THE LINUX FOUNDATION | Research

Enterprise Transitions to the Agentic Era

A Qualitative Study of How Open Source AI Innovators and Enterprise Practitioners are Navigating the Shift from Cloud-first to Agent-first

Hillary Curran, *Guru Technologies*
Hilary Carter, *The Linux Foundation*
Anni Lai, *Futurewei; Generative AI Commons, LF AI & Data*

Foreword by
Guillaume Blaquière, *Group Data and Agentic Architect, Carrefour*

September 2026

Sponsored by:

FUTUREWEI Technologies
GURU
LF AI & DATA
tokenomics Foundation
Agentic AI Foundation

Enterprise Transitions to the Agentic Era

Model capability is no longer the gate; **trust, governance, and cost** now set the pace of enterprise agent adoption through 2030.

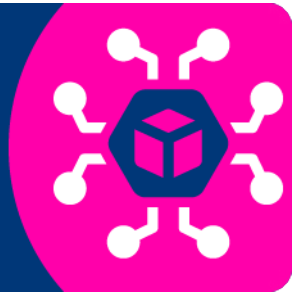

Agents are becoming an **orchestration layer** for the enterprise, an evolution in cloud architecture that enables greater coordination across existing systems.

As agents become more autonomous, **accountability for their errors remains with the organizations that deploy them** and the people who authorize their use.

Per-seat pricing is under pressure as agents reduce the number of people needed to do the work; pricing must shift toward usage or outcomes.

Put the control layer in early by placing a governance and observability gateway between agents and every model and tool they call.

Manage volatile token economics by metering usage at the gateway, so you can see which team and **which workflow is spending what.**

Measure what the agent completes rather than what it consumes, tracking completion and resolution rates **against a fully loaded human-process baseline.**

Position organizational **knowledge and context** as core enterprise infrastructure, because an agent's effectiveness is strictly bounded by what it knows about the business.

Convert shadow AI into governed adoption with a shared platform, since **forgotten tools stay callable** and credentialed long after anyone remembers them.

Build trust into the plumbing from the start, because it constrains agentic adoption more tightly than any remaining limit on model capability.

Enforce policy at the gateway and tool boundary rather than asking the model to police itself, because retrospective audit cannot keep machine pace.

Engage in **open source** communities to secure long-term digital **sovereignty, interoperability,** and full system **auditability.**

Contents

Foreword

When we look back at the early waves of enterprise artificial intelligence, the overarching concern was almost entirely technical feasibility: Can the model actually perform the task? Today, that foundational question has decisively shifted. Model capability is no longer our primary gate. The real challenge facing global enterprises as we transition into the agentic era lies in navigating operational friction, legacy constraints, and integration pain points that dictate the true pace of enterprise adoption.

One of the major challenges centers on data access in the real world versus cloud native assumptions. Protocols like the Model Context Protocol (MCP) assume that enterprise data is already cleanly reachable through modern APIs. Unfortunately, it's not so straightforward. Many large organizations still operate complex legacy environments whose modernization is necessarily a multi-year journey. Building real time data change replication systems and pipelines into cloud mirror databases is operationally feasible, but at the same time, introduces latency and carries significant cost. Building agentic systems around legacy architectures is inherently expensive precisely because those foundational systems were never designed for automated agentic interactions.

Another challenge is recognizing that sophisticated model architectures and open frameworks do not automatically drive user adoption. Enterprise leaders effectively face a fork in the road: expend resources on traditional change management or invest heavily in deep agent integration within existing applications, daily workflows, and core operational processes. In our experience at Carrefour, successful adoption depends as much on people, processes and ways of working as it does on technology. . Consequently, the ultimate success of any agentic deployment depends on allocating budget to this integration layer—an aspect that remains widely underestimated across the industry.

This study offers a timely, grounded look at these and other operational realities facing organizations today. By synthesizing insights from open source innovators and global enterprise practitioners worldwide, this report provides executive leaders, software architects, and digital transformation teams with a view into the current landscape, and a practical roadmap for building responsible, high impact intelligent enterprises. Accounting for the trade offs between speed, governance, and infrastructure is paramount. And achieving durable scale requires confronting our legacy realities, prioritizing deep system integration, and leveraging open source ecosystems where they can strengthen auditability, interoperability and technological sovereignty. . I invite leaders across industries to engage with these insights as we collectively shape the future of the agentic enterprise.

Guillaume Blaquière, *Group Data and Agentic Architect, Carrefour*

Executive summary

Enterprises have deployed artificial intelligence for decades. Speech recognition, document classification, demand forecasting, and fraud detection were in production long before the current wave, and a generation of companies was built on them. What changed a few years ago was narrower and sharper: the transformer architecture, the large language models it made possible, and the question of whether this new class of system could do real enterprise work at all. Drawing on in-depth interviews with open source AI innovators and enterprise practitioners worldwide, this study finds that the question has changed again. Interviewees broadly agreed that model capability is no longer the primary constraint. The harder and less-resolved problems are organizational: establishing enough trust, governance, and accountability to let an agent act, and doing so at a cost the enterprise can sustain.

That shift, from "can it work" to "can we trust it to act," and increasingly "can we manage the associated costs", organizes the four findings of this study. First, on the technology stack, interviewees accepted that agents are becoming a coordinating layer of the enterprise but generally declined to endorse the stronger claim that the agent has become the enterprise "kernel"; most described an evolution of existing cloud architecture rather than its replacement. Second, on governance, they were nearly unanimous that accountability for an agent's actions stays with the deploying organization and its people, and that guardrails must be enforced beneath

FIGURE 1

THE QUESTION HAS CHANGED

The constraint on enterprise AI has moved from capability, to trust, to cost.

the model rather than requested of it. Third, on business models, they agreed that per-seat pricing is under terminal pressure while disputing what replaces it, and they rejected the popular claim that established SaaS vendors will simply disappear. Fourth, on return on investment, they reframed high pilot-failure rates as a normal feature of experimentation and converged on measuring outcomes delivered rather than activity or tokens consumed.

Trust recurs as the binding constraint in every pillar. As Mazin Gilbert, Executive Director of the Agentic AI Foundation, put it, "Trust isn't a feature an enterprise can add later. It has to be in the plumbing." But when interviewees were asked directly to name the single biggest bottleneck to full adoption by 2030, their answers split by seat: people working on standards, foundations, and open ecosystems most often named trust infrastructure, while operators running AI inside a single company were more likely to name cost, data quality, organizational change, or the human capacity to review agent output. One interviewee declined to name a single bottleneck at all, arguing that the binding constraint moves as systems mature. This paper therefore treats trust as a connective theme rather than declaring it the sole determining factor.

A cross-cutting theme runs beneath all four pillars: enterprises are in a "shadow AI" phase in which bottom-up experimentation is running ahead of formal integration. Interviewees described ungoverned tooling proliferating outside official oversight, durable adoption emerging when leadership enables rather than mandates, and open source functioning less as a cost lever than as the substrate for interoperability, auditability, and sovereignty. The through-line for enterprise leaders is that the infrastructure the agentic era most needs, including protocols, identity, policy, and observability, is being built, and is most credibly built, in the open.

Introduction

The agentic enterprise is commonly framed as a transition in which agents become the organizing layer of the technology stack, work is increasingly executed autonomously, and human roles are augmented by digital coworkers. This report treats that "agent-as-kernel," execution-first premise as the industry narrative it tests, not as its own claim.

A note on terms is warranted, because "AI" carries decades of history that this report is not about. Machine learning has been in enterprise production since long before the current moment. The discontinuity examined here is specific: the transformer architecture, the large language models built on it, and the agents now being built on those. Where this report refers to AI without qualification, it means that lineage rather than the broader field.

The findings in this report reflect analysis from 25 interviews conducted from April to June, 2026, with volunteers sourced from LF AI & Data Foundation, the Agentic AI Foundation, and Linux Foundation Research participant communities. Across the interviews, a more qualified reality emerged: interviewees broadly held that capability is no longer the gate, and that the pace of adoption through 2030 will be set less by advances in the models than by organizational factors: trust, governance, cost, data, and the management of enterprise knowledge.

