

Best practice stops living in a binder and starts running through your team.

Most capital improvement starts and stops at one artifact: an estimate someone validates, a schedule someone reviews, a gate package someone reads. That work matters, and we do it as deeply as anyone in the industry. But a review fixes one document. It does not fix the system that produced it.

- One system runs the work and grades it.

— WHERE MOST OWNER PROGRAMS ARE TODAY

The system that produced the package is still the system after the findings close.

These three survive a good review, and they are the ones that move the cost and the in-service date.

THE GATE PROCESS TODAY

One gate ladder that big projects tolerate and small projects route around

ACCOUNTABILITY TODAY

Owners are named in a matrix, not on the deliverable that is due

THE RECORD TODAY

Gate decisions and the basis behind them live in slides and inboxes

01 The standard you own

CAP·CORE · THE CAPITAL PROJECT MANAGEMENT SYSTEM

One process, run at four depths. A major project carries full front-end rigor and every gate; a replacement in kind turns down to sixteen activities and two gates. A seven-factor calculator scores which depth a project climbs. Underneath it is the documented best-practice version of a complete owner's capital program, governance through closeout, built on public frameworks (AACE, CII front-end planning practice, PMI, DCMA, GAO, API). It is the yardstick every review measures against, and it is **yours: transparent and editable, not a black box you rent and cannot see inside.**

47

Standards across the full capital lifecycle

156


Activities at Major; 16 at replacement in kind, same process

3+3

Funding gates, then delivery gates; four outcomes and a record at every one

4

Tiers, scored by a seven-factor calculator, so the light track stays light



02 Made live, and it coaches the team

CAP·PATH · THE CAPITAL PROJECT MANAGEMENT PLATFORM

CAP·PATH deploys the standard to every project at its tier as a dated, gated plan anchored to the sanction gate and the field start, with a named person responsible for every activity and a named person accountable for every acceptance. The governing standard and the expected deliverable sit one click away on the exact step they apply to. A project manager or a controls lead does not just get told they own something. **They see what they own, when it is due, what good looks like, and the standard behind it.** Running the process is how the team learns the standard, project after project.

It coaches the person who owns the work, then it keeps the record. On the exact step someone owns, the governing standard, what good looks like, and the date the gate needs it by are already there. The deliverable and the decision behind it are memorialized on one transparent record, and the disciplines who need to weigh in are brought in while the work is done rather than at the gate. Then the control: nothing closes without the accountable person's acceptance against its acceptance test, and a gate closes only on a recorded proceed, recycle, hold, or cancel. It also lifts above the single project: one portfolio carries every project across your sites gate to gate, with turnaround and project cost stacked by year.

ON THE STEP
The standard

RESPONSIBLE
Build it

KEPT
The record behind it

ACCOUNTABLE
Acceptance

RECORDED
Gate decision

The standard, what good looks like, and the due date sit on the step itself. Nothing closes without the accountable person's acceptance against its acceptance test, and a gate closes only through a recorded proceed, recycle, hold, or cancel, on an append-only audit trail.

GOES DEEPER Two companion pieces go deeper at site-leadership altitude: **The Capital Portfolio** one-pager, and the **CAP-PATH + CAP-CORE** sheet for the six capability groups and the standard.

03 Assured on the same record

CAP-AR · CAP-SAR · CAP-SPR · INDEPENDENT ASSURANCE

Because the work runs on one project record, independent assurance sits right beside it, grading the same standard. **CAP-AR** is the gate arc: a feasibility review ahead of G1, a definition review ahead of G2, a sanction readiness review ahead of the final investment decision, and a closeout review that files the benchmark record, each scoring definition completeness, the estimate class and its basis, schedule integrity, and team alignment. **CAP-SAR** works the integrated schedule in full, procurement logic and the tie-in window included, and hands back a prioritized correction plan before the baseline locks. **CAP-SPR** joins the sanction baseline to the as-built and tests every explanation for the slip against your own records. CAP-AR4 is the look-back over the whole project at G6, business case, final cost, lessons, benchmark record; CAP-SPR is the schedule forensics inside it, or standalone when you want the why before closeout.

Where it counts, the assurance is AI-native. One-on-one interviews run by agent, by voice or text, at each person's convenience, across every discipline seat, with no site visit to build a week around. Every finding is anchored to a quote in the transcript and adversarially verified before it reaches you. **The interviews gather. The consultant decides. The gatekeeper still decides the gate.**

WHO YOU ARE WORKING WITH

The guide who has sat in the seat, on both sides of the table.

Owner-side experience across refining, petrochemicals, industrial gas, LNG, and power, as the owner accountable for the capital outcome and as the consultant brought in to protect it. The capital-projects practice of STO Intelligence LLC, on one spine with its turnaround practice, so a project and the window it rides are one record.

Independent · Owner-side · Veteran-owned · Public-framework methods

WHERE THIS IS GOING

Every project writes to one independent data spine.

Live today: the standard, the platform, and the project record every benchmark will be built from. Next: the benchmark itself, for the sustaining and small-capital portfolios nobody else measures, then forecasting of cost and schedule exposure by tier and estimate class before sanction, and assurance between the gates rather than only at one.

No benchmark dataset exists yet. Yours would be among the first records in it.

THE LOOP Most lessons learned change nothing. Here each becomes an **owned, dated action** someone other than its owner verifies, the ones worth keeping are promoted into the standard, and they come back on the next project's matching step.

ONE REVIEW AT A TIME HOW IMPROVEMENT USUALLY WORKS

- ✗ The review closes and the process that produced the finding is unchanged.
- ✗ The gate ladder lives with two or three senior project managers, and small projects route around it.
- ✗ Sanction decisions are defensible only as far as someone's memory.

ONE SYSTEM, LIVE AND GRADED THE FULL VALUE LINE

- ✓ Findings close against a standard your team can see and edit.
- ✓ Every deliverable has an owner, a date, and what good looks like, at every tier.
- ✓ Every gate decision sits on an append-only record, and lessons promote into the standard.

In one sentence Own the best-practice standard, run it live so it coaches every person into their role at the right tier, and grade it independently on the same record, so the sanction holds and the capability stays with your team.

Start at any rung.

Many sites begin with a portfolio challenge or a sanction readiness review, because the payback is immediate. That first engagement is the front door to the rest of the line.

Book a walkthrough

capintelligence.ai

Call or text +1 832 900 2592