✦ **Holaris** ADVISORS

✦ AN OPERATOR BRIEF

AI implementation *for operating leaders.*

Choosing vendors, asking the right questions, and setting honest ROI expectations. Includes the twelve questions every operator should ask an AI implementation consultant before signing a master services agreement.

WRITTEN FOR CEOs, CFOs, COOs, and owner-operators of $5M–$100M companies

READING TIME 12 minutes

PREPARED BY Charlie Moberly, Managing Partner, Holaris Advisors, Reno, Nevada

✦ 01 · The Setup

Most AI initiatives in the middle market fail before the contract is signed.

By the time you're evaluating implementation partners, the outcome has largely been decided. The vendor was picked for the wrong reasons, the questions weren't asked, and ROI was promised in a range that nobody on either side actually believed.

This brief is built for operating leaders of small- and middle-market companies of your size, so you can run that evaluation differently. It covers the four kinds of vendors competing for your budget, the twelve questions to ask before you sign, and what honest ROI looks like in the first 12 months.

73 %	30 %	90 days
Of \$5M–\$100M companies have adopted AI in some form by 2026.	Report successful implementation. The gap is not access to technology, it is execution.	Is the right window to prove a Quick Win against a measured baseline.

✦ 02 · The Four Kinds of Vendors

Know who is sitting across the table.

Most AI implementation pitches come from one of four categories. Each has a structural bias. Recognizing the category clarifies the bias, so you can weigh the evidence accordingly.

TYPE 01

The Big 4 and large systems integrators

Deloitte, Accenture, EY, KPMG, and the global SI tier. Built for Fortune 500 complexity. Engagements typically run $400K–$1.5M with 12+ month commitments and large junior staff teams.

WHAT YOU'LL HEAR

Enterprise transformation language, multi-phase roadmaps, governance frameworks, and a partner who shows up at the kickoff and disappears afterward.

TYPE 02

AI tool vendors and reseller networks

Direct sales from model providers, plus the partner networks that resell them. Compensation is tied to seats and consumption, not to whether the tool gets adopted or generates revenue.

WHAT YOU'LL HEAR

Demos before diagnosis, "platform plays," roadmap features that aren't yet shipped, and recommendations that conveniently align with the vendor's own catalog.

TYPE 03

Solo and boutique consultants

Independents charging $200–$400 per hour, often with strong individual expertise but no bench. Quality varies widely. The best are excellent; the median is uneven.

WHAT YOU'LL HEAR

Pricing by the hour rather than by the result, strong opinions on tools without an underlying framework, and capacity constraints that surface mid-project.

TYPE 04

Generalist consulting firms with an AI practice

Firms that bolted AI onto an existing operations, technology, or marketing line of business. The intent is right; the depth often isn't. Long client lists usually mean junior staff doing the work.

WHAT YOU'LL HEAR

AI framed as a feature of a larger engagement, retrofitted frameworks from other practice areas, and senior partners pitching while staff consultants deliver.

✦ 03 · The Twelve Questions

What every operator should ask before signing.

These are the questions a CFO would ask a CFO. None of them are about technology. All of them are about how the engagement will actually run, so you can tell who has done this work before and who is still selling it.

> *Use this as a live scorecard during a vendor pitch. A strong partner will welcome these questions and answer most of them without notes. A weak one will deflect, generalize, or promise to "come back to that."*

01 Walk me through prior AI implementations you have personally delivered, and what each one did to the operating metric that mattered.

Specifics are the test, not logo counts. Names can be confidential, but the operating story (the function, the baseline, the result) should be vivid and consistent. A partner with a handful of fully documented implementations tells you more than one with a long client list and thin detail.

LISTEN FOR *Concrete metrics on the prior baseline and the post-implementation number, and a clear answer on which functions the partner has actually shipped (operations, finance, customer service, sales) rather than industry-by-industry claims.*

02 Who specifically will be doing the work, and what is their utilization on my project?

The pitch team is rarely the delivery team. Get the name, hours per week, and the résumé of the person who will be in your operations on a Tuesday afternoon.

