


THE AI INSTITUTE / OPERATING INTELLIGENCE

INFLUENCE & PARTNERSHIP RADAR · 02 OCT / 2026



Who Owns the Learning in Your AI Partnership?

Agree what can improve, who can reuse it and what your organisation keeps.

Boards, C-suite and senior management

Research period: September 2026 – October 2026

First published: 2 October 2026

Current to 2 October 2026

PUBLICATION RECORD

An Institute thesis and decision framework

Treat the learning produced by an AI partnership as a negotiated business asset: define permitted reuse, access to improvements and the expertise your organisation must retain.

WHAT LEADERS SHOULD TAKE FROM THIS PAPER

- 01
- Separate model training from other ways a service improves.
- 02
- Agree who can reuse jointly developed methods and tests.
- 03
- Give practising experts a funded role in the partnership.
- 04
- Keep enough evidence and know-how to evaluate the next supplier.

Research method

Original Institute synthesis of the Synopsys announcement, new arXiv research and official Singapore consultation material, checked 2 October 2026. Commercial questions and the learning-asset register are Institute recommendations. No supplier contract, customer data or private conversation was inspected or quoted.

First published 2 October 2026 · Updated 2 October 2026 · Research period September 2026 – October 2026 · Version 1.0 · Suggested citation: The AI Institute, *Who Owns the Learning in Your AI Partnership?* (2026).

CONTENTS

Inside this paper

01 Buy capability without giving away the learning
The decision

02 Learning is not one technical process
Define the asset

03 Make the learning arrangement explicit
Commercial terms

04 Give experts time and authority
Whose evidence matters

05 Test the knowledge where it will be used
What does not travel unchanged

06 Decide what remains after the pilot
Before the next signature

M Method and limitations

R References

How to use this brief

Start with the decision and use the framework to challenge the current plan. Figures retain their original populations and definitions. Institute analysis is labelled, and limitations are recorded before the references.

Buy capability without giving away the learning

Before signing an AI development partnership, ask what each party will learn from the work and who will be able to use it afterwards. That question is broader than whether confidential data trains a model. It includes the methods, tests and practical judgement developed while making the system useful.

On 30 September, Synopsys and OpenAI announced a multiyear agreement to develop GPT-Synopsys for chip-design workflows. The arrangement combines model development, licensed engineering tools and a shared revenue framework. Synopsys describes early customer engagements and a planned bundled service; it does not establish generally available performance or a measured customer return.¹

The announcement also states that customer data is not used to train the model. That is an important boundary, not a detail to omit in pursuit of a more alarming story.¹ The wider leadership question is how any specialist AI partnership distributes the value of what the parties build together. This article proposes questions for that negotiation, not allegations about undisclosed terms.

Learning is not one technical process

A service can improve without changing its underlying model. A team might refine an instruction, add a verification step, organise a document collection or write a better test. Alternatively, a separately agreed project might train a specialist model. Those activities have different inputs and produce different assets. A general promise about learning can conceal those distinctions.

The current research pipeline illustrates the difference. AREX-2 studies training agents on extended improvement sequences, while a separate self-evolving-harness paper studies changes to the surrounding software with the model held fixed. Both are preprints, with benchmark evidence rather than proven commercial outcomes.²³ They do not describe the Synopsys product or establish its design.

For a business sponsor, the implication is practical: name the improvement mechanism before negotiating it. Ask whether the pilot will produce reusable prompts, workflow code, evaluation cases, model adaptations or operating documentation. Identify what is supplied by the customer, what the partner already owns and what the team expects to create. Avoid treating all of it as an undifferentiated service.

Make the learning arrangement explicit

Start with a register of the assets the pilot is expected to produce. For each, ask who can inspect it, change it and use it after the pilot. Then ask whether the partner may use it for other customers. The commercial answer may differ by asset: a generic test method need not receive the same treatment as a confidential design rule.

Agree what happens when the supplier improves its service using permitted material. Does your organisation receive that improvement within the existing price, through a new licence or not at all? Can it retain the tests needed to compare alternatives? These are negotiating questions, not a recommendation that every customer demand ownership of a supplier's model.

Distinguish the right to retain something from the ability to use it. An exported workflow may depend on a licensed tool, an unavailable model version or undocumented knowledge. Before calling it transferable, ask a second team to explain and run an agreed example. A portability promise is more useful when its dependencies and costs are visible.

Counsel and procurement should translate the agreed business arrangement into terms appropriate to the jurisdictions and parties involved. This briefing does not determine ownership under any particular law. Its purpose is to prevent the operational questions from arriving only after the legal drafting is complete.

