

WHITEPAPER

The quote is the constraint. The estimator is the risk.

For owners of specialty manufacturers in the Delaware Valley

Signet, a division of Circle Square Consulting

August 2026

COVER LINE

This document decides whether your shop takes the intake, extraction and first draft of a quote off one estimator's desk this year, what that costs, and what the agent will never be allowed to do.

Executive summary

Your RFQ inbox moves at the speed of one or two people. Everybody in the building knows their names. Day to day that is a bottleneck; over five years it is a succession problem, and it is the same man.

Break a quote into its nine steps and two are judgment: estimating the routing, and setting the price and the date. The other seven are reading a PDF, pulling part data, finding what the last similar job actually cost, checking material, formatting, sending and logging. Those seven fill the afternoon so that the skilled work happens Thursday.

We build an agent that does the seven. It reads the RFQ and the attachments, extracts part numbers, revisions, quantity breaks, material and finish callouts, required dates and any flow-down buried on page two, pulls your own prior job cost alongside, and assembles a draft with the price and lead time left blank. Your estimator does his two steps and signs. Every quote leaves with a human name on it.

The line, fixed in writing before anyone builds: the agent does not price the job, does not commit a lead time, does not accept a change order, and does not touch the shop floor, a machine, a traveler in process, or any quality disposition. You approve that list before we start.

Price. The Quote Clock is free, thirty minutes. Agent Blueprint **\$4,900**, two weeks. Agent Build **\$10,000**, or **\$14,900** for both. Each additional agent **\$9,900**. Where the ERP is the risk, Claude Setup Connect is quoted separately at **\$3,900**. Agent Care **\$750 per agent per month**. No build without a Blueprint, and no fixed fee against an ERP API we have not tested.

The situation

Two of your three levers are gone.

Input costs. In NAM's Manufacturers' Outlook Survey for Q2 2026 (215 respondents, fielded May 12 to 28, 2026), 83.1% named raw material costs their top challenge, up 25.6 points from 57.5% in Q1. 94.4% expect input costs to keep rising over the next twelve months and 59.6% expect increases above 5%. Trade uncertainty was cited by 71.8%, healthcare costs by 76.8%, energy by 72.0%.¹

Labor. In the same survey, attracting and retaining talent was cited by 46.95%, average unfilled job rate 4.25%. The CADDi/SME 2026 American Manufacturing Survey puts it harder: 79% call the skilled-labor shortage the single biggest challenge heading into 2026, and 69% plan capital investment in machinery and robotics while software investment declines.² Methodology is not fully published. Directional.

Meanwhile demand is up. ISM's Manufacturing PMI hit 55.6% in July 2026, up from 53.3% in June, the highest since May 2022, seventh straight month of expansion, New Orders 56.7%, Employment 52.8% and the first expansion in 33 months.³ More RFQs, same office, same estimator.

Two more things arrived while you were busy. Manual quoting for a single part can take up to two hours.⁴ And cyber and export-control clauses are being written directly into supplier quality agreements as AS9100, IATF 16949, NIST 800-171 and CMMC converge.⁵ That one lands on your Quality Manager, not on IT.

What most shops miss

Most owners think the constraint is estimating. It is not. The constraint is everything wrapped around estimating.

Watch the clock rather than the man. Sort last month's RFQs by sit time and you get a fast group, a middle group, and a tail sitting eight, twelve, nineteen days. Cut the same list by part family and the number that should stop you is the median on **repeat** parts: parts you have quoted before and made before, with cost history on file. If a repeat part sits three days, nobody was estimating it for three days. It was in a queue. Queue time is not skill, and buying skill will not fix it.

Here is the position, and you are welcome to disagree with it. **Most of what is sold into this market automates the wrong half.** The quoting tools that lead the category exist to suggest a price: a should-cost model, a number in a box the estimator can accept with one click. That is the half you should never give away. It is the thirty years you are actually paying for, it is the part a customer will hold you to, and a suggested price with a click-to-accept button is a price nobody decided. Give away the reading and the typing. Keep the judgment. A shop that automates pricing and keeps the typing has it exactly backwards.

The second thing missed is what the asset is. The agent is software and software depreciates. The asset is the context pack: how your shop actually quotes, written down. Which customer always says 500 and orders 200. Which material certs that aerospace account rejects. Which print has a tolerance drawn tighter than the part needs because a young engineer copied the block. Which flow-down clauses matter and which are boilerplate. None of that is written anywhere today. It is in one head, and that head is sixty-three.

The third: nobody has a number. **There is no independently audited industry average for RFQs per shop per month, for quote win rate, or for the dollar cost of slow quoting.** It does not exist. Anyone handing you one is quoting a vendor's survey of that vendor's own customers. We would rather say that than sell you a benchmark you cannot check.

What good looks like

A standard you can hold any vendor to, including us. Ten items. You can implement most of it without hiring anybody.

1. The clock is measured before anything is bought. Thirty days of RFQs: in, out, sit time, who quoted it, whether it went out at all. Report the median, the 90th percentile and the worst five. The average is dragged around by the tail and tells you nothing you can act on.