The study is built on a four-pillar interview framework: technology strategy and the "reasoning kernel"; governance, risk, and the digital coworker; business model shifts; and return on investment and the agentic roadmap. It also includes a fifth, cross-cutting section that consolidates recurring themes around shadow AI, adoption culture, and the role of open source. A short sidebar considers embodied, physical-world AI. Findings are drawn from approved, attributed interview quotes; some participants are identified by role only, and the analysis reflects a qualitative synthesis rather than a statistical survey.

The transition to agentic AI will not follow a single path. Each enterprise's journey will depend on its cloud and data maturity, legacy technology landscape, regulatory requirements, and available budget. Cloud native organizations may be able to deploy agents relatively quickly, while legacy-intensive enterprises may first need to invest in data replication, integration layers, and workflow redesign. For many established companies, agentic transformation will therefore be a gradual process that must coexist with legacy systems for years rather than a direct move to an idealized cloud-native environment. This study's participants sit disproportionately at the cloud-native end of that range.

FIGURE 2

HOW THE STUDY IS ORGANIZED

Four pillars, cross-cutting themes above, and trust as the constraint beneath them all.

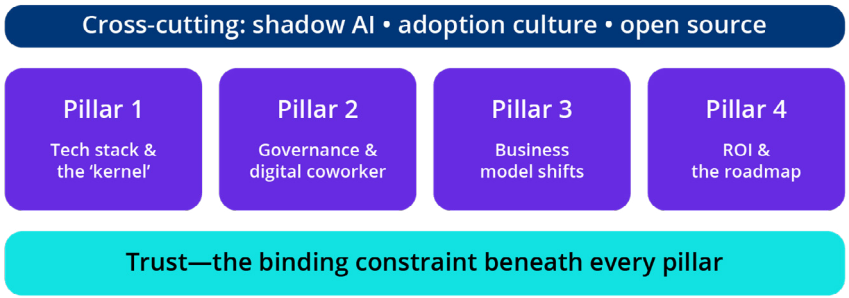

A further aim shapes how the report should be read. Much of the preexisting literature on agentic AI reflects a single vantage point (a vendor, an analyst, or a region), and carries the biases of that seat. This study is deliberately constructed as an aggregation across roles and geographies: it brings together contributors to open source projects and standards bodies, independent technology advisors, and practitioners running agentic AI inside enterprises, spanning North America, Europe, and Asia-Pacific.

Many participants are hands-on operators, developers, and open source contributors who rarely publish, and whose working experience is largely absent from the public conversation. The intended contribution is therefore not another forecast from one perspective, but a balanced synthesis of how a deliberately varied group of practitioners describe the same transition, often affirming, and sometimes complicating, what the most visible industry voices are saying.

FIGURE 3

WHO WAS INTERVIEWED

25 attributed participants contributing through semi-structured interviews and written responses, April–July 2026. Each mark is one participant.

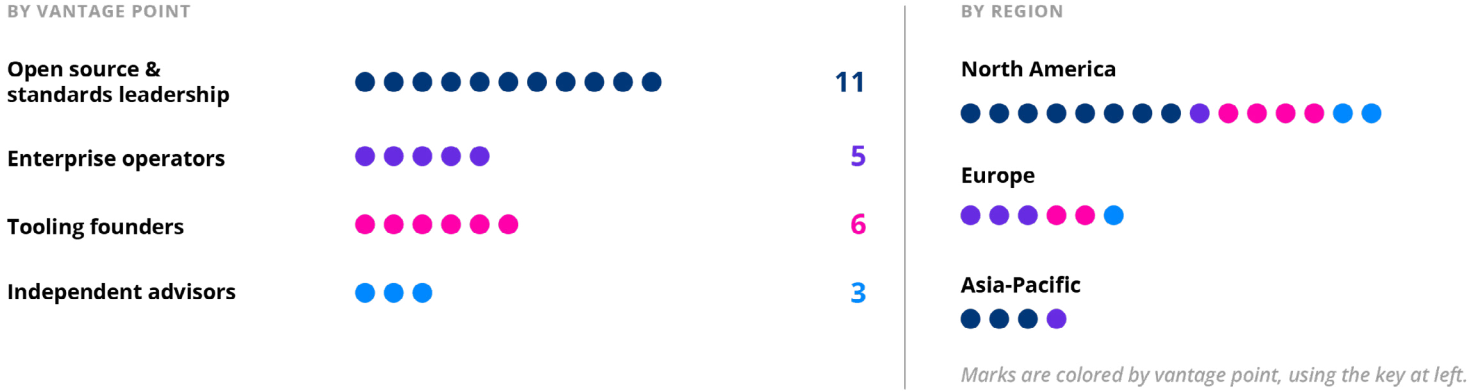

The framing hypothesis of the study is that enterprises are in a shadow-AI phase in which experimentation precedes formal integration, with open source acting as accelerator. Open source agentic AI projects, among them the Model Context Protocol, Goose, AGENTS.md, and agent-to-agent protocols, recur throughout as the layer most relevant to the open source audience, and are examined pillar by pillar below.

The industry backdrop: what the major players are saying

To situate these findings, it is worth noting how the largest platform vendors and frontier labs were publicly framing the agentic transition over the same period. The remarks below are drawn from public keynotes and published interviews rather than from this study's primary interviews, and are offered as context that the practitioner findings in the chapters ahead both affirm and complicate.

The dominant message from the major cloud vendors in 2026 was that the experimental phase had ended. Opening Google Cloud Next 2026, CEO Thomas Kurian framed the shift as already underway.

> *"The era of the pilot is over. The era of the agent is here."*
>
> Thomas Kurian, CEO, Google Cloud — Google Cloud Next 2026

> *"You have human agents and AI agents, and they basically govern the same way."*
>
> Jensen Huang, CEO, NVIDIA — ServiceNow Knowledge 2026

That vision extended to how agents would be managed. Speaking at ServiceNow Knowledge 2026, NVIDIA's Jensen Huang described a future workforce in which digital and human workers are overseen through the same structures.

From inside a frontier lab, Fiona Fung, who manages Anthropic's Claude Code and Cowork teams, offered a more operational note that anticipates the study's central theme, that capable models still require human judgment.

The chapters that follow test these framings against the experience of practitioners building and deploying agents. In several places the interviews affirm the industry narrative — on human accountability and verification in particular — and in others they complicate it, most notably on whether the "pilot era" is truly over and on how quickly the transition will run.

> *"It's all about trust, but verify. The models are really good, but there are definitely a lot of areas that still need verification. And wherever you need deep subject matter expertise, that's an area to definitely still invest in."*
>
> Fiona Fung, Manager of the Claude Code and Cowork teams, Anthropic — Lenny's Newsletter, June 2026

Pillar 1: Technology strategy: The agent as coordinating layer

The kernel metaphor, complicated

Interviewees broadly agreed that agents are beginning to coordinate enterprise resources in new ways. Most, however, declined to endorse the strong claim that the agent has become the enterprise "kernel," describing an evolution of existing cloud architecture rather than its replacement. Xavier Shay, whose remit spans Up Bank and Bendigo Bank, located the change in architecture rather than in a new controlling entity:

> **"The true enterprise shift is architectural, not technological. We are heavily focused right now on what the orchestration layer for agents actually looks like. We need to determine the equivalent orchestration framework for running agents across an organization in a similar way to what we did with cloud native."**
> **– Xavier Shay, CEO, Up Bank; Chief Digital Officer, Bendigo Bank**

The strong version of the claim was argued most directly by interviewees working on standards and platforms. Matt White, former Global CTO of AI at the Linux Foundation, held that "the agent is becoming the new enterprise kernel," reasoning that "just as the operating system abstracted hardware, the agentic layer is abstracting applications, data, and workflows; and like every kernel that came before it, the one that wins will be open." Vincent Caldeira, VP of AI (APAC) at Red Hat, offered a similar reading: The corporate core, in his account, "is shifting from a centralized database model to an active execution layer," within which agents "function as an operational operating system kernel, dynamically orchestrating applications, accessing federated data products, and coordinating digital workers to execute complex goals."

