


# Agentic quantification of federated BIM: an architecture for live model-to-decision linkage

*Field notes from a UK hospital engagement · 6 min read*

The project director asks: “How confident are you in September completion? And what will the curtain-wall supplier cost us if they fall behind a fortnight?”

On most projects today, the honest answer is **nobody knows**. The model lives in Revit. The programme lives in P6. The procurement schedule sits in someone’s inbox. The specs are PDFs scattered across SharePoint. The cost plan is a spreadsheet that was true the day it was issued and has been drifting ever since. By the time someone reconciles it, the answer is two weeks out of date and the next variation has already landed.

This paper sketches the architecture of an agentic layer that sits between these systems — one that quantifies a federated BIM model, ties every element to its procurement, programme and specification context, and keeps that linkage current as the model evolves. We describe the approach, summarise field results from a recent hospital engagement, and outline the two engagement structures we offer.



*Federated hospital BIM model in the pipeline view — a 23,466-element federated model from a 4,500-file project folder, processed end-to-end in under four minutes. Project identifiers anonymised.*

## 1. The quantification problem

Quantifying a BIM model — turning its geometric content into a classified, counted, costed view that downstream cost planning and risk analysis can read — is one of the largest hidden costs in digital engineering today. A quantity surveyor takes a federated Revit model, exports element schedules, manually maps them to procurement packages and NRM2 categories, reconciles against an incoming BOQ, and updates the result with every variation. The work is slow, the output is obsolete within weeks, and the cost compounds: each party in the supply chain that wants a costed view of the model rebuilds it from scratch.

Most of the recent industry investment in 5D BIM has gone into better authoring tools rather than better quantification — the bottleneck has moved, but it hasn't shrunk. Our work focuses on closing that gap once, for the federated model, and keeping it current as the model and its surrounding documents change.


## 2. Field results: a recent hospital engagement

On a recent hospital engagement, we processed 23,466 model elements from a 4,500-file project folder in under four minutes. The deterministic rule layer alone classified nearly half of them before any AI was asked to look. The agentic layer committed on a further large slice with median confidence above 0.9.

**The remainder — items the agent could not confidently resolve — went to a tightly-scoped review queue rather than being silently guessed at.** That abstention behaviour is the design feature. A reviewable shortlist of a few thousand items is operationally useful; a confidently-wrong classification of all of them is operationally dangerous.

Every committed classification carries an auditable citation back to its source — the property field, the spec page, the procurement schedule row, or the prior decision on a similar element from earlier in the project. Nothing silent. Nothing hallucinated. Nothing the QS could not defend in a meeting.

The trick is not the AI. The trick is what surrounds it: deterministic rules that narrow the work before any token is spent, schemas that reject malformed answers, citations enforced as code, and a project-by-project knowledge layer that accumulates. Within a single project we already observed that layer resolve roughly a third of elements on rerun without firing the agentic stack at all. Compounding that signal across engagements is the design intent — and the structural advantage the architecture is built to capture.



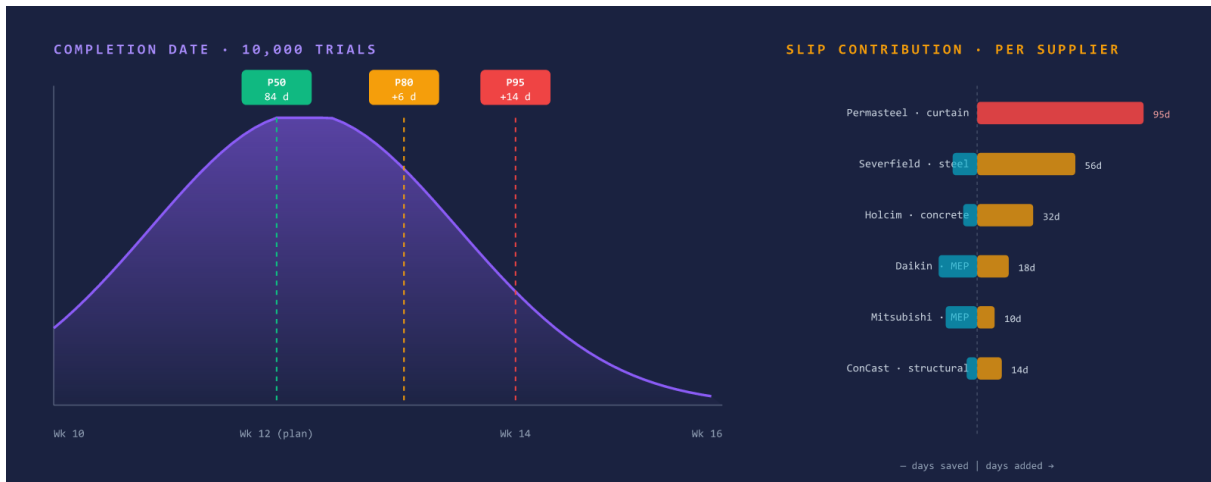
*Element-level audit in the BIM pipeline view. Every committed classification surfaces its source — here, a compound ceiling element with its Revit category, parent and design option visible alongside the geometry. The same trace stands behind every line of the materialised view.*

### 3. Three problems this is built to address

**The reconciliation problem.** A quantity surveyor spends the better part of a week each month reconciling drawings to BOQ. The structure we deliver is a model-grounded BOQ where every element ties to a priced package with audit trail. Re-runnable on every model revision; the priced delta of a variation is minutes of compute rather than days of manual reconciliation.