2. RFQs arrive at a shared mailbox. If they land in one estimator's personal inbox, that is a single point of failure with a name and a birthday, and it is a twenty-minute fix.

3. Two of nine steps stay human. Routing and price-plus-date. Write down which steps are being automated and which are not, before anyone quotes the work.

4. **The escalation list exists in writing and the agent follows it without discretion.** At minimum: a price discrepancy against your quote or the blanket, a requested date earlier than you promised, a part number you have never made. Each is a commercial decision dressed as data entry. Add your own: a new customer with no credit terms, a revision letter you do not hold, ITAR language.
5. **The output is a draft, not a send.** Nothing reaches a customer without a person's name on it.
6. **There is a named reviewer inside your company.** Not a queue. A person. If nobody internal will own the review step, do not buy the agent.
7. **Two strikes and it is law.** Anything the reviewer corrects twice gets written into the context pack permanently, so the same correction is not made a third time.
8. **The ERP question is settled first, in writing.** Public API documentation, an OpenAPI spec, a sandbox, documented authentication, webhooks or polling. Five checks, five yes-or-no answers.
9. **Everything is logged and you approve the refusal list.** What it did, what it escalated, what it could not read.
10. **The Blueprint is the off-ramp.** If it is not worth automating, you stop there and you keep the map and the context pack.

A shop that does items 1, 2, 4 and 6 with no software at all is better run than one that buys a tool and does none of them.

How we do it

ENGAGEMENT	DETAILS
The Quote Clock: free, thirty minutes.	Your last thirty days of RFQs, from your inbox and your ERP: how many came in, how many went out, how long each sat, how many never got quoted. Then we tell you whether an agent could take the intake and the first draft off the estimator's desk, or whether it could not. You keep the numbers. No second call unless you ask.
Agent Blueprint: \$4,900, two weeks.	We map quoting as it actually runs, sit with the estimator until the answers stop being obvious to him, write the extraction rules per customer, define the escalation and refusal lists, test what your ERP will accept, and produce a build specification with a fixed price on it. If the ERP will not cooperate, we tell you and you stop, and you keep the map.
Agent Build: \$10,000, or \$14,900 with the Blueprint.	The agent, a runbook, a governance record, a named engineer, thirty days of tuning. Second agent (order entry, or Books for AP/AR and month-end): \$9,900. Agent Care \$750 per agent per month afterwards.

THE ERP, HONESTLY, BECAUSE IT CHANGES THE PRICE.

Epicor Kinetic has documented, actively maintained integration tooling and NetSuite has a mature API; on either, integration is a known quantity. **E2 Shop System from ECI received a third-party Grade D API assessment:** no public endpoint reference, no OpenAPI spec, no sandbox, no documented authentication flow, and a QuickBooks integration repeatedly reported broken.⁶ **JobBOSS², also ECI, shows the same pattern.** Global Shop, Infor, ProShop and Made2Manage were not part of that assessment and we will not invent a grade for them. Where integration is the risk we quote **Claude Setup Connect at \$3,900** separately, so you see the number instead of paying for hidden risk inside a build fee. **We will not quote a fixed fee against an API we have not tested.** Not for E2, not for Kinetic.

NOT INCLUDED.

Shop floor automation, machine vision, anything on the OT side. Scheduling your floor. ERP replacement or implementation. CMMC certification assessment (we are not a C3PAO). Marked up model, platform or OCR consumption; pass-throughs bill at cost.

Questions

QUESTION	ANSWER
"Why not our existing IT provider?"	Often the right answer. If they can hand you an escalation list, a refusal list and a named reviewer, use them. Most cannot, because keeping a network up and writing down which decisions software is forbidden to make are different jobs. We do both: Circle Square is the MSP, Signet is the side that builds and signs for the agents.
"Our last ERP implementation was a disaster. Why is this different?"	Reasonable. That is why the build is a fixed fee against a spec you approve, and the Blueprint has an off-ramp. We are not replacing your system of record. We are reading from it.
"What if the AI gets a quote wrong and I lose the customer?"	It will get things wrong. That is why it does not price, does not promise a date and does not send. Your estimator gets a draft with the sources shown and a list of what it could not read. A wrong extraction is visible in ten seconds. A wrong price is not.
"Should we buy a machine instead?"	If your constraint is spindle capacity, yes: buy the machine. This is the wrong purchase and we will say so. If it is RFQs sitting in an inbox until Thursday, no machine fixes it.

Mike Cervino · Circle Square Consulting · Radnor, PA · *AI does the typing. A named human makes the judgment calls.*

Footnotes

- 1 NAM Manufacturers' Outlook Survey, Q2 2026, published June 2026.
- 2 CADDi/SME 2026 American Manufacturing Survey. Methodology not fully published. Directional.
- 3 ISM Manufacturing Report on Business, July 2026.
- 4 *Modern Machine Shop*, November 1, 2021, "When it comes to RFQ, response time is money."
- 5 *Quality Magazine*, 2026, on ISO 9001, AS9100, IA9100, IATF 16949, NIST and CMMC convergence.
- 6 supergood.ai API Report Card, E2 Shop System.