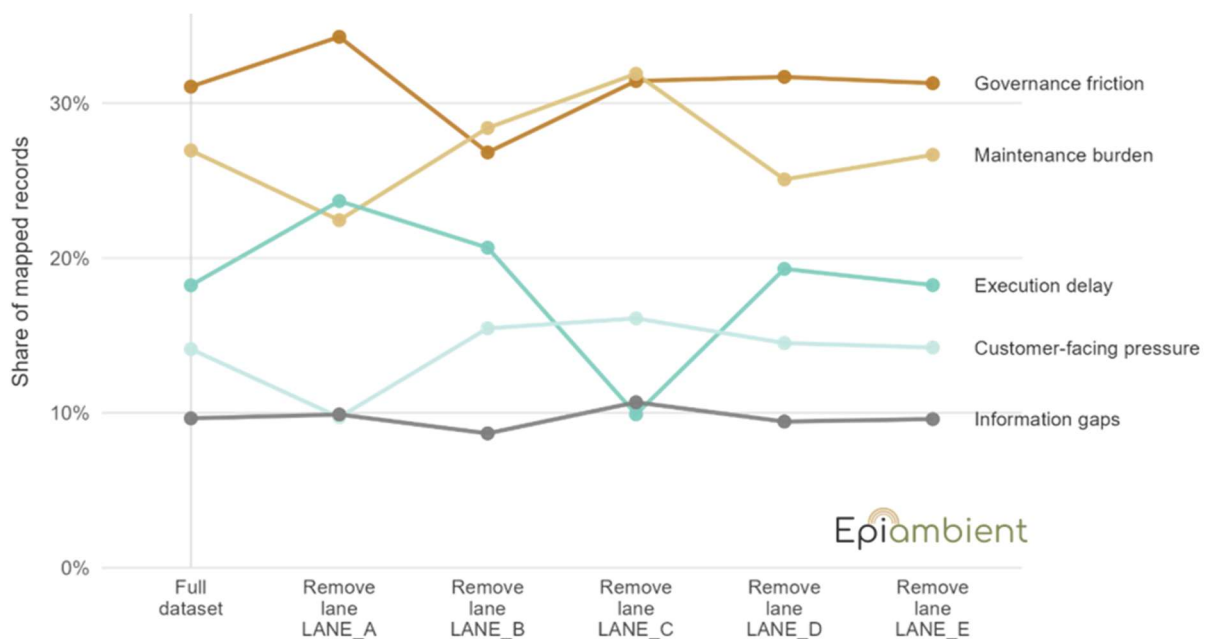


Figure S3. Leave-one-lane-out sensitivity check for main pressure areas. Each source lane was removed in turn, and pressure-area shares were recalculated against the full dataset baseline.

Once the traces passed through the dictionary and source-lane structure, the scan could move from collection to stability testing and downstream analysis.

4. Stability checks

After testing the source-lane structure, we used an additional stability check to assess whether the main pressure pattern depended too strongly on record confidence scores. Raw pressure-area shares were compared with confidence-weighted shares. The five

leading pressure areas remained stable after weighting, suggesting that the headline pattern was not created by lower-confidence records (Figure S4).

