

7 Questions to Ask Before You Choose an Enterprise AI Platform

Written for teams building a shortlist. Every question here applies to any vendor under consideration, including us.

Most platform evaluations are decided on the feature list, because that is what vendors put in front of you. The commitments that are hardest to reverse sit underneath it: where your data has to live, who you have to hire, what you renew into, and what you keep if you change direction.

- 1 Where will our data have to live?**
- 2 Who owns what we build?**
- 3 What does our team have to learn before anything reaches production?**
- 4 Who staffs and supports this after go-live?**
- 5 What does this cost at production volume, not pilot volume?**
- 6 What happens when a model is deprecated or replaced?**
- 7 What happens in year three if we want to switch?**



1

Where will our data have to live?

WHY IT MATTERS

Migration is where timelines and budgets move first, and it decides which security, residency and privacy reviews you trigger before anything is built. A platform that requires ingestion into its own environment is a different project than one that runs against data where it already sits.

WHAT A GOOD ANSWER SOUNDS LIKE

A specific line between what works against data in place and what requires ingestion into the vendor's environment. A vendor who has done this before can draw that line without hedging, and can name which capabilities degrade if you keep data where it is.

▲ RED FLAG

"We support both." Ask what stops working in the version where your data does not move.

2

Who owns what we build?

WHY IT MATTERS

Data export rights are not the same as owning the intelligence layer. The ontology, the object types, the pipelines and the operational logic are where the value accumulates over three years, and on most platforms they are built inside the vendor's architecture. IP built on top of a licensed platform is licensed IP.

WHAT A GOOD ANSWER SOUNDS LIKE

Each artifact named separately, with ownership stated for each: the data, the ontology or semantic model, the trained models, the application code. Stated in the contract rather than the sales deck, along with the format each one leaves in.

▲ RED FLAG

"Your data is always yours." True, and not the question you asked.

3

What does our team have to learn before anything reaches production?

WHY IT MATTERS

Training is staff time, and it rarely appears in a pricing conversation. It also attaches to people rather than to the organization, so new hires train too and turnover means retraining. Palantir, as one example, runs role-based training tracks and separate professional certification exams that carry a cost, and notes that most courses require access to a Foundry environment provided by your employer.

WHAT A GOOD ANSWER SOUNDS LIKE

How many people, in which roles, for how many hours each, at what cost, and whether anything can reach production before that is complete. Ask what happens to that number as your use cases grow.

▲ RED FLAG

Being pointed at a free introductory badge or foundations quiz. Ask about the credential your team actually needs to operate the platform unsupervised.



4

Who staffs and supports this after go-live?

WHY IT MATTERS

Training is the entry cost. The standing team is the recurring one, and a data engineering group does not cover it alone. Palantir's published adoption documentation is the clearest example: its first phase names a Head of Program, Program Manager, Head of Data Governance, Data Lead and Ontology Lead, with the team growing across four phases as usage expands. Palantir presents these as suggestions, but the shape is the point. Several of those roles sit outside a data engineering group.

WHAT A GOOD ANSWER SOUNDS LIKE

An org chart with headcount, split into what the vendor's team covers and what yours does, based on comparable customers rather than a reference architecture.

▲ RED FLAG

“Your existing data team can handle it.” Ask which of the named roles they mean, and who owns data governance.

5

What does this cost at production volume, not pilot volume?

WHY IT MATTERS

Consumption pricing is compared at pilot scale and paid at production scale. The unit price looks competitive during evaluation, and the full cost only appears after deployment. Industry analyses consistently find that implementation costs, integration, data preparation and change management add substantially to stated platform pricing.

WHAT A GOOD ANSWER SOUNDS LIKE

The vendor models your top workflows at your expected volume, and again at three times that volume, using your numbers, in the room. Base fee and consumption named separately. Implementation costs drawn from real customer projects, with a day rate and a fixed-scope estimate in the proposal.

▲ RED FLAG

A per-unit price with no volume model offered, or “it depends on usage” without a proposal to work it out with you.

6

What happens when a model is deprecated or replaced?

WHY IT MATTERS

The model you evaluate is not necessarily the model you will be running in two years. Versions are deprecated and behavior changes, and either can degrade a production system without a single change to your code or your data.

WHAT A GOOD ANSWER SOUNDS LIKE

A committed notice period and a parallel evaluation window where both versions run, written into the contract. Ask what happened to customers the last time a version was retired.

▲ RED FLAG

Being pointed at a public policy page. Policies change without your consent. Contracts do not.



7

What happens in year three if we want to switch?

WHY IT MATTERS

Your leverage is highest before signature and lowest at renewal, once the platform is operationally embedded. Rebuilding an intelligence layer on a new system is a multi-month project, and the cost lands in a year when it was never budgeted.

WHAT A GOOD ANSWER SOUNDS LIKE

In writing: what exports and in what format, what has to be rebuilt, and roughly how long that takes for a customer your size. A vendor confident in the answer will put it in the contract as an exit provision.

⚠️ RED FLAG

An answer that is only about data export. The data was never the hard part.

IF YOU ONLY ASK ONE

Ask the last one, and ask for the answer in writing. Every other question here is a version of it, asked early enough to still have a choice.

WHO WROTE THIS

Cyberhill Partners builds rapidly deployable enterprise AI. It is ontology-driven, engineered for your environment on open standards, and deployed on the infrastructure you already run. Your data does not move. The ontology, code and models are yours.

We build this way because of where we learned to. Selected from 78 companies to lead AI at national scale for the Department of Defense, we have spent eight years building where every answer gets traced, challenged, and defended. The same standard runs through more than 1,000 complex enterprise implementations across the Fortune 500.

Wherever you are in the evaluation, we will show you what is working across enterprise AI right now, what is stalling, and how we would approach your use case.

Learn more at cyberhill.ai

Need more support? [Schedule an AI Strategy Session](#)