Others qualified the metaphor instead of rejecting it. Several likened the agent to a scheduler rather than an operating system: it owns no resources but prioritizes access to them based on inferred intent. On this view, data becomes more important as agents proliferate, and harder to manage. Hao Li, Founder and CEO of HAOEASY AI, framed the change as continuity, holding that agents "are not replacing enterprise systems, they are becoming the execution layer that orchestrates them." Richard Bian made the continuity explicit, holding that agents "do not replace the fundamentals of software engineering; they add a new execution layer," with "cloud infrastructure engineers … naturally evolving into agentic infrastructure engineers."

MCP: Today's standard, contested durability

Interviewees near-universally identified the Model Context Protocol as the default means by which agents reach tools and data. Mazin Gilbert of the Agentic AI Foundation offered the most-quoted characterization:

> **"MCP is the USB port of the agentic world — one unified protocol for any agent to reach databases, knowledge, and tools. But USB alone didn't make the internet."**
> **– Mazin Gilbert, Executive Director, Agentic AI Foundation**

Angie Jones, who helped build Block's Goose agent prior to its contribution to the Agentic AI Foundation, and before MCP existed, described why a standard became necessary:

> **"We started building an in-house agent at Block, called Goose (now open source), before MCP existed. The problem was that for everything we wanted to connect it to, we had to write the integration code ourselves. We realized very early this wasn't going to scale. You'd ship it with five connections, and anytime you wanted to expand to more connections it was a whole project. So we started thinking this should be a standard, and we talked to Anthropic about it. They said they were working on something, and they showed us the spec for MCP. It was clear a standard was the only way to go, and we became design partners."**
> **– Angie Jones, VP of Developer Relations, Agentic AI Foundation**

Fewer interviewees expected MCP's primacy to persist unchanged. Several placed it in the lineage of earlier open-standards efforts such as open banking, valuing it for interconnectivity across ecosystems rather than siloed ones.

Ania Musial, Head of AI Platforms (Product) in Bloomberg's Office of the CTO, tied that interoperability back to the cloud-native precedent: "The transition toward agentic AI strongly mirrors the cloud native revolution. Just as Bloomberg was among the first to embrace Kubernetes, we see developer workflows being dramatically reimagined today using agentic AI." She cast open standards as the safeguard against fragmentation, arguing that "embracing open standards like MCP across the technology ecosystem is critical to avoid creating agentic silos, to ensure responsible evolution of technology, and to drive cross-functional collaboration across the enterprise."

Guillaume Blaquiere, Group Data and Agentic Architect at Carrefour, went furthest in questioning the protocol's durability: "MCP is today a standard like the API before. However, with the quality of models, APIs and skills would soon be sufficient and I'm not sure MCP will have values, and might disappear." Several interviewees pointed instead to "skills" as the more accessible abstraction, a point Lin Sun, VP of Open Source at solo.io, made directly: "Writing a skill is so much easier than writing an MCP tool. MCP involves JSON-RPC and writing code, while skills are mainly just a markdown file. Skills have a much lower learning curve."

MCP can standardize how agents interact with tools and data sources, but it does not, by itself, make enterprise data accessible. Many established organizations still rely on legacy systems that lack modern APIs and cannot be replaced quickly. These enterprises may need intermediate solutions such as change data capture, replicated databases, or custom integration layers to make legacy data available to agents. Such approaches introduce additional cost and complexity, and they typically leave the agent working from data that is minutes old rather than live.

Beyond MCP: The larger stack

The more architecturally focused interviewees regarded MCP as necessary but insufficient, pointing to agent-to-agent communication, identity, and discovery as the layers that will ultimately decide whether agents from different vendors can interoperate. In their account MCP is one layer of a larger stack: A2A standardizes how agents talk to other agents, while discovery, identity, and observability layers make those connections trustworthy at scale. Some called the result an "Internet of Agents."

FIGURE 4

MCP IS ONE LAYER, NOT THE WHOLE STACK

The agentic enterprise needs five complementary layers; MCP handles only agent-to-tool.

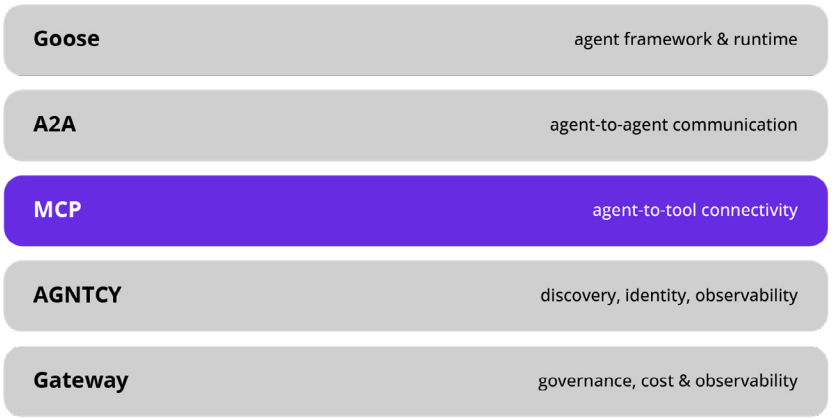

The intermediary gateway layer

Among the strongest points of agreement in the dataset, practitioners across diverse organizations reported placing an intermediary layer between agents and the models they call, and characterized it not as an optimization but as a baseline requirement for production. Lin Sun described the pattern concretely:

> **"We put an agentgateway between the agents and the large language model. Through the gateway, as both a governance point and a metrics provider, we were able to see which engineer was using how many tokens."**
> **— Lin Sun, VP of Open Source, solo.io**

Jérôme Oufella, Vice-President of Technologies at Savoir-faire Linux, reported building "an MCP gateway for VulnScout … anticipating that people will use it in agentic workflows."

FIGURE 5

THE GATEWAY PATTERN

The strongest point of agreement in the sample. Practitioners placed a governance and observability layer between their agents and everything those agents call, and described it as a baseline for production, not a later optimization.

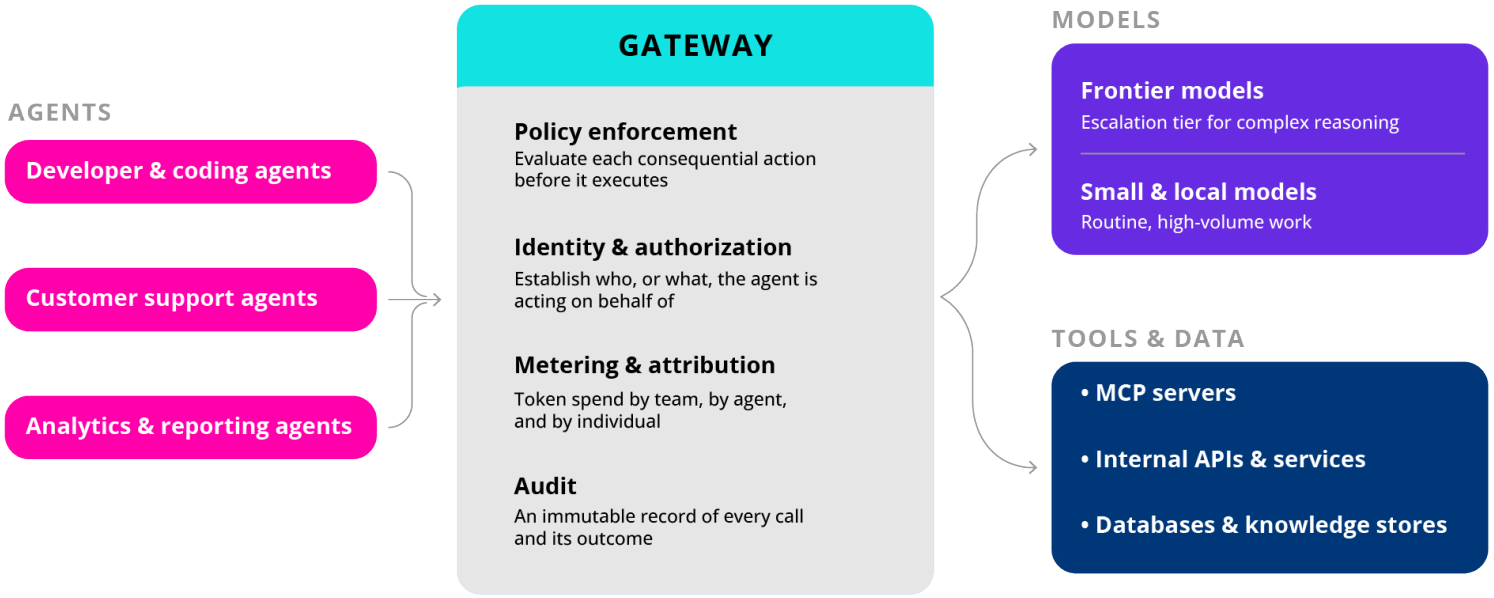

Shared roadmap priorities, and the “Claw” divergence

Maintaining private forks is clearly a common strategy for organizations to bridge gaps between their needs and OSS project roadmaps. However, sustaining private forks over time creates significant maintenance burdens that evolve into technical debt. Our survey analysis generates several insights into the nature and scale of these hidden costs.