**The drift problem.** A static contract programme is obsolete a week after it is saved. We re-parse the programme, re-read the model, and rebuild the activity-to-element links on demand — designed to run nightly in production so the PMO dashboard is closer to a live instrument than a quarterly artefact.

**The risk-reporting problem.** Monte Carlo on a programme is normally run once, at tender, and forgotten because the inputs go stale. We structure those inputs as live data: every activity carries a duration distribution, every supplier a lead-time distribution, every element ties to both. A critical-path kernel reads that view as a function and produces P50 / P80 / P95 windows plus a tornado attributing slip to specific suppliers. Re-runnable as often as you are willing to schedule it.



Probabilistic completion distribution after 10,000 trials. Tornado chart attributes slip risk to specific suppliers — so the team knows where to push, not just that there is pushing to do.

**A note on what is shipped and what is designed.** The classification, BOQ-link and procurement-link layers are demonstrated on a real project. The nightly Monte Carlo cycles and the cross-project knowledge compounding are designed for, and partially validated within a single project — the production cycles themselves come with the engagement, not before it.

#### 4. Beyond quantification: live federated state

Once a model is quantified and linked to its procurement and programme context, the next move is to make the federated model itself a live system. Each element gains an extensible state vector that the project actually cares about:

- planned
- specified
- design approval
- procured
- material on site
- installed
- inspected and signed off

Each state transition becomes a write back to the federated model with a citation — a delivery note, an inspection record, a supplier confirmation, a snag closure. The result is what many UK projects have wanted for years and very few have actually achieved: a single source of truth that says *this column has been ordered, this beam has arrived, this wall has been installed* — with an audit trail every project board can defend.

This is the layer we are building next. The quantification layer is delivered today; the live federated state layer is in active development; the integrations that keep them current are tuned per engagement.

## 5. Where this applies

Contractors with procurement schedules and contract programmes that have drifted out of alignment with the model. Architects whose project folders have grown organically across teams and consultants. Developers who want to know the *actual* risk on their programme rather than the version that fits on a board paper. PMOs who have stopped trusting their own status reports.

If you have already built some internal BIM tooling and it is not quite reading the model the way you had hoped, this work plugs in alongside what is there rather than in place of it.

## 6. What this approach is not

This is not a platform. It is not another BIM authoring tool. It does not replace the QS, the planner or the BIM manager. It gives them a live, model-grounded view of the project they currently compile by hand — and steps aside once that view is in their workflow.

Project data does not leave the client tenant. What accumulates between engagements is the pattern of work — the anonymised classification fingerprints — not the data itself. Opt-out is available, and priced accordingly.

## 7. Engagement structures

Two modes, deliberately distinct rather than two price points on the same product.

**Per-project engagement (1–3 weeks).** For teams that need a single model quantified, reconciled against contract documents and structured for risk and cost analysis on a specific job. Deliverables: a dashboard, an audited materialised view, and the Monte Carlo input layer. The model and the documents remain in the client tenant; only the resulting analysis leaves. **Indicative scope: £8–30k** depending on model scale and document richness. Typical turnaround: two weeks.

**Strategic embedded engagement (4–6 months).** For organisations that want the pipeline as a permanent part of their delivery capability. We tune the architecture to the client's tools (Primavera P6, Aconex, the client's BIM standards, the client's procurement workflow), train the team, and embed the system in the delivery process. The code, the architecture and a clone of the cross-project knowledge layer transfer with the engagement. **Indicative scope: £80–180k** depending on integration breadth and rollout horizon, with an optional managed-service tail for the first quarter post-handover.

Both modes contribute to the same accumulating knowledge layer. Per-project engagements feed it and partially benefit from it; strategic engagements own a tuned clone of it. The shape of that compounding moat is, we believe, the most interesting architectural property of the system — and the one hardest for an off-the-shelf BIM tool to replicate.

## Contact

If you have a Revit or IFC model — and the contract documentation around it (programme, cost, specification and standards) needs reconciling, sooner rather than later — that is the conversation we want to have.

---

**LEVERET SYSTEMS**

Agentic AI engineering for AEC and manufacturing.

**leveretsystems.co.uk** · [ejoubert@leveretsystems.co.uk](mailto:ejoubert@leveretsystems.co.uk)