LISTEN FOR *Hedging, "we'll staff appropriately," or a partner whose calendar disappears the week after kickoff.*

03 What is your diagnostic framework, and can you show me a sample output from a prior client?

If they don't have a framework, every engagement is improvised. If they have one but can't show it, the framework is marketing rather than methodology.

LISTEN FOR *A named, repeatable diagnostic with visible scoring axes, not a slide titled "Our Approach."*

04 How will we measure success, and what is the baseline today?

Before any tool is selected, you need agreed-upon metrics, a baseline, and a measurement window. If those three things aren't in the SOW, the ROI conversation is theater.

LISTEN FOR *Willingness to instrument the baseline before kickoff, and a measurement plan that survives staffing changes on either side.*

05 What is your fee structure, and what triggers a change order?

Hourly billing rewards drag. Fixed fee rewards scope discipline. Either model can work, but you need to know which one you're buying, so you can plan accordingly.

LISTEN FOR *A clear answer on scope boundaries and a written change-order process, not "we'll figure it out as we go."*

06 Are you compensated by any tool vendor whose product you might recommend?

Reseller commissions distort recommendations. The question isn't whether commissions exist; it's whether the consultant will tell you about them before you ask.

LISTEN FOR *A direct yes or no, and a written disclosure of any partner relationships in the engagement letter.*

07 Who owns the models, data, prompts, and integrations after the engagement ends?

IP terms in AI engagements are often left vague on purpose. Vague IP terms favor the vendor. Clarify ownership of training data, derived models, prompt libraries, and any custom code.

LISTEN FOR *Plain-English IP clauses in the MSA, including portability if you switch partners.*

08 What does your engagement look like if a Quick Win fails to hit baseline?

Mature partners have a written answer to this. The answer might involve fee adjustment, scope reset, or a structured post-mortem, but it exists. If they've never thought about it, they've never had one fail, which is itself a flag.

LISTEN FOR *A defined off-ramp, not "that won't happen with us."*

09 How do you handle change management with the team whose work the AI will change?

Most failed AI implementations are adoption failures, not technology failures. If the answer is "training" and nothing else, the engagement will produce a tool nobody uses.

LISTEN FOR *Specific adoption mechanisms, named change sponsors, and a plan for the months after go-live.*

10 What insurance do you carry, and what is your indemnification posture on AI errors?

Errors & omissions, cyber, and clear indemnification language matter more in AI than in conventional consulting, because AI errors are harder to attribute and easier to spread.

LISTEN FOR *Real policy limits ($2M+ E&O is standard for serious firms) and a contract that doesn't carve out AI as an exception.*

11 What references can I call without going through your sales team?

Curated references rehearse the same answers. Ask for two clients whose engagements ended within the last 18 months, and call them directly. The conversations you have without a chaperone are the ones that tell you the truth.

LISTEN FOR *A willingness to share more than two names, including at least one where the engagement was hard.*

12 What will you walk away with at the end of the engagement that you could not have walked away with before?

This is the most important question, and the one most consultants haven't thought through. The honest answer names a capability your team will own afterward, not just a tool that will be running.

LISTEN FOR *An answer that ends with a sentence about your team, not the consultant's team.*

✦ 04 · Honest ROI Expectations

What good looks like, and what is being oversold.

Operating leaders are not asking for the moon. They want to know what the next twelve months will actually return, so they can plan budget, headcount, and timing around something real.

WHAT HONEST ROI LOOKS LIKE

Modeled in operating units, not vendor language.

✦ A measured baseline captured before the engagement starts, signed off by both sides.

✦ A single Quick Win targeted at one workflow with a measurable output (cycle time, error rate, cost per transaction).

✦ A 90-day window to first ROI, with monthly check-ins against the baseline and an honest mid-point reassessment.

✦ A range of expected outcomes, not a single number. Year-one returns in the 2x–4x range on a well-scoped Quick Win are credible. Higher claims require receipts.

✦ Savings expressed as cost avoided, hours redeployed, headcount avoided, or revenue generated. Each one should map to a line on a P&L that your CFO recognizes.

WHAT GETS OVERSOLD

Marketing math dressed up as operating math.