As shown in Figure 2, downstream users rely equally on two strategies: contributing upstream and maintaining private forks. In Figure 5, we summarize the key reasons cited by organizations for maintaining private forks: the need for custom features (54%), integration with internal operations (44%), and adding security patches or compliance requirements (37%). While many integration costs can simply be seen as a necessary cost incurred to use open source software, there are likely cases where the upstream projects would benefit from contributions in the form of new features or security patches.

RECOMMENDATIONS

FOR
Enterprise Leaders

INVEST IN THE INTEROPERABILITY LAYERS, NOT ONLY IN MODELS.

Direct roadmap spend toward agent-to-agent communication, identity, and discovery—the layers interviewees expect to decide interoperability—rather than betting on any single model or app. Data accessibility and legacy integration should be treated as foundational components of agentic readiness.

FOR
Platform & infrastructure teams

TREAT AN INTERMEDIARY GATEWAY AS A BASELINE.

Place a governance-and-observability gateway between agents and models from the outset. Interviewees treated it as a baseline for production, not a later.

FOR
Builders & developers

MATCH THE ABSTRACTION TO THE TASK.

Favor skills where they suffice, for maintainability and a lower learning curve; reserve MCP tools for cases that need them, and expect the two to converge over time.

WHAT OPERATORS SAID ABOUT ACCOUNTABILITY

"You're all managers now... A manager is accountable for what their team produces, whether that output came from a person or an agent. If it isn't right, you don't narrate its origin story. You correct it before it reaches anyone else.

Rick Nucci - Co-founder & CEO. Guru

"At the end of the day, AI is just a tool... the ultimate accountability for ensuring a product works, provides customer value, and complies with regulators still comes down to the individual teams."

Rick Nucci - Co-founder & CEO. Guru

"You are the leader and you're now a leader of agents. You are responsible at the end of the day for whatever happens with those agents."

Grace Francisco • Tech Advisor, m2i

Pillar 2: Governance, risk, and the "digital coworker"

Accountability stays with humans

This was the most widely shared finding in the dataset. Regardless of an agent's degree of autonomy, interviewees located responsibility with the deploying organization and the people who configure and authorize the agent. The idea that a failure is the fault of AI alone would not hold up in a regulatory or reputational context.

Angie Jones described how the principle held up at engineering scale, and why employees welcomed it. Her team's rule of thumb was that "a human had to be held accountable"; even when an agent opened a pull request, "your byline is still attached, so you're accountable for it," and you "must review it before it's merged in." She found that "employees actually really liked that. They wanted their name associated with the actions."

Diego Gosmar offered the framing several interviewees preferred over "autonomous agent," a "delegated digital coworker with bounded authority." In smaller organizations the same principle appears as direct practice; Arnaud Giuliani, co-founder of Kotzilla, reported that "each execution is always human validated and never pushed directly to production."

The open frontier: accountability across interacting agents

While responsibility for a single agent was treated as settled, interviewees identified accountability across interacting agents as the unresolved frontier:

> **"The hardest question is not who owns a single agent: it is who bears responsibility when a failure emerges from the interaction between agents, where no single decision caused the outcome, but the combination of legitimate decisions produced an illegitimate result."**
> **— Christian Maitre, President, Caspera.lab**

The governance question shifts accordingly. It is no longer only "where is the human in the loop" but where agents hand off to other agents, and how those handoffs are governed. Lin Sun located the same difficulty in access control: The common case is an agent that acts on one person's behalf, but "the more complex case is when an agent works on behalf of multiple people ... that's where access control and audit become really interesting and important."

Runtime guardrails, enforced beneath the model

Interviewees agreed that retrospective audits are insufficient when agents act at machine speed, and that policy should be enforced before an agent acts, at the runtime level rather than through model self-regulation.

> **"Operating in a highly regulated space, we rely heavily on domain expertise and keeping humans in the loop to shape system quality. To add deterministic controls to non-deterministic systems, we contributed Interceptors to the MCP standard, allowing us to connect guardrails and validate messages at runtime."**
> **— Ania Musial, Head of AI Platforms (Product), Office of the CTO, Bloomberg**

FIGURE 6

GUARDRAILS BELONG BENEATH THE MODEL, NOT INSIDE IT

Policy is enforced before, during, and after an agent acts, not by the model itself.

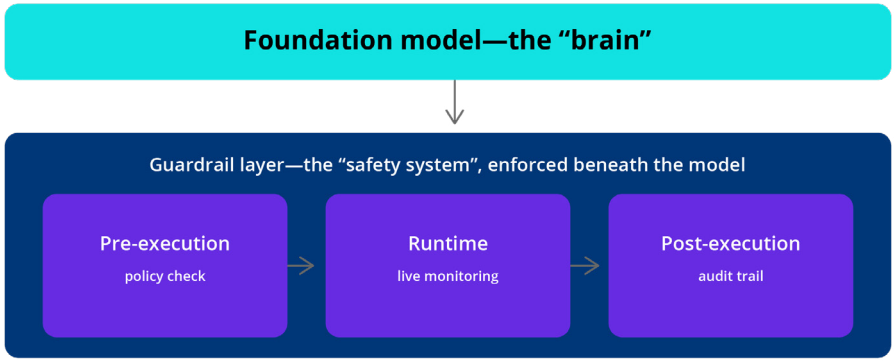

Vincent Caldeira described the stakes of getting this wrong: "When a classic chatbot hallucinates, it returns a wrong answer; when an active agent compromises its logic, it executes unauthorized production actions," requiring a "multi-layered, defense-in-depth stack, assuming breach by default." Emily Chen summarized the layered model: "If foundation models are the brain of an agent, guardrails are its safety system."

François Méro located the same principle inside the agent's execution loop, and cast it as a precondition for adoption in data-heavy work:

> **"The point of a runtime guardrail is to move checks into the agent's loop, so it fires while the work is being done rather than after the damage is out. This is crucial for enterprises that want to adopt AI in their data science practice without compromising on trust."**
> **— François Méro, CEO, Probabl**

Grounded risks, and the limits of consensus

The governance discussion was grounded in observable failure modes rather than hypothetical risk. Andreas Traczyk, an open source software consultant at Savoir-faire Linux, warned that "a lot of companies may unknowingly be breaching their own security protocols and exfiltrating their code bases by using an agent to perform a daily task," and Grace Francisco cautioned that "your data is the gold mine of everything," with the real risk being agents "making decisions you would never have made." Jérôme Oufella framed the operator's balancing act as the need to "bridge the need to help our teams benefit from using agents while managing cybersecurity and regulatory risk." The consensus on guardrails proved easier to state than to operationalize: managing policy across thousands of agents would itself require automation.

Recommendations

FOR Enterprise leaders, legal & risk

ASSIGN ACCOUNTABILITY BEFORE DEPLOYMENT.

Treat agents as delegated coworkers with bounded authority and a named human owner before deployment, so lines of accountability are clear to leadership, legal, and regulators.

FOR Platform & security teams

ENFORCE GUARDRAILS BENEATH THE MODEL, IN THE OPEN.

Evaluate policy at the runtime and operating-system level, not by asking the model to police itself, and favor an inspectable, open guardrail layer to earn institutional trust.

FOR Builders & developers

MATCH THE ABSTRACTION TO THE TASK.

