



Case Study: A SR&ED claim from a third-party codebase

How a Canadian SaaS company filed 18 months of R&D activity without interviewing a single developer — and how HammerLedger surfaced a mid-year contractor change from the data alone.

The setup

A Canadian SaaS company building a leadership development platform needed to file its 2024-2025 SR&ED claim covering the fiscal year of September 2024 through August 2025.

The challenge was unusual but increasingly common: **the engineering work had been built by a third-party development contractor, not in-house staff.** The company had no internal developers to interview about what they had built, why, or how much of their time went where.

Worse, partway through the fiscal year, the development contractor was changed. The codebase had two different contributor patterns layered on top of each other, and neither was available for the traditional retrospective interviews specialists rely on.

By the time tax preparation rolled around, the company faced a familiar SR&ED problem in a much harder form: **18 months of R&D activity, no developers to ask, and a mid-year handoff to explain.**

Why the traditional approach failed

Most SR&ED specialist firms charge 15–25% of the recovered credit to perform four tasks:

1. Interview developers (4–8 hours each) to reconstruct what they worked on.
2. Write technical narratives in CRA-expected language.
3. Categorize labour and time against eligible projects.
4. Assemble a supporting evidence binder.

When developers aren't reachable — because they were contractors who have since moved on — step 1 cannot happen. And without step 1, steps 2 and 3 lose their primary data source.

The codebase was the only remaining source of truth. So we pointed HammerLedger at it.

What HammerLedger extracted

The full ledger was created and analyzed in **44 minutes** — from upload to a complete first-pass evidence chain.

From the source files alone, the platform identified:

What it found	How much	What that means in plain terms
Files reviewed	605	Every meaningful file in the codebase was read
Files with R&D activity	335 (55%)	Just over half the codebase contained eligible work — the rest was setup, configuration, dependencies
R&D findings	516	Specific eligible R&D activities — algorithms designed, performance challenges solved, architectural choices made under uncertainty
Components	17	Major product areas (e.g., authentication, billing, AI features) — the same categorization that ends up on the T661
Technologies referenced	823	Distinct frameworks, libraries, and tools used — the foundation of the technical narratives
Contributors identified	40	Every git author who touched the codebase — including two AI code-generation tools correctly flagged as non-human

The contractor change, surfaced from the data

Most importantly: **HammerLedger detected the mid-year contractor change automatically.**

The handoff wasn't documented anywhere in the project specs. It showed up in the data:

- Distinct commit-style fingerprints (file structure, commenting conventions, dependency choices) before and after a specific window.
- A clear drop-off in activity from the original contributor cluster.
- A new cluster picking up with characteristic ramp-up patterns.

This matters for SR&ED because the work product on either side of the handoff has to be allocated correctly — labour categorization that a specialist would normally derive from invoices and interviews. HammerLedger surfaced it from the code itself.

Confidence and evidence quality, explained

A first-pass extraction is only useful if you can tell which findings to trust.

Finding confidence is the AI's self-reported strength of each finding, on a 0-to-1 scale: "how clearly does this code or commit demonstrate R&D activity?"

Confidence	Count	Plain-English read
Strong (≥ 0.80)	212	The AI is confident this is R&D — the code articulates the experimental work itself
Moderate (0.60–0.79)	295	Clear R&D signal, but with context the human reviewer should confirm
Below moderate	9	Surfaced for review but flagged as weaker — not silently included

98% of findings landed at moderate confidence or higher. The 2% below the threshold went into a human-review queue — visible to the reviewer, not invisibly stitched into the filing.

Evidence quality runs at the claim-section level — the bundles of evidence assembled to support each T661 narrative. 266 sections were drafted (133 paired Technological Uncertainty + Proof of Experimentation sets, the structure CRA expects).

Evidence quality	Count
Strong	8 sections
Moderate	192 sections
Weak (flagged for additional support)	66 sections

Every section traces back to specific files, commits, and line numbers. Nothing is inferred from outside the source material.

The human gate

A first-pass extraction is **not a filing**. Findings, allocations, and contractor attributions were surfaced for review — the client (who knew the business context the code didn't explicitly state) signed off on what was eligible.

Client-verified accuracy on the team and project allocations: 85%.

The 15% gap reflected work that didn't show up in commits — design discussions, planning meetings, documentation produced outside the repository, prototype work that never landed in main. HammerLedger is honest about that limitation by design: the chain of

custody only extends as far as the source material reaches. The remaining 15% was added by the client in review, with a clear distinction between machine-extracted-and-sealed and human-added-context.

What this case demonstrates

	Without HammerLedger	With HammerLedger
Developers to interview	None (contractor — unreachable)	None needed
Time to first-pass evidence	Weeks of reconstruction	44 minutes
Project allocation source	Specialist estimation from invoices	Code + commits, 85% client-verified
Contractor change tracking	Manual reconstruction	Detected from commit patterns automatically
Traceability per finding	Narratives from memory	Every finding → file → commit, cryptographically sealed

R&D tax credit eligibility, project allocations, and contributor attributions were extracted from 605 files of third-party code — in under an hour, without a single developer interview, and with an honest accuracy ceiling the client could review and complete.

What this means for buyers

If your developers are contractors, or have left the company, or simply don't remember 18-month-old work: the traditional SR&ED workflow breaks. HammerLedger doesn't depend on interviews — it depends on the code or data, which is already there.

If you're an accountant whose clients face this situation: the evidence package is defensible enough that you can file the T661 in-house without referring to a SR&ED specialist firm.

If you're a CFO paying 15-25% of your credit to a specialist today: the math on what they're doing for that fee — and what could be automated — is here: hammerai.ai/calculator

Try HammerLedger free for 7 days: hammerledger.app/sign-up

Open-source CLI: hammerai.ai/cli

Reach us: support@hammerai.ai