EXHIBIT / SEPARATE THE PARTNERSHIP'S LEARNING ASSETS

01 Inputs What may the project use, and for which purpose? Institute recommendation	02 Improvements Who may inspect and reuse the resulting methods? Institute recommendation	03 Retained capability What can your team still explain and operate? Institute recommendation
---	---	---

Give experts time and authority

The missing constituency is often the practitioner whose judgement makes the workflow reliable. A senior engineer, nurse, claims specialist or teacher may know which apparently minor exception changes the answer. Invite that expertise into the design of the tests, not merely the final demonstration.

Give those contributors paid time, an accountable manager and a way to record disagreement. If their corrections become part of a shared workflow, make that process visible. Do not describe ordinary operational participation as permission for unrelated reuse of personal or confidential information.

A mutually useful partnership can give the supplier a clearer product problem and the customer a better service, stronger tests and more capable staff. The balance need not be symmetrical. But both parties should be able to state what they gain, and the customer's experts should not disappear from the operating plan once the pilot has been approved.

Test the knowledge where it will be used

A method learned in one setting can be inappropriate in another. A new preprint on context confusion reports that training on acceptable behaviour in one domain can induce unsuitable behaviour elsewhere. Its experiments are bounded research, not evidence that every adapted model becomes unsafe.⁴ The useful question is which adjacent tasks need retesting after a change.

The US-origin chip-design announcement concerns a specialist commercial ecosystem with established engineering tools. It is not evidence of equivalent access, bargaining power or results in every country. A smaller institution without its own evaluation team may need a university or independent practitioner to help specify what it should retain.

Legal uncertainty also varies. Singapore's current consultation on AI and intellectual property runs to 22 October; it is a request for views, not an enacted allocation of rights.⁵ Do not export a contract assumption from that debate to Europe, the UK or another market. Local language, professional practice and access to tools can change whether retained knowledge is usable even when contractual wording is clear.

Decide what remains after the pilot

Ask the commercial sponsor and a practising domain expert to bring one proposed partnership to the next investment review. Alongside the promised result, show the expected learning assets, permitted reuse and the capability that will remain inside the organisation.

Fund a bounded experiment where those arrangements can be tested. Review useful performance and the cost of human correction together. A strong result may justify deeper collaboration; poor transfer or unclear access may justify a narrower scope. Neither outcome should depend on the partner's brand alone.

The next generation of AI partnerships may produce better specialist systems. The executive task is to ensure that a successful collaboration also leaves the organisation able to understand what improved, use what it paid for and make an informed choice about what comes next.

RESEARCH RECORD

Method and limitations

Method

Original Institute synthesis of the Synopsys announcement, new arXiv research and official Singapore consultation material, checked 2 October 2026. Commercial questions and the learning-asset register are Institute recommendations. No supplier contract, customer data or private conversation was inspected or quoted.

Limitations

The announced service is forward-looking; no independent product benchmark, full contract or general availability date was established. Research papers are preprints and do not evaluate GPT-Synopsys. No global adoption estimate or comparative regional outcome is claimed. Legal arrangements require jurisdiction-specific review.



Sources and limitations

Source facts link to the original publication. Institute analysis reconciles definitions and turns the material into a management proposition. Frameworks are decision aids, not claims of universal causality. Read every important statistic with its population, date, definition and caveat.

Research cutoff: 2 October 2026. Review cycle: six months, or earlier after a material change in the underlying facts, standards or policy.

REFERENCES

Evidence behind the thesis

01 Synopsys: OpenAI and Synopsys announce GPT-Synopsys

30 September 2026. Company announcement; early engagements and forward-looking service. Explicitly states customer data is not used to train the model.

02 Qian et al.: AREX-2

v1 submitted 29 September; new listing 1 October 2026. Preprint on long-horizon reflective training and benchmark transfer.

03 Xu: Self-Evolving Harness on Multiple Tasks

New submission in the 1 October 2026 batch. Preprint; frozen model, evolving surrounding software, held-out benchmark evaluation.

04 Bakman et al.: Aligned Data Can Induce Misalignment via Context Confusion

v1 submitted 29 September; new listing 1 October 2026. Preprint with bounded cross-context experiments, not product evaluation.

05 Singapore MinLaw and IPOS: AI and intellectual property consultation

26 August–22 October 2026. Consultation, not enacted law. Retained only for its specific jurisdictional status.

THE AI INSTITUTE

Make the decision.
Then measure what changes.

Research · Readiness · Capability

theai.institute