Prepare for multi-agent accountability now. Because responsibility for outcomes arising from interacting agents is unresolved, begin instrumenting agent-to-agent handoffs and audit

Pillar 3: Business model shifts

Per-seat pricing under pressure, with no settled successor

There was broad agreement that pricing tied to human seats becomes untenable as agents perform the work. Interviewees judged the per-seat model structurally out of step with an agent-first world, faulting it for pricing the human seat rather than the work the agent does. Yet the successor is unsettled, and Jonathan Rochelle of Lutely named the efficiency paradox directly:

> **"If you're charging on the outcome alone, the person who nailed it on the first try pays the same as the person who ran it 12 times. I'd love a model where there's an advantage to being more efficient, but I just don't have a clean answer."**
> **— Jonathan Rochelle, Co-founder & CPO, Lutely**

One operator who has already moved his own company onto consumption pricing offered a working rule rather than a formula:

> **"Seat-based pricing is on a death march. Consumption pricing is what's replacing it. But swing too far toward it without a plan, and you've traded one broken model for another. … Every billable event has to be explainable in one plain-English sentence. If you cannot say what the customer got for it, do not charge for it."**
> **— Rick Nucci, Co-founder & CEO, Guru**

Zach Lloyd, Founder and CEO of Warp, pointed to the margin problem beneath the pricing question: "When you resell inference, you can't take a software margin on that. It used to be like 90% margin on software. On inference it's not like that at all." Most interviewees sketched the same near-term outcome: a mix of platform fee, consumption, and occasional outcome bonuses, likely for several years. The more telling result is the gap itself. There was broad agreement that the prevailing model is declining, and little agreement on what replaces it.

"SaaS is dead," rejected

In contrast to the prevailing narrative that incumbents will be displaced, interviewees generally expected established vendors to adapt rather than disappear. Grace Francisco was blunt: "Everyone's running around going, 'SaaS is dead.' … No, it's not. SaaS will decline a bit because boutique solutions are getting rebuilt with AI. But for the bigger solutions like Salesforce? People are too embedded." Richard Bian reframed the acronym rather than burying it:

> **"SaaS may increasingly mean service as software: software that delivers outcomes more directly through agents and automation, while the business model remains anchored in value delivered."**
> **— Richard Bian, Head of Product and Operations, InclusionAI**

Xavier Shay pointed to a countervailing force reshaping the build-versus-buy calculus, as non-engineers begin to "'vibe code' entire custom operating systems for their businesses in a matter of months, successfully cutting out external vendors," a shift he noted "raises the operational floor for what non-technical teams can achieve" even as it introduces new risks.

FIGURE 7

PER-SEAT PRICING IS ENDING. ITS SUCCESSOR IS NOT SETTLED.

Interviewees agreed on the direction of travel and diverged on the destination. Agreement declines at each stage.

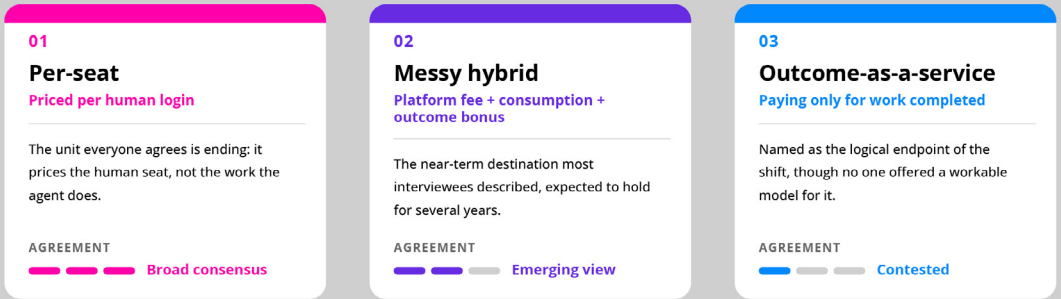

The interface persists, but changes function

The graphical interface was expected to shift from the primary site of work to a surface for oversight, configuration, and exception handling, with durable advantage moving toward APIs, data, and embedded domain knowledge. In this reading the API contract itself becomes the brand, where machine-readable reliability, latency, and semantics are the new UX. Christian Maitre held that "data custody is the new lock-in … The front door still matters, it simply no longer defines the address."

Richard Bian offered a counterpoint to the strongest disintermediation claims: because agents "augment people rather than fully replace them, human-facing interfaces will continue to matter," and the winning products will be those that "discover the right new interaction model for human-agent collaboration."

Toward hybrid model portfolios and smaller local models

The interviews pointed toward a hybrid model portfolio in which small, often local models handle high-volume routine work and frontier models are reserved for complex reasoning, driven by cost and reinforced by data-residency needs. Brad Topol, VP of Engineering at IBM, reported that "Agents are benefitting from improvements in small open-weight models. Early adopters are realizing that running locally results in predictable costs and huge savings." Micaela Stump described this split running in production at Adyen:

> **"For 'low-level, long-execution tasks,' we provide open source models hosted locally on our own on-prem data centers. We also give access to high-reasoning models, like Claude, for complex, high-reasoning tasks."**
> **— Micaela Stump, Team lead of Product Management, Platform Engineering and AI, Adyen**

One interviewee estimated that 80 to 90 percent of calls by volume could run on small or local models, with frontier models reserved for escalation. The trajectory is clearer than the operational readiness to capture it. Running small models in-house is still hard while GPUs are scarce, and Arnaud Giuliani observed that local models are "not yet ready just to run on [any] laptop."

RECOMMENDATIONS

FOR
Commercial & pricing leaders

PLAN FOR A HYBRID PRICING TRANSITION.

Expect a mix of platform fees, consumption, and outcome components for several years; design contracts and instrumentation that attribute value without penalizing efficient use.

FOR
Software vendors & product teams

COMPETE ON DATA AND DOMAIN MODEL.

As the dashboard becomes an oversight surface, invest in API reliability, data custody, and the embedded domain model - but keep designing the human-facing interaction, which interviewees expect to persist.

FOR
Platform & infrastructure teams

BUILD A PORTFOLIO OF MODELS.

Route high-volume routine work to small or local models and reserve frontier models for complex tasks, while budgeting for the real operational cost of running models in-house.

Pillar 4: ROI and the agentic roadmap

Pilot failure, reframed as normal

Interviewees who addressed the widely cited failure statistics generally regarded such failure as an expected feature of experimentation rather than evidence of systemic problems. Jonathan Rochelle held that "50 percent of pilots failing is probably normal. And if it's not, then you're not experimenting enough." What experimentation "should impose on us is measurement": each pilot should state what it is trying to achieve and record it.

Diego Gosmar argued that "most failures aren't model failures, they're measurement and integration failures; pilots that automated a demo rather than a process." Xavier Shay located the stall point precisely: "The barrier to building an impressive prototype is extremely low, but scaling to a repeatable, production-grade workflow you can trust is where organizations stall." Christine Rimer, an independent technology advisor, placed goal clarity upstream of deployment: "you need to be thoughtful about what you're solving for."

A North Star of outcomes, not activity

While interviewees proposed different specific metrics, they shared a common principle: success should be measured by completed work and value delivered rather than by tokens consumed or agents deployed. Christine Rimer noted that token measurement "will collapse fairly quickly as it's measuring a game not outcomes." Proposals for a replacement varied but not in kind: a fully loaded cost per completed outcome against the human-process baseline, an autonomous resolution rate within human-in-the-loop bounds, and a notion of "delegation confidence." Diego Gosmar added a further point about what is hardest to measure and most valuable:

> **"The real ROI question in agentic supply chain is not "how fast can the agent act?" but "how well does the agent know when not to act?" Calibrated restraint is the hardest capability to build, and the most valuable."**
>
> **— Diego Gosmar, Head of AI, TESISQUARE**

Xavier Shay tied the metric back to the customer: "When organizations solve a customer's problem in three minutes instead of leaving them on hold for forty, consumer sentiment shifts permanently. If you have to explicitly tell a user to look at your 'AI feature,' you've built the wrong thing." As a concrete proof point, Lin Sun reported that an internal AI support agent "helped us determine problems three-to-five times faster."

