Exec Pitch Kit

The pitch outline, the ROI one-pager skeleton, and the objection handlers. Everything here assumes the Measurement Workbook (Module 6) is feeding you numbers.

The pitch, in seven slides

1. **The problem in their language.** Ticket costs rising; pressure to "do AI"; and [if applicable] a KCS investment whose assumptions — humans write, humans read — expired. One slide, no methodology talk.
2. **The one-sentence model.** *AI drafts, humans approve, machines and customers consume, and the system tells us what's missing — without asking agents to become writers.*
3. **The economics.** The shift-left spectrum with YOUR channels and YOUR unit costs (Channel Inventory tab). The ask lands here: one captured solution answers hundreds of times at pennies.
4. **What we measure.** Blended cost per answer (down) and shift-left ratio (up) — plus the metric inversion warning, delivered *now*, in advance (Inversion Brief tab, verbatim if you like).
5. **The 90-day pilot.** Scope, people, success bar from the signed charter. Emphasize what agents are asked for: 30 seconds per case. Nothing else.
6. **Governance is the moat.** One slide on the gate: nothing AI-written enters the record without a human. This is the slide that separates you from "we bought a bot."
7. **The ask.** [Curator hours, sponsor time, $X tooling]. Small numbers, next to slide 3's leverage.

The ROI one-pager (send-ahead or leave-behind)

Pull four numbers from the workbook and stop: current blended cost per answer · projected at target shift-left (Scenario Modeler) · monthly savings at scenario · curator investment (Governed CPA tab — the pennies line). One chart: blended CPA trend. One sentence of method: *"all-channel answer economics; no deflection attribution assumptions."* Footnote the illustrative-vs-actual status of every figure honestly — credibility compounds.

Objection handlers

- **"We tried KCS; it fizzled."** It fizzled because of the capture tax — asking solvers to write mid-queue. That labor is now near zero: AI drafts from the closed case; humans approve in minutes. The part of KCS that was right (knowledge as a byproduct of solving) is the part we keep. *(If they're certified: nothing you invested is wasted — the principles carry; the machinery updates.)*
- **"Can't we just buy an AI bot?"** A bot on an ungoverned KB is garbage in, confident garbage out — polished wrong answers, cited, at scale. The vendors sell the smart edge; the value is the boring core it stands on. We're building the part no vendor ships.

- **"Agents won't do it."** We're asking for a 30-second flag, not authorship. The pilot's charter caps solver effort and *kills the pilot* if it creeps — put that in front of the skeptical team lead and watch the objection dissolve.
- **"We can't afford curators."** Governed CPA tab: curation runs pennies per answer against the dollars each shifted answer saves. The alternative isn't zero cost — it's uncurated answers whose worst confident error lands on a customer.
- **"Our KB is too far gone."** That's a Level-1 placement, and the plan starts there: type it, own it, retire the fifth (20%) that's dead — weeks, not quarters (Modules 2, 3, 5). "Too far gone" is the reason to start, not to wait.
- **"TTR is getting worse — the program is failing."** *(You'll hear this at month 4–6.)* The Inversion Brief, verbatim: the easy volume left the queue; that's the strategy working. This objection is why slide 4 delivered the warning in advance.
