

The AI Ownership Scorecard

For Field-Service SaaS Leaders (\$3M–\$15M ARR)

Six questions. Five minutes. One honest answer about who's actually deciding what gets built.

Most field-service SaaS companies in this range have somebody unofficially handling AI decisions, usually the most technical person in the room, squeezed in between their real job. That's not a strategy. It's a Slack thread with a mandate problem.

Score each answer below (A=1, B=2, C=3, D=4), then check where you land on the back.

1 When someone suggests a new AI use case, what happens next?

- | | |
|--|--|
| A) Dies in Slack (1) | B) Gets a scrappy prototype on someone's own time (2) |
| C) Goes to whoever's unofficially "in charge" (3) | D) Runs through a standing process tied to a number you track (4) |

2 Where has AI actually shipped, not just been discussed?

- | | |
|---|--|
| A) Nowhere (1) | B) A customer-facing chatbot (2) |
| C) A mix of tools, added as they came up (3) | D) Dispatch, scheduling, or field ops, where the margin lives (4) |

3 Has your engineering process changed because of AI, or just the tools sitting on top of it?

- | | |
|--|---|
| A) Same process, added a Copilot license (1) | B) A few engineers use it individually (2) |
| C) Rebuilding parts of it, inconsistently (3) | D) Development and discovery genuinely run on AI now (4) |

4 If your board asked for AI ROI, could you answer with a number?

- | | |
|--|---|
| A) No, we'd talk around it (1) | B) We'd point to adoption stats, not dollars (2) |
| C) Rough sense of impact on one or two things (3) | D) Yes, tracked like any other investment (4) |

5 Who owns the decision to fund or kill an AI initiative?

- | | |
|---|---|
| A) Whoever's most excited that quarter (1) | B) The founder, when they have time (2) |
| C) The VP of Eng, as one more thing on their plate (3) | D) Someone with an actual mandate to own tech strategy (4) |

6 How often does someone formally revisit your AI priorities?

- | | |
|---|---|
| A) Never, only when an idea comes up (1) | B) Occasionally, triggered by a competitor or a board question (2) |
| C) Roughly quarterly (3) | D) On a fixed cadence, regardless of what's in the news (4) |

Where You Land

SCORE	BAND	WHAT IT MEANS
6–11	Reactive	You're funding AI on volume, not impact. Expensive, not unusual at this stage.
12–16	Aware	Somebody's paying attention. No consistent process yet. One busy quarter from drift.
17–20	Systematic	Real structure exists. The gap is surviving a hiring freeze or a rough quarter.
21–24	AI-Native	Rare at this size. Your job now is protecting the discipline as you scale.

Before you fund anything else labeled "AI": pick one number already living on someone's dashboard, cost per work order, technician utilization, average handle time, and start tracking it now. Without a baseline, every future claim about AI's impact is a guess wearing a suit.

If you scored Reactive or Aware: the gap between knowing that and fixing it isn't more research. It's someone with the mandate to turn the research into decisions. That's the work I do for field-service SaaS companies at your stage, real ownership of the AI-review function, without the overhead of a full executive hire.

[Book Your Free Consultation →](#)

Bring your score. We'll spend 30 minutes on where you actually stand.