Why faster than cloud

Interviewees offered a consistent explanation with several converging strands. Agents run on cloud infrastructure that enterprises have already built and already trust. Development itself moves faster, with open source compressing into weeks what proprietary roadmaps once took years. Adoption is also bottom-up and competitively pressured: unlike the cloud era, a business that does not transform quickly may not remain a business. And the barrier to entry is far lower than it was for cloud.

Ania Musial put the pace down to accessibility: "The agentic transformation is unfolding at a breathtaking pace compared to the cloud revolution because of how accessible this technology's natural language interface is to just about anyone, not only software engineers and cloud enthusiasts." Micaela Stump made the market-size version of the same point: "The total addressable market for AI is significantly larger than what we saw with cloud native. The barrier to entry is way lower, the use cases are larger, and the broad applicability of AI is driving rapid adoption across the board." Zach Lloyd argued the change runs deeper than a story about delivery speed: "This is intelligence for sale basically. Things that only people could do, now computers can do."

The dissent, and the split on the biggest bottleneck

A dissenting view held that agentic adoption will take roughly as long as the cloud transition did, constrained less by technology than by organizational change:

> **"Agent adoption will also take a decade … the slowest move is not technology, it's human change management."**
>
> **— Guillaume Blaquiere, Group Data and Agentic Architect, Carrefour**

When interviewees were asked to name barriers to adoption, trust and governance recurred most often across the full set of conversations, but it was far from the only concern raised.

FIGURE 8

WHAT ARE THE BARRIERS TO AI ADOPTION?

Barriers raised across the interviews, many named more than one.

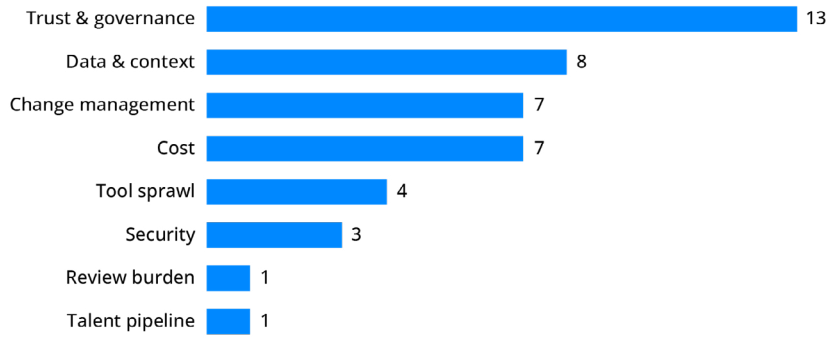

Asked specifically to name the single biggest bottleneck to full adoption by 2030, interviewees divided along role lines. Figures working on standards, foundations, and open ecosystems most often named trust infrastructure; operators running AI inside a single company more often named cost, data quality, organizational change, or the capacity to review agent output.

Vincent Caldeira captured the operator's compound concern as "a dual challenge of managing volatile token economics and a critical deficit in organizational trust," adding that closing the "Day 2" gap requires replacing "reactive human validation with proactive, machine-speed execution stacks."

François Méro of Probabl named trust the binding constraint, while adding a capability caveat that complicates the study's broader premise that capability is no longer the gate — in his field, he argued, it still is:

> **"The biggest bottleneck is trust. In data science specifically, agents can already code far faster than humans, but they lack the core skills to produce methodologically sound data science that we can confidently trust."**
>
> **— François Méro, CEO, Probabl**

FIGURE 9

WHEN FORCED TO PICK ONE, THE ANSWER DEPENDS ON THE SEAT

Single biggest bottleneck to adoption by 2030, by camp.

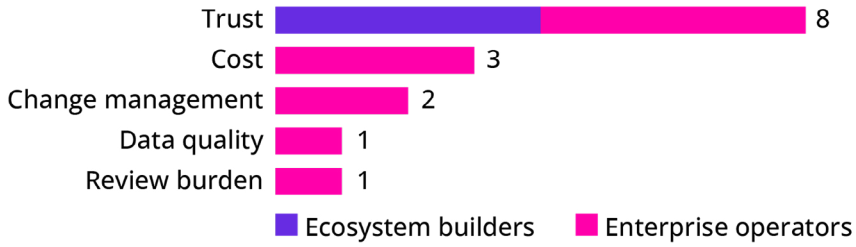

The uneven progress across disciplines was, in his account, the deeper point: “Agents excel at coding, but not at data science because data science is more than just coding. Data science demands skills in methodology and statistical reasoning, which agents currently lack.” One interviewee separately declined to name a single bottleneck at all, holding that the binding constraint shifts as systems mature. This divergence is the basis for treating trust as a connective theme rather than designating it the single determining factor.

RECOMMENDATIONS

FOR
ROI & finance owners

MEASURE COMPLETED OUTCOMES, NOT ACTIVITY.

Define a completion- or resolution-rate metric and a fully loaded cost-per-outcome measure against the human-process baseline; avoid token counts and agent headcount as proxies.

FOR
Program & platform leads

BUDGET FOR EXPERIMENTATION, AND INSTRUMENT IT.

Treat pilot failure as a normal cost of learning, but require each pilot to state its objective and record results, so failures produce insight rather than only expense.

FOR
Strategy & transformation leads

SEQUENCE FOR YOUR REAL BOTTLENECK.

Identify whether trust, cost, data, change management, or review capacity is your real bottleneck, and sequence investment accordingly.

Cross-cutting findings: Shadow AI, adoption culture, and the role of open source

This section consolidates themes that recurred across all four pillars and relates them to the study's initial hypothesis: that enterprises are in a shadow-AI phase in which experimentation precedes formal integration.

Shadow AI, named

Several interviewees independently described ungoverned agent tooling proliferating outside formal oversight, evidence that supports the study's framing hypothesis and reframes governance as an adoption challenge as much as a risk concern. Christian Maitre described "shadow amnesia: tools built, deployed, forgotten, still callable, still connected, but outside any governance perimeter." A broader pattern, sometimes called "shadow agentics," is that pure bottom-up experimentation without governance produces fragmented, ungovernable tool proliferation.

FIGURE 10

HOW SHADOW AI ACCUMULATES

Interviewees described ungoverned tooling following a predictable path. Christian Maitre named the end state "shadow amnesia"; Diego Gosmar named the broader pattern "Shadow agentics."

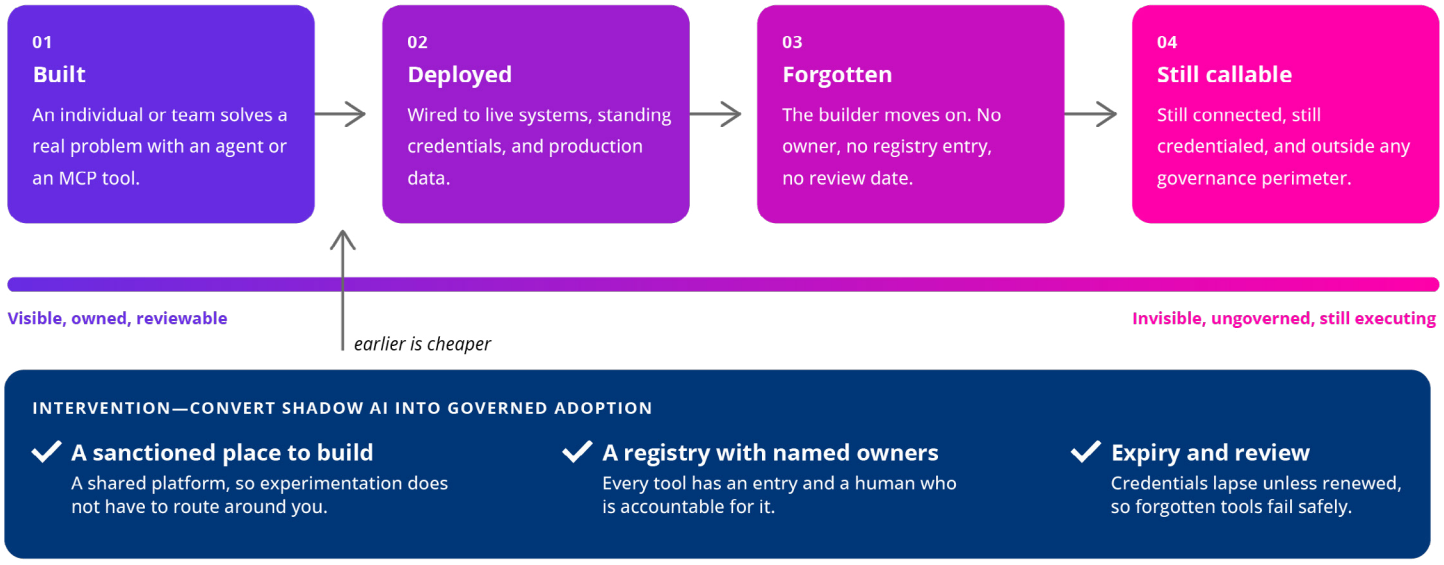

"Tools built, deployed, forgotten, still callable, still connected, but outside any governance perimeter." – Christian Maitre, President, Caspera.lab

