

ScopeLintel | Know what your bid carries.

Preconstruction scope intelligence. Early-stage concept / synthetic demonstration / October 2026.

THE PRODUCT: A planned Claude-centered review layer for U.S. commercial general contractors. Turn messy subcontractor proposals into source-linked scope intelligence, then return reviewed results to the existing Excel bid tab.

THE ENTRY POINT: Mechanical packages for occupied office renovations. Follow a requirement across proposal clauses, revisions and clarifications. Preserve unresolved obligations instead of silently treating them as included or free.

THE EXAMPLE: Harbor Point Office Renovation. Three synthetic HVAC proposals arrive as a formal PDF, a spreadsheet and a revised email. Submitted bases: A \$486,000; B \$512,000; C \$475,000.

THE REASONING TEST: C increases its price by \$17,000 for standalone controls. That amount is already included in its revised \$475,000 base. BMS gateway and existing-system tie-in remain excluded. The correct result preserves that open obligation; it does not add the \$17,000 twice.

THE COMPARISON: Quantified subtotals are A \$527,000, B \$512,000 and C \$505,500. A and C still have unresolved scope. These are not approved award totals, and the lowest incomplete subtotal is not a recommended winner.

60-SECOND REVIEW: Open the package; compare the matrix; inspect E07 and its source; inspect ambiguous TAB; download the Excel workbook. All sources are supplied.

Claude at the core. Decisions stay human.

CLAUDE, PLANNED: multi-document reasoning, structured extraction, semantic scope normalization, qualification interpretation, ambiguity detection and clarification drafts.

CODE: validate source anchors and output structure; calculate monetary adjustments; prevent revision double-counting; map approved values and open issues into Excel.

ESTIMATOR: approve the scope baseline, resolve ambiguous clauses, accept or reject proposed adjustments and own the award decision.

WHY THIS IS NOT CLAIMED AS UNIQUE: MeltPlan, Bridgeline, Procore and Autodesk already address parts of AI leveling. Excel export, evidence and revision history alone are not differentiation. Our proposed test is whether requirement-level obligation continuity and decision review provide material additional value in the chosen workflow.

CURRENT STAGE: landing page, five-view concept interface, synthetic sources, hand-authored expected outputs and formula-based workbook. No live Claude inference, customer traction, measured accuracy or production security certification is claimed.

NEXT PROOF: instrumented Claude runs; expert-annotated historic packages with permission; competitor comparisons; estimator review-time and material-omission evaluation.

VISION: connect requirements, offers, clarifications and award scope to a durable customer-owned record of what was promised. Expand only when buyer demand and measurable quality justify it.

SOURCE BASIS: Official Anthropic startup FAQ and 6 October 2026 expansion; current first-party competitor pages. The full application report contains dated references and readiness gaps.