VERIDECT · EXECUTIVE BRIEF · TWO MINUTES

The five questions every executive asks — plain answers

Most buyers ask the same five things in the first ten minutes. Here they are in your words, with the short answer to each — and where each one is proven live at verident.ai rather than claimed.

1. "How does this help my company?"

The AI you're deploying now answers your customers and takes actions — approving, paying, sending, changing records. The exposure isn't that AI makes a mistake; it's that nobody independent checked before the mistake went out under your name. Veridect is that independent check, built in at the moment it matters: it **verifies what your AI says** (four frontier models from four vendors cross-examine the answer and return a confidence you can defend), **governs what your agents do** (before an action executes it is cleared, held for a person, or blocked — out-of-scope actions blocked, risky-but-authorized ones sent to a human), and **proves both** in a signed, tamper-evident record anyone can check. What that buys you: AI allowed into the workflows that matter — money, customers, patients, regulators — because there is an independent check and a record. *See it: veridect.ai/governance-suite governs a real action in its first minute.*

2. "How do I integrate it?"

One REST endpoint. Your engineers replace one AI call with a call to Veridect — typically a 10–20 line change — and the answer comes back verified, with the record attached. Agents connect the same way: a gateway in the path of MCP tool calls, framework adapters, or one call wrapped around each high-stakes action in a custom framework. It runs where you say — AWS, Azure, GCP, or on-premises — on your own keys and the model providers you already use. Nothing is ripped out, and swapping a model later is a configuration change, not a re-engineering effort. *Read it: the public integration guide, linked from veridect.ai/site#docs.*

3. "How long does it take?"

Today: run the live system yourself — nothing to install, no one's approval needed. **A private sandbox** — your own tenant, your policy thresholds, your data — takes a day once you say go. **Production is typically about two weeks**, and that clock runs mostly on your side: security review, key provisioning, access approvals. Our side is ready on day one. Once we're working together, a mutual NDA opens a private integration document built to your specs, followed by hands-on working sessions.

4. "How will this make me money?"

Three ways — and we won't invent a number for a business we haven't seen.

- **Revenue you've been holding back.** The highest-value agent work — anything that touches money, customer data, or a record — only ships when someone can sign off on it. An independent check plus a signed record is what lets it ship.
- **Cost you stop paying.** Routine calls settle on a fast quorum; the hard rules — money over a limit, personal data, health data, protected-class factors — run in plain code at zero model cost; your people review the exceptions, not every call. Four models isn't four times the bill: a few cents of extra checking on the one action in a thousand that could really hurt you — and one engine across departments replaces a vendor per function, so total AI spend nets down, not up.
- **Losses you don't take.** The out-of-scope action blocked before it executes; the risky one held for a human; the incident, fine, or breach that doesn't happen. When the audit or security review asks *was this genuinely checked?*, the evidence is already built — a signed record, not weeks of reconstruction.

Your own number comes out of the sandbox, on your own workflows.

5. "Why you, and not a big, well-known company?"

Because the big names can't stand where this layer has to stand. A model vendor can't be the independent check on its own model — that's a model grading its own work, and bundled "governance" is self-governance. The large security names check the edges — was the prompt safe, was the response safe — not whether the answer was right, whether the agent was allowed to act, or whether there's proof afterward. Nobody sold all of it as one control plane — verify, govern, prove, and see risk across a whole fleet — so we built it. The engine behind it is proven in education and in production under Fortune 100 contracts in regulated industries. The size question, answered by design: it runs in your environment, on your keys and your providers; the models are swappable, and so are the agent frameworks — the fleet and fabric layers govern many agent harnesses at once — one policy, one ledger, for every action sent through Veridect; every decision record verifies offline against a published key, with no call to us — your dependency on any vendor, including us, is bounded on purpose. And don't take the claim on trust: a credentialed Fortune 100 model-risk reviewer ran live adversarial scenarios against the production system and every one matched pre-stated behavior; a separate deterministic test suite holds 4,697 policy scenarios at 100% conformance; the open findings are published with the results at veridect.ai/validation-summary.

Bring these five to the call — we'll answer them on your workflows, in an environment you control. Live now: **veridect.ai/governance-suite** · Online: **veridect.ai/executive-brief** · Lisa Russell, CEO · lrussell@veridect.ai