

AI-Assisted Capital-Plan Sequencing

Version: 1.0

Date: 2026-08-16

Type: Educational methodology — generic hotel operations

Purpose

A brand's Property Improvement Plan (PIP) arrives as a list of scope: rooms to renovate, systems to replace, finishes to bring to standard. What it never arrives as is a plan — the sequence, the phasing, the answer to *when each piece happens and what it costs to be closed*. Historically, turning scope into sequence took a consultant engagement or a general manager's evenings for a month.

This document describes how a property can build that plan itself using AI over three records it already owns. No property-specific data appears here — the method is generic, and any figures used are illustrative.

The Three Inputs

Everything the sequencing needs already exists at the property:

1. **The scope document (the PIP itself).** What must be done, to which spaces, to what standard.
1. **The payroll and roster.** Which trades exist in-house, what capacity they truly have, and when the labor year swells and shrinks — including any collective-agreement rules about seasonal movement of staff.
1. **The demand calendar.** Daily occupancy and revenue history, at least one full year. This is the input most properties have and least use.

If any of the three is missing, build it first. A sequence computed on a guessed demand calendar is a guess wearing a schedule.

The Method

Step 1 — Extract the scope into work units

AI reads the PIP and returns a structured register: one row per unit of work, each carrying the space it touches, the trade it needs, what it depends on, and a duration field that starts as **UNKNOWN**. Resist the urge to fill durations with estimates at this stage — an unknown that is labeled unknown can be measured later; a guess gets mistaken for a fact.

Step 2 — Read labor reality from the payroll

The roster answers what the scope document cannot: which trades are on staff versus contracted, what the true maintenance capacity is per week, and when seasonal labor rules open or close the option of shifting people. AI cross-reads the register against the roster and flags every work unit whose trade the property does not employ — that flag list is the contracting plan.

Step 3 — Price the closure, not the construction

Every room taken offline costs the revenue it would have earned **that night** — which the demand calendar knows precisely. A room closed in a 60%-occupancy shoulder month often costs nothing at all, because the house had empty rooms anyway; the same room closed in a 95% peak month costs a full night's rate. AI joins the register to the calendar and prices each candidate window in displaced room-nights. This inverts the usual instinct: the question is not "when is construction convenient" but "when is the building cheapest to be partially closed."

Step 4 — Sequence

With scope, labor and closure-cost as data, sequencing becomes a solvable arrangement problem instead of a debate: phases that respect trade dependencies, stay inside real labor capacity, and sit in the cheapest closure windows. AI drafts the sequence and — critically — shows its reasons per phase, so a human can challenge any placement with a better fact.

Step 5 — Measure the first block, then re-sequence

The first completed block of work is the only real measurement of what this scope costs and how long it takes *at this property, with these trades, on this supply line*. Feed those measured durations back into the register and let the sequence recompute. Every phase after the first is then built on measurement; only the first is built on estimate.

The Rules That Keep It Honest

- **The register is the plan.** One table of work units, each row carrying its status and its evidence. Documents are printed *from* the register; the register is never updated from a document.
- **UNKNOWN is a permitted value.** A schedule that forbids unknowns

manufactures fictions to fill them.

- **AI drafts; a named person decides.** Every phase in the sequence has an owner who can be asked "why is this here?" and can answer from the data.
- **Closure cost is a nightly fact, not a season-level vibe.** Price windows from the daily calendar, never from "summer is busy."

What This Replaces

A sequencing exercise of this shape has traditionally been bought as consulting or extracted from an operator's scarce attention. With the three inputs in order, AI performs the reading, joining and arranging in hours, repeatably, and re-runs the whole sequence any time a fact changes — a lead time slips, a block finishes early, a season books differently. The judgment stays human. The labor of turning records into a plan is what the machine removes.