

Can you defend your AI spend?

Twelve questions. Three minutes. No email gate.

Score each line **0** = no / don't know · **1** = partly, or it's out of date · **2** = yes, and it's written down. Be honest — nobody sees this but you. Total at the bottom.

Area	Question	What a 2 looks like	0	1	2
Inventory	Can you produce, today, a list of every AI tool, seat and bundled vendor feature you pay for — with a cost and a named owner on each line?	One list, under a month old, every row has a name.			
Inventory	Do you know which of those tools have a model provider in the path of a customer-facing or revenue-affecting workflow?	Each workload is mapped to its provider and model version.			
Shadow AI	Do you know which unapproved AI tools your people are using, and what data is going into them?	A current map exists and someone reviews it.			
Shadow AI	Is there a sanctioned path that does what the shadow tools do — so people don't have to route around policy to get work done?	Yes, and usage of the sanctioned path is growing.			
Value	For your three largest AI line items, can you show measured savings or revenue — not the vendor's projection, not a survey of user sentiment?	Measured against baseline time and cost, in writing.			
Value	Has anyone re-run the build-vs-buy or rent-vs-own math on your AI infrastructure in the last twelve months?	Yes, with the assumptions written down.			
Accountability	For each AI system in production, is one person named as owning the outcome — not the vendor relationship, the outcome?	A name per system, reviewed on a cadence.			
Accountability	Is it written down who may ship an AI system despite a security or quality finding, by workload type?	One page, signed, by workload type not job title.			
Resilience	If your primary model provider went dark for four hours this morning, which workflows stop — and do you know that from a drill or from a guess?	Tiered list, drilled in the last quarter.			
Resilience	Have you read the actual SLA for the AI service you call — what uptime it promises, what it excludes, and what a breach pays you?	Yes; gaps are listed and were raised with the vendor.			
Decision	Do you have a keep / expand / kill decision on record for each AI workflow, with the arithmetic?	Yes, reviewed at least twice a year.			
Decision	Is there a sequenced 30/60/90 plan for the AI estate with an owner on every action?	Yes, and last quarter's plan was mostly executed.			

Total (out of 24)

—

—

Score	Reading	What it means
0–8	Exposed	You are running AI on trust. Nobody could answer your CFO or your examiner today. Start with the inventory — it takes a month and changes every conversation after it.
9–15	Partial	You have pieces — usually an inventory and some ownership — but no evidence chain from spend to outcome, and no tested answer to absence. This is where most mid-market organizations sit.
16–20	Governed	The structure exists. The gaps are usually value measurement and resilience drills. Close those and the AI estate becomes defensible at board level.
21–24	Defensible	You can show what AI costs, what it returns, who owns it, and what happens when it's gone. Keep the cadence; most of the value now is in re-running the math annually.

Scored under 16? That is what the AI Value Diagnostic fixes — in two weeks, for \$3,500.

Every question above maps to one of the six artifacts you get back: the inventory, the shadow-AI map, the value scorecard, the keep / expand / kill list, the accountability structure, and the 30/60/90 plan. Fixed scope, fixed fee, no production access, about five hours of your team's time. If the answer is that you don't need help, it says so in writing.

yasinjabbar.com/diagnostic.html · hello@yasinjabbar.com · Charlotte, NC · on site regionally, remote nationwide