Bottom-up experimentation, enabled from the top

Usage typically originated with individuals and teams, and interviewees described leadership's role as funding the platform, setting policy, and removing obstacles rather than mandating use. Zach Lloyd observed that the field is "kind of past the phase of CEOs having to say 'use AI.' People who try it just get it." Others emphasized leadership's enabling role. Hao Li held that "meaningful adoption is primarily top-down," and Emily Chen put the pattern plainly: "Agent innovation often starts at the edges of the organization, but agent transformation requires leadership to scale it." Several described the same dynamic as bottom-up demand meeting top-down enablement.

Angie Jones described a concrete mechanism for making adoption stick: "the 1-9-90 rule, where 1% of a given community are the creators, 9% will tinker, and 90% just consume." She found the 1% in engineering, put them in a cohort with their leads' buy-in to "learn it, experiment, fail, pivot," and bring the lessons back. She watched that group shift to "shipping features two to three times faster than they were two months ago."

Interviewees also noted the coercive edge of mandated adoption, cautioning that top-down mandates without bottom-up proof points produce compliance without conviction. Jonathan Rochelle described executives "letting people go who aren't eating the dog food," a dynamic this paper names rather than smooths over.

Tool selection as the underrated barrier

Interviewees identified the volume and rapid turnover of available tools, rather than the technology itself, as a leading impediment to getting started. Jonathan Rochelle called it "a signal problem," noting that even the most fluent users struggle to keep up, and argued that organizational constraints on tool choice can accelerate adoption: once a company "put some guardrails in," decision paralysis gave way to progress.

Open source: Interoperability, auditability, sovereignty

Few interviewees cited cost savings as the primary rationale for open source, emphasizing instead interoperability, auditability, and sovereignty. Diego Gosmar cast open interoperability standards as "the architectural immune system that prevents today's integrations from becoming tomorrow's lock-in." Mazin Gilbert described the foundation's role as maintaining "a neutral home … so that no single company dictates or owns the technology," drawing the analogy that "the web works because no one owns TCP/IP or HTTP." For clients in ASEAN, MENA, Francophone Africa, and Latin America, several interviewees said the appeal of open source stacks was architectural strength and sovereignty, not cost.

Knowledge and context as the organizational ceiling

Interviewees observed that an agent's effectiveness is bounded by the context available to it. Christine Rimer held that organizations "are only going to be as effective as how well we solve the context and knowledge management problem," and Angie Jones made the point most directly, calling data "the moat": without "the context of your business, how you operate, what's important, what you're building, what your customers care about," teams "still miss some leverage with the agents."

Thomas Garbay observed that a solution's quality depends on "the extent to which human knowledge of the company is part of the agentic building process."

WHEN AGENTS LEAVE THE SCREEN (EMBODIED AND PHYSICAL-WORLD AI)

Two interviewees looked past software to the physical world. Emily Chen offered a striking reframing: “The future of robotics is not Robot-as-Hardware, but Robot-as-Agent. Customers will not buy machines; they will buy autonomous intelligence that delivers outcomes.” She added that “The endgame of agentic AI is not a chatbot on a screen, but an intelligent agent capable of acting in the physical world.” Zach Lloyd framed embodiment as the most compelling frontier: “The thing I’m most interested in is how this translates to the physical world. Robots, medicine, doing things in labs, because it’s such a powerful primitive.” The prominence of this angle in the final report remains an open question for the research team.

RECOMMENDATIONS

FOR Platform & IT leaders

CONVERT SHADOW AL INTO GOVERNED ADOPTION.

Define a completion- or resolution-rate metric and a fully loaded cost-per-outcome measure against the human-process baseline; avoid token counts and agent headcount as proxies.

FOR Executive sponsors

ENABLE FROM THE TOP, ADOPT FROM THE BOTTOM.

Treat pilot failure as a normal cost of learning, but require each pilot to state its objective and record results, so failures produce insight rather than only expense.

FOR Open source contributors & advocates

FRAME OPEN SOURCE AROUND INTEROPERABILITY.

Identify whether trust, cost, data, change management, or review capacity is your real bottleneck, and sequence investment accordingly.

FOR Knowledge & data owners

INVEST IN KNOWLEDGE AND CONTEXT.

Because an agent’s effectiveness is bounded by the context available to it, treat knowledge and context management as core infrastructure, not an afterthought.

Conclusion and actionable insights

Across all four pillars, interviewees converged on a single spine: Capability is no longer the gate. The work that will take place from 2026 to 2030 is not making models more capable, but building the trust, governance, and knowledge infrastructure that lets agents act. Much of that infrastructure is being built, most credibly, in the open. The findings are not uniformly settled, however, and it is worth being explicit about which conclusions rest on consensus and which remain contested.

FIGURE 10

HOW SETTLED IS EACH FINDING

Interview findings grouped from broad consensus to actively contested.

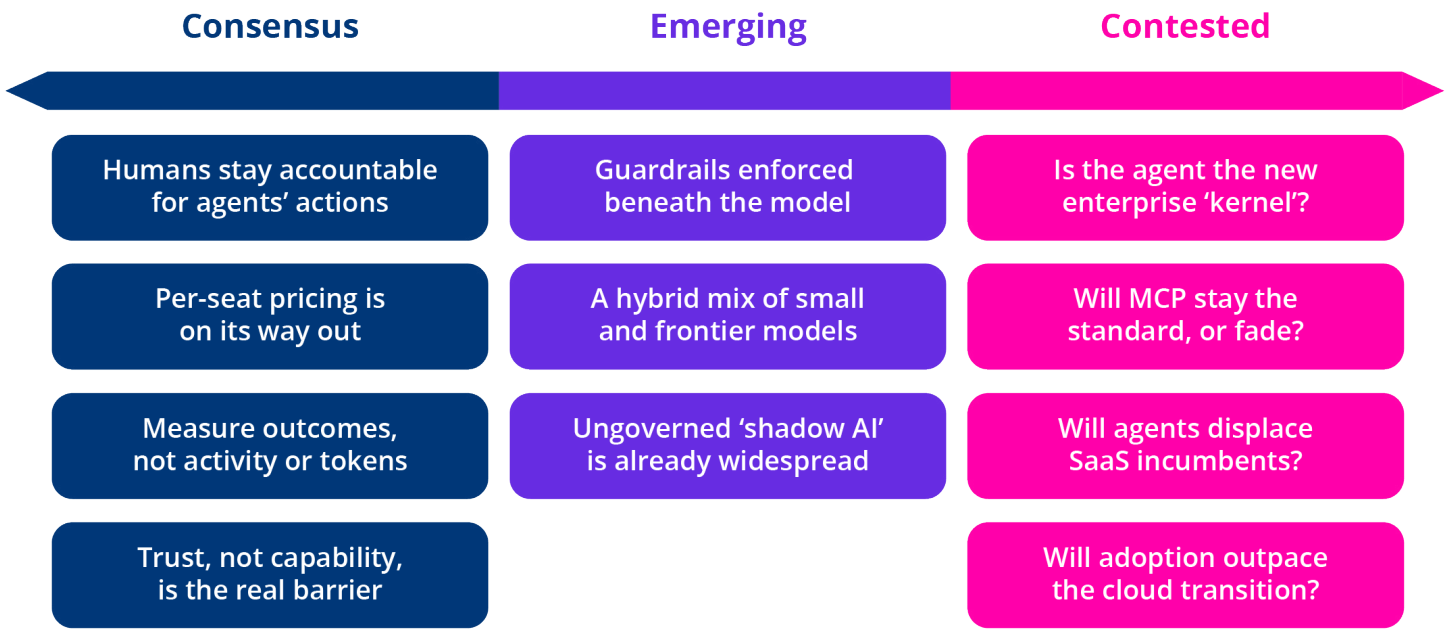

The most settled findings, that accountability stays with humans, that per-seat pricing is giving way, that outcomes should be measured over activity, and that trust is the binding constraint, can be acted on with confidence. Others, including whether the agent is truly a "kernel," how durable MCP will prove, whether established SaaS is genuinely threatened, and how fast the transition will run, remain open, and leaders should hold them loosely.

