

PERSPECTIVE

The AI Diffusion Ecosystem

What it takes to turn frontier intelligence into trusted, scalable outcomes

June Manley
September 2026

FOUR CONDITIONS

Economics · Infrastructure
Governance · Frontier intelligence

AT THE CENTER

Redesigned work



- 01 The catalyst
- 02 The AI Diffusion Ecosystem
- 03 Outcome-based economics
- 04 Sustainable and resilient infrastructure
- 05 Proportionate governance and guardrails
- 06 Safe, reliable frontier intelligence
- 07 Where diffusion actually happens

The AI Diffusion Ecosystem is the interdependent system required to turn frontier intelligence into trusted, scalable and economically valuable outcomes.

01
Outcome-based economics
Cost per useful outcome

02
Sustainable and resilient infrastructure
Optionality and control

03
Proportionate governance and guardrails
Controls scaled to consequence

04
Safe, reliable frontier intelligence
Evaluation and assurance

The ecosystem creates the conditions. The operating model converts them into value.

01 The catalyst

When Dario Amodei called on the AI industry to "pace the frontier," Sam Altman agreed within hours. Elon Musk needed three words: "Dario is right."

Then The Information reported that Anthropic, OpenAI and Google had already been discussing an industry standards body, and Anthropic committed to giving outside evaluators employee-level access to its models and training pipeline.

The consensus is narrower than it first appeared. Some industry leaders favor coordinated pacing, independent evaluation and stronger public oversight. Others put more weight on company-level engineering discipline, testing and existing incentives.

There is growing agreement on the need for control, reliability and accountability. The disagreement is over the mechanism.

One version of that disagreement has already reached federal court.

So the questions arrive quickly. Who funds the evaluators? Are they independent? Can labs coordinate without an antitrust waiver? Can any of this hold if China does not participate?

Those are serious questions. They are also questions about one condition in a much larger system.

The debate about speed is really a debate about something bigger: what has to be true for increasingly capable intelligence to diffuse safely, sustainably and economically.

Pacing governs the velocity of frontier capability. It does not govern whether that capability can move into the economy and into work.

02 The AI Diffusion Ecosystem

I think of it as an ecosystem.

The AI Diffusion Ecosystem is the interdependent system required to turn frontier intelligence into trusted, scalable and economically valuable outcomes.

AI diffusion is not a new term. Microsoft uses it to measure the share of the working-age population using generative AI, and the OECD has studied the adoption and diffusion of AI across firms. What I am proposing is an architecture for the conditions that make that diffusion possible.

There are four. Outcome-based economics. Sustainable and resilient infrastructure. Proportionate governance and guardrails. Safe, reliable frontier intelligence.

They are interdependent. Weakness in any one can constrain diffusion regardless of how capable the underlying model becomes.

A model that cannot be trusted does not get deployed. A model that cannot be afforded at scale does not get used. A model with no governance around it does not get permission.

And the four conditions surround one thing.

The ecosystem exists to enable how work gets redesigned.

I'll come back to that. First, the conditions.



FIGURE 1 The four conditions of the AI Diffusion Ecosystem surround redesigned work.

03 Outcome-based economics

Does more intelligence create more value?

That sounds like an obvious yes. It is not, because the participants in the ecosystem are compensated against different units. Labs monetize intelligence. Clouds monetize consumption. Software monetizes usage or seats. Integrators often monetize effort. The enterprise has to produce an outcome.

None of those participants is indifferent to customer value. But each is rewarded for something other than the outcome the customer is accountable for. When every layer of the stack optimizes its own unit, the enterprise absorbs the gap.

The metric I keep coming back to is cost per useful outcome. Not cost per token. A customer issue resolved. A contract reviewed accurately. A pipeline risk surfaced before the quarter closed. Tokens measure an input. Businesses care about what the system accomplished.

Cost per useful outcome started as a way to right-size the intelligence applied to a task. It has a bigger job now. It asks whether the economics of the entire ecosystem point toward value creation or toward consumption.

For diffusion to scale, those incentives have to start converging. Outcome-based pricing, shared risk between vendor and customer, and services economics tied to results rather than hours would move those incentives closer to the customer's outcome.




FIGURE 2 Each layer of the AI stack monetizes a different unit. The enterprise is measured on cost per useful outcome.

04 Sustainable and resilient infrastructure

Intelligence has to be computed every time it is used.

That makes AI infrastructure look less like a software license and more like industrial capacity. Compute, energy, cooling, grid access and hardware all sit underneath every inference, and the organizations that control that infrastructure control the cost curve.

I have written at length about the compute divide, so I will keep this short.¹

Intelligence has to be economically, physically and operationally sustainable before it can diffuse at scale.

