AFCEA submission intro (≤70 words):

> *Federal proposals cost six figures to write, yet most companies win only ten to twenty percent of what they bid. The reason usually isn't weak writing — it's chasing bids they can't win and never reading their own proposal the way an evaluator will. The answer isn't to bid more or write faster, but to bid the right things and score them honestly against the criteria.*

Submission image: `image-2-industry.png` (red-pen markup of a project proposal)

You Can't Grade Your Own Proposal — But Someone Will

By David Kindle, Director of Engineering, Falcon Logic

Federal proposals are expensive to write and hard to win. A single major bid can consume six figures in bid-and-proposal cost and weeks of senior talent, and the commonly cited industry win rate still hovers somewhere between ten and twenty percent. Do that math across a full pipeline and an uncomfortable truth emerges: most of what a company spends on proposals it will not win. The instinct is to fix this by writing better or bidding more. Both miss the point. Most federal bids are lost before a word is written — in two decisions that happen long before submission.

The first loss: bidding the wrong things

The first is a discipline problem. Faced with a full pipeline, many companies treat every opportunity as a pursuit and spread their capture and proposal resources evenly across all of them. The result is predictable. The B&P budget thins out, the capture team burns out, and the proposals that had a real chance get the same tired attention as the long shots that never did.

The contractors who win consistently do the opposite. They say no early and often, and they reserve their best effort for the handful of bids they can actually win. Probability of win is not a vanity metric; it is a capital-allocation decision. Every hour spent chasing a fifteen-percent opportunity is an hour stolen from a sixty-percent one. A focused pipeline of genuinely winnable pursuits beats a bloated pipeline of hopeful ones every time — and the discipline to walk away is what makes the pursuits you keep worth winning.

The hard part is that bid/no-bid decisions are too often made on enthusiasm rather than evidence. A team that is excited about an opportunity will find reasons to pursue it. A disciplined process forces the opposite question: what does the record actually say about our fit, our past performance, our incumbency disadvantage, and the competition — before a dollar of capture money is committed?

The second loss: never reading your own proposal the way it will be read

The second decision is subtler and, in my experience, more costly. Companies write proposals from the inside looking out. The evaluator reads them from the outside looking in.

That difference is everything. The proposal team knows what it meant to say, so it reads its own draft charitably, filling gaps with context that lives in the room but never made it onto the page. The government evaluator has none of that context. They have Section M, a rating scale, and a stack of competing proposals to get through under deadline. They credit what is on the page against the stated criteria — nothing more.

This is exactly where awards are won and lost, and the government's own record proves it. Year after year, the leading grounds on which the GAO sustains a bid protest are unreasonable technical evaluations and unreasonable cost or price evaluations — assessments that could not be reconciled with the solicitation's criteria (GAO Bid Protest Annual Report to Congress for Fiscal Year 2025, GAO-26-900695). The same gap that gets an evaluation overturned is the gap an offeror should have caught first: a claim that isn't substantiated, a requirement addressed in spirit but not in the words on the page, a strength the author sees clearly and the evaluator never will. The offeror who reads their own draft the way the board will is the one who finds those gaps while there is still time to fix them.

Why the easy AI answer is the wrong one

It is tempting to hand a draft to a general-purpose AI and ask whether it is any good. Don't confuse that with an evaluation. A model that returns “this reads well” without mapping each claim to a specific requirement in Section M is producing optimism, not analysis — the very thing that already sinks too many proposals. Worse, feeding a proprietary win strategy into an uncontrolled service trades a competitive edge for a convenience.

Used well, AI on the industry side should be held to the same standard a real evaluator is held to. Score the proposal against the actual evaluation criteria, not a general impression. Read it outside-in, simulating how the board will weigh it, and surface the weaknesses while the deadline still allows a fix. Tie every identified gap to the specific passage that causes it, so the finding is actionable rather than vague. Ground the bid/no-bid decision in evidence and document the reasoning. And keep proprietary material inside a controlled environment, never in a service whose data handling you cannot inspect.

The goal isn't to bid more — it's to bid right

The measure of a capture operation is not how many proposals it produces or how fast it produces them. It is how reliably it commits scarce resources to the bids it can win, and how well those proposals survive first contact with an evaluator who owes them no benefit of the doubt. AI can help on both fronts — but only if it is disciplined enough to tell you what you need to hear rather than what you hope to.

You will never get to grade your own proposal. Someone else always will. The winning move is to see it through their eyes first.