What this means for enterprise AI owners

For the person who owns AI strategy inside an enterprise, the four pillars resolve into a sequence of concrete moves.

1. **Begin with the bottleneck that is actually yours, rather than the one the market assumes.** Trust recurs across every conversation, but operators just as often named cost, data quality, organizational change, or the human capacity to review agent output as the constraint gating them. Diagnose which it is before you invest, and sequence the work accordingly.

2. **Put the control layer in early.** The interviewees who had reached production treated a governance-and-observability gateway, placed between agents and the models and tools they call, as a baseline rather than a later optimization. They enforced policy beneath the model, at the runtime and operating-system level, rather than asking the model to police itself.

3. **Assign human accountability before an agent ships:** Treat agents as delegated coworkers with bounded authority and a named owner, and begin instrumenting agent-to-agent handoffs now. Responsibility for outcomes that emerge between agents is the frontier no one has yet resolved.

4. **Measure what the agent completes, not what it consumes.** Define a completion- or resolution-rate metric and a fully loaded cost per outcome against the human-process baseline.

5. **Treat pilot failure as the normal cost of learning,** but require every pilot to state its objective and record its result.

6. **Treat the organization's knowledge and context as core infrastructure,** since an agent's effectiveness is bounded by what it can be told about how the business actually works.

7. **Plan for a hybrid transition on both pricing and packaging and models.** Expect a mix of platform fee, consumption, and outcome components for several years, and route routine work to small or local models while reserving frontier models for the reasoning that needs them.

8. **Budget for integration, not only for change management.** Technical quality alone does not ensure agent adoption. Enterprises must also fund the organizational changes required for people to adopt new ways of working, the technical integration needed to embed agents into existing applications and workflows, or, most commonly, a combination of both. This integration and adoption budget can be decisive in whether an agent progresses from proof of concept to sustained production use.

Your innovation strategy is an open source strategy

One theme sits beneath all of these moves, and that is the role of open source as a viable ecosystem of collaboration through which organizations can navigate pathways to AI success. . This collaboration includes the protocols, the identity and discovery mechanisms, and above all the frameworks and guardrails that make an agent safe to trust. For an enterprise, this turns open source from a procurement decision into a strategy forum. Innovation and problem-solving accelerate for the organizations that stay engaged. And the return compounds when engagement runs both ways. An enterprise uses open components, and it contributes back: the edge cases, the fixes, the hardening that shape the standards it will end up depending on. The organization that only consumes inherits whatever the ecosystem produces. The organization that participates helps to influence and define the outcomes.

This is not theoretical. Financial institutions are among the most regulated and risk-averse enterprises there are, and they are already collaborating in the open on exactly these problems. The breadth of enterprise engagement in neutral, vendor-independent ecosystems, such as the Agentic AI Foundation, FINOS, and the LF AI & Data Foundation, is a signal in itself. These are the venues where the interoperability, trust, and guardrail layers are being defined. Participating in them is how an enterprise shapes those layers to its own constraints, rather than inheriting someone else's. On the evidence of this study, that is a prudent path forward.

Two interviewees framed the deeper stakes. Christine Rimer held that "the companies that do well will be the ones with the courage to stop, take a breath, and ask what they're actually solving for, instead of running for the sake of running." And Christian Maitre reframed the entire exercise: "The deeper strategic question is therefore not 'How do we deploy agents?' but 'What kind of organization do we need to be to govern them well?'" The answers to these questions, and so many others relating to organizational agentic AI strategy, invariably are best answered outside of the boundaries of the firm, from the collective wisdom that open source collaborative communities foster.

Methodology

Research design

This study is a qualitative, exploratory inquiry into how organizations are approaching the transition to agentic AI between 2026 and 2030. It was designed to surface the reasoning, priorities, and points of disagreement among the people building and deploying agentic systems, rather than to produce statistically representative estimates. Its findings are directional: they describe the shape of expert opinion and the questions the field is actively working through, and are best read as a map of the current debate rather than a measurement of the market.

Participants and sampling

Participants were selected against a single threshold: direct, current responsibility for agentic AI systems, whether building them, deploying them inside an enterprise, governing them, or stewarding the open standards on which they run. Within that threshold, the sample spans four vantage points: open source and standards leadership, enterprise operators, tooling founders, and independent advisors. It is deliberately international, with the majority of participants based outside the United States and a concentration of European and Australian enterprise practitioners whose regulatory and cost constraints differ materially from their US counterparts. This trades statistical generalizability for evidentiary specificity: where participants' accounts of the same problem diverge along role or regional lines, we report the divergence rather than resolving it.

Data collection

Evidence was gathered through semi-structured interviews and written questionnaire responses conducted between April and July 2026. Each conversation followed a common interview guide organized around four thematic pillars together with a cross-cutting set of questions on shadow AI, adoption culture, and the role of open source. The guide was applied flexibly, allowing interviewees to expand on the areas most relevant to their experience.

Attribution and consent

All participants were invited to contribute quotes for possible inclusion and to review and approve the wording attributed to them. Only approved, attributed quotations and participants' own written responses appear as quotations in this report; no attributed quotation has been paraphrased. Where the analysis summarizes a widely shared view in its own words, it does so without attributing that paraphrase to any individual.

Analysis

Interview material was analyzed thematically. Recurring topics were coded and grouped into the findings presented under each pillar, and each finding is characterized by its degree of agreement across the interviews: from broad consensus, through emerging views, to genuinely contested questions (see Figure 11). Where the report quantifies how often a theme was raised (for example, the barriers to adoption in Figure 8) the counts reflect interpretive coding of the full interview set rather than a closed-response survey; a single participant may be counted under more than one theme, and the underlying attributions are documented in the study's working materials. Throughout, the broad question of which barriers were raised at all is kept distinct from the narrow question of the single biggest bottleneck (Figure 9).

Limitations

As a qualitative study, this report reflects the perspectives of a purposively selected group at a specific moment in a fast-moving field, and its findings should not be read as market-wide measurements. The sample's emphasis on open source practitioners, the self-selection inherent in agreeing to be interviewed, and the interpretive judgment involved in thematic coding all shape the results. Third-party figures discussed in the report, such as published pilot-failure rates, are presented as interviewees raised them and were not independently verified.

Disclaimer

LF Research uses generative AI tooling in varying degrees for research-related tasks including aggregating secondary sources, synthesizing data, supporting report drafting, copyediting, and translation. All AI-generated content is reviewed and revised as needed.

References

Secondary sources cited for industry context (public keynotes and published interviews):

Thomas Kurian, opening and closing remarks, Google Cloud Next 2026, Las Vegas, April 22, 2026. As reported by Capacity, "Google Cloud declares the pilot era over as agentic AI takes centre stage at NEXT 2026."

Jensen Huang, remarks on governing AI and human agents, ServiceNow Knowledge 2026, Las Vegas, May 2026.

Fiona Fung, interview on Lenny's Newsletter podcast, "Building the most AI-pilled engineering team in the world," June 21, 2026.

Acknowledgments

The research team thanks the study's sponsors, peer reviewers, distribution partners, and above all, the 25 interviewees who contributed their time, expertise, and candor.

Interviewees:

- **Richard Bian** — Head of Product & Operations, InclusionAI
- **Guillaume Blaquiere** — Group Data and Agentic Architect, Carrefour
- **Vincent Caldeira** — VP of AI (APAC), Red