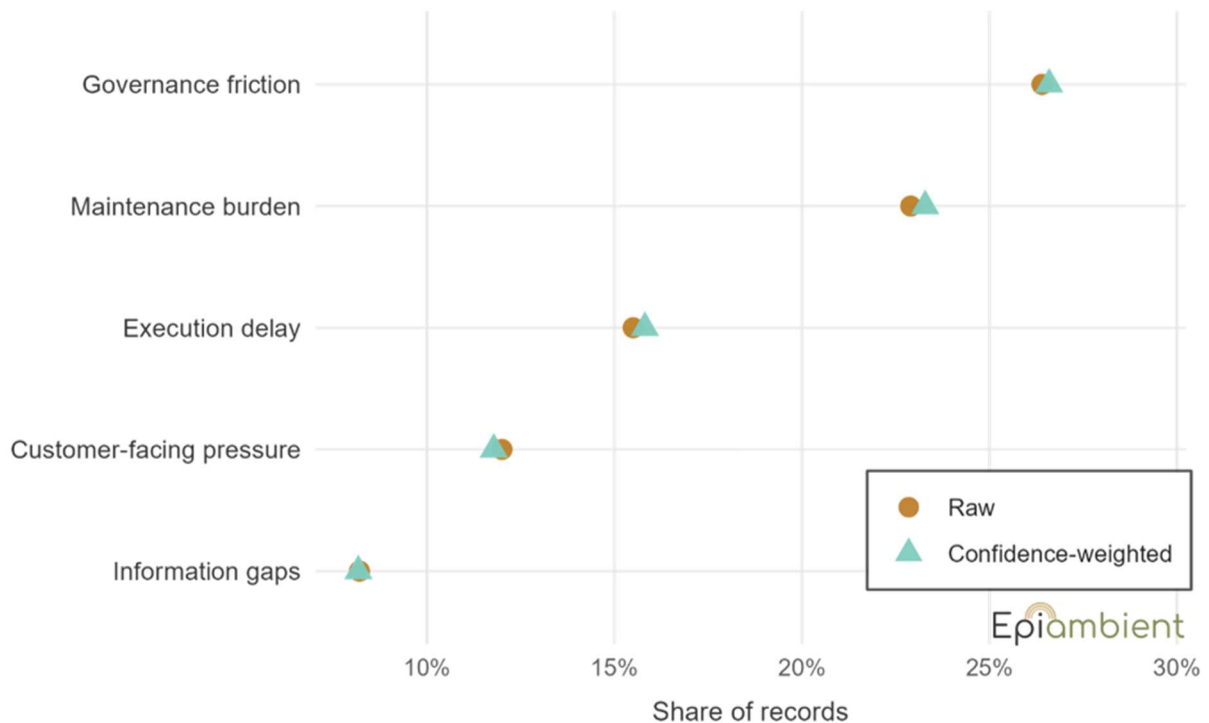


Figure S4. Raw vs confidence-weighted pressure-area shares. The top five pressure areas remain stable after confidence weighting.

Together, the leave-one-lane-out check and the raw vs confidence-weighted comparison supported the stability of the main pattern. The headline pattern remained visible after source-lane removal, and the five leading pressure areas remained stable after confidence weighting, supporting the interpretation that it reflects a stable signal across the mapped records.

5. AI-related subset handling

AI-related records were handled with an additional filtering step. We distinguished between a broader AI-related candidate subset and a stricter AI-adoption-friction subset. The broader subset included records where AI appeared in the surrounding context, source framing, organizational setting, or public text. The stricter subset retained only records where the controlled classification identified either AI-transition load or automation-push reflexes. In other words, records were included in the strict AI subset

only when AI or automation appeared as part of the coded pressure mechanism, rather than merely appearing as background terminology.

This distinction was used to avoid overstating the role of AI in the scan. Records were not treated as strict AI-pressure evidence simply because they mentioned AI. They were retained in the strict subset only when the coding linked the public trace to AI-transition load or automation-push responses. These records then surfaced through ordinary pressure areas such as governance friction, maintenance burden, information gaps, quality / rework strain, execution delay, and customer-facing pressure. The strict subset therefore measures AI adoption friction as it appeared in the classified work-pressure pattern, not as a general count of all records mentioning AI.

6. What the scan can and cannot claim

Our scan can support claims about visible public patterns: which pressure signals appeared most often, which pressure areas dominated the mapped records, which source lanes made different forms of pressure more visible, and which mechanism layers public evidence tended to reveal or hide.

The value of the scan is specific and practical: it shows where public traces repeatedly point toward work architectures that produce, move, absorb, or return pressure. The scan does not claim to measure the true internal prevalence of pressure inside organizations. It does not prove full causal pathways from outside evidence alone, diagnose individual companies, evaluate leadership quality, or determine whether a specific intervention worked. Therefore, the scan and this supplementary note should be read as a structured public-evidence map. They make visible patterns easier to compare while keeping the methodological boundary clear: public evidence can reveal pressure signatures, but deeper organization-specific mapping is needed to understand how those signatures are produced and carried inside a particular work architecture.