



Five Questions to Ask Before You Put AI on Your Data

A field guide for the people who sign the check.

The market is pricing AI as if it all works. It does not yet, and not the way most organizations are doing it. Before you fund AI on top of your data, ask these five questions. If the honest answer to any of them is “we are not sure,” you are about to industrialize a problem, not solve it.

1. Ask five people to define our core term. Do we get one answer, or five?

Pick the entity your work turns on: customer, patient, asset, claim, case. Ask five people what it means.

Why it matters. AI cannot act correctly on a term the organization itself defines five different ways. The disagreement does not disappear when you add AI. It gets automated at scale.

A good answer. One definition, bound to the data as a machine-checkable constraint, not living in a slide deck or in someone’s head.

How SDC answers it: meaning is bound to the data when it is modeled, so one shared definition replaces the private ones people carry in their heads.

2. For any answer this AI gives us, can we trace it back to the source?

Take one output. Ask which source record, and which version of it, produced that answer.

Why it matters. An output you cannot trace to its source is an assertion, not evidence. It will not survive an audit, and it cannot be defended when someone challenges it.

A good answer. Every output carries its provenance back to the source record, and a third party can verify it without taking the vendor's word for it.

How SDC answers it: provenance is bound to the payload with a tamper-evident receipt an auditor can check independently.

3. Is our AI governance enforced at the moment of decision, or is it a document we hope people follow?

Why it matters. Only 28% of enterprises have a formal AI governance framework in production. The rest govern informally, which means inconsistently, or not at all. A policy document is not a control.

A good answer. A deterministic allow-or-deny at the moment of each decision, with a record. The same input yields the same verdict every time.

How SDC answers it: deterministic PERMIT, DENY, or INDETERMINATE verdicts at the point of use, each with a receipt. Governance is enforced, not asserted.

4. If we changed AI vendors or models next quarter, would our data keep its meaning and its governance?

Why it matters. If meaning and governance live inside a vendor's platform, your governance is on rented land. A sovereign or regulated environment cannot accept that.

A good answer. Meaning and governance are bound to the data in open standards, so they survive a change of vendor, a change of model, even a change of the underlying reference model.

How SDC answers it: an open-standards substrate built on W3C and ISO, sovereign-capable, and permanent. The data stays governed wherever it lands.

5. Are we building governance into the data now, or paying for it after the first incident?

Why it matters. Governance tends to get funded after the first incident, at roughly three times the cost of building it proactively. Putting AI on broken data does not fix the data. It industrializes the problem.

A good answer. The data substrate is in place before the AI goes on top of it, not bolted on after something breaks.

How SDC answers it: fix the foundation first. SDC is the substrate the AI stands on, built in, not painted on.

The point

In a gold rush you can buy the picks, or you can be the ground everyone has to dig through. The data layer is the least crowded position in AI and the one every disciplined buyer now knows they need. These five questions are how you find out whether you have it.

At Axius SDC we build the substrate that lets the answer be yes.

More information at <https://axius-sdc.com>