That means more than efficiency.

IT MEANS

model routing	deliberate workload placement	data sovereignty	portability across accelerators and providers	the ability to change course when pricing or capability shifts
---------------	-------------------------------	------------------	---	--

Infrastructure and model optionality preserve the enterprise's ability to control its own cost curve.

¹ June Manley, [Software Ate the World. Hardware Must Not Starve the Dream](#), September 2026.

Without them, diffusion is not democratized intelligence. It is democratized access to concentrated infrastructure.

05 Proportionate governance and guardrails

The most capable model is not automatically the most deployable model. Trust decides permission to deploy.

As The Information has reported, Palantir has withheld Anthropic's Fable model from its platform pending an irrevocable zero-data-retention guarantee, Nvidia has limited Fable to less sensitive work, and Booz Allen has restricted its use on proprietary cybersecurity projects. A large utility canceled a pilot using the model on core grid infrastructure over the same issue.

Capability did not change. Trust did.

Governance is what turns trust into permission, and it should be proportionate.

Controls should scale with consequence.

A contract review agent should not be governed like an agent operating the power grid.

Some risks require frontier-level evaluation. Some require a sandbox. Some require identity and permissions. Some require a human approval step. Some require a contract clause. Some require the board.

Much of the discipline organizations need already exists.


Ciaran Martin, who ran Britain's National Cyber Security Centre, argues that the common factor in the agent incidents he examines was a failure to control the testing environment, and that established security practices would have gone a long way toward containing them.

We should not invent new governance machinery where security, risk and accountability practices we already understand will do the job.

Where we do need new governance, someone has to own the decision.

Evaluation without authority is information. Governance determines what happens next.

Controls should scale with consequence



At the frontier, ownership is harder because there is no agreed governance model. Some leaders want government and the public to have a formal stake. Others put the responsibility on companies to test rigorously and withhold systems that are not ready.

The mechanism can vary. What cannot be ambiguous is authority and accountability: who can challenge a consequential decision, who can stop it and who owns the outcome.

06 Safe, reliable frontier intelligence

Capability without reliability does not diffuse.

This is where the pacing debate finally lands, and it lands in a narrower place than the headlines suggest.

The Economist's reading of Amodei's essay is that a slower pace is less about halting progress and more about creating room for operational

excellence, alignment and interpretability to catch up. That framing is the right one.

Frontier intelligence has to become more reliable as it becomes more capable, or it stays in the lab.

Evaluation and control have emerged as the clearest points of convergence. The harder question is who does the evaluating and what authority follows from the result.

Who evaluates, and who pays them?

The evaluator ecosystem is small and young, and critics have already questioned the funding and personnel ties between watchdog groups and the labs they assess.

Independence has to exist structurally, not as a label.

What access do they get?

An evaluator who sees a finished model but not the training pipeline cannot judge whether the lab responded to its own safety findings.

What happens when they find something serious?

If the answer is nothing, this is oversight in name only.

At the frontier, we may need forms of independent assurance that are not yet mature or institutionalized.

In the enterprise, much of the necessary governance discipline already does.

Pacing is one mechanism for creating the space to build that assurance. It is not the end state.

07 Where diffusion actually happens

None of the four conditions creates diffusion on its own.

Economics can converge. Infrastructure can preserve optionality. Governance can be proportionate. Frontier models can be reliable and independently verified.

And a company can still be running pilots on the side of the business, measuring productivity on tasks, waiting for transformation to arrive.

Diffusion happens when the operating model changes.



Human expertise moves from production toward judgment.

The scarce resource is no longer the ability to produce more work. It is context, critical thinking and the ability to recognize when an answer that looks perfectly reasonable is wrong.

The ecosystem creates the conditions for diffusion. The operating model converts those conditions into value.

That is where diffusion begins.

Ultimately, this is about how organizations absorb what the AI industry builds.

The labs will decide how fast to build. The ecosystem will determine whether that intelligence can be trusted, afforded and sustained. Every organization will still have to decide how to work.

01

Outcome-based economics

02

Sustainable and resilient infrastructure

03

Proportionate governance and guardrails

04

Safe, reliable frontier intelligence

Published September 2026. Enterprise examples reflect reporting as of September 19, 2026.

Sources: The Information; The Economist; Microsoft AI Economy Institute; OECD; Associated Press, "Lawsuit says Anthropic, OpenAI, SpaceXAI and Google made illegal agreement on AI slowdown," September 20, 2026; Ciaran Martin, former CEO, UK National Cyber Security Centre; June Manley, ["Software Ate the World. Hardware Must Not Starve the Dream," September 2026.](#)

www.junemanley.com