✦ "10x productivity" framed as a return without naming the workflow, the baseline, or the measurement method.

✦ Industry benchmarks ("companies see X% lift") applied to your company without adjustment for size, data quality, or operating model.

✦ Soft savings (employee satisfaction, brand uplift, future optionality) presented as if they were hard dollars.

✦ ROI projected from pilot data extrapolated linearly across the enterprise, ignoring the cost curve of scale.

✦ Numbers that were never written down before the engagement started, which makes the post-engagement number impossible to dispute.

✦ 05 · A Realistic Twelve Months

What a well-run engagement looks like, month by month.

There is no universal timeline, but a well-run middle-market AI engagement typically follows the same phases. If the partner's proposed timeline skips any of them, the proposal needs another draft.

PHASE 01 Weeks 1–4	**Diagnostic and baseline.** Identify the candidate functions, score them against a structured framework, instrument the baseline metric, and record it. No tools selected yet.
PHASE 02 Weeks 5–8	**Quick Win design and vendor selection.** Solution architecture for the single highest-priority workflow. Vendor evaluation against criteria written before the demos begin. Integration plan.
PHASE 03 Weeks 9–12	**Pilot and production launch.** Implementation, integration, training, pilot, production cutover, and first reading against baseline. This is when the 90-day ROI claim either holds up or doesn't.
PHASE 04 Months 4–6	**Quick Win #2 and infrastructure for the Strategic Bet.** Launch a parallel Quick Win in a second function. Begin investing in data and infrastructure for the higher-value, higher-difficulty initiative identified in the diagnostic.
PHASE 05 Months 7–12	**Strategic Bet launch and capability handoff.** Strategic Bet enters production if Quick Wins are performing. Internal team takes ownership of operating, monitoring, and iterating. Consultant scope contracts.

✦ 06 · Red Flags

Eight signals to walk away.

If you encounter any of the following during the evaluation, treat it as information. The earlier you walk, the cheaper the lesson.

✦ *Flag 01*

The pitch leads with technology rather than the operating problem.

A partner who can't articulate your operating pain in your language hasn't done the homework, and won't.

✦ *Flag 02*

There is no named methodology, or the methodology was clearly built for this meeting.

Without a repeatable framework, every engagement is improvised, and the senior partner is the only thing keeping it together.

✦ *Flag 03*

References are curated to one or two clients, with sales staff on the line.

If the firm only has two clients willing to talk, you are the third experiment.

✦ *Flag 04*

ROI is quoted as a single round number without a baseline.

"3x return" is not a return. It's a press release.

✦ *Flag 05*

The proposed timeline skips the diagnostic and goes straight to tool selection.

This is how vendors disguise themselves as advisors. The tool was picked before the meeting.

✦ *Flag 06*

The MSA is silent on IP ownership, data portability, or AI-specific liability.

Silence in the contract is not neutrality. It is a default to the vendor.

✦ *Flag 07*

The team in the pitch is not the team that will be doing the work.

This is normal at the Big 4 and unacceptable in a middle-market engagement.

✦ *Flag 08*

The partner cannot describe what your team will own after the engagement ends.

If they can't see past the contract end, neither can you.

✦

"The best AI engagements leave your team with a capability they didn't have, not a dependency they can't unwind."

CHARLIE MOBERLY, MANAGING PARTNER, HOLARIS ADVISORS

✦ ABOUT HOLARIS

AI for operating leaders.

Holaris Advisors is an AI advisory firm built for the CEOs, CFOs, COOs, and owner-operators of small- and middle-market companies of your size. We help operating leaders identify, implement, and prove out the AI initiatives that actually move operational quality, cost, and margin. Every engagement is anchored by the AI Opportunity Map™, our proprietary five-axis diagnostic that prioritizes initiatives by operating impact rather than technology preference.

Led by Managing Partner Charlie Moberly, the firm is built on the conviction that AI buying is an operating decision before it is a technology decision. Our voice is an operator talking to an operator.

▸ **BOOK A 30-MINUTE CONSULTATION** · **REQUEST A FREE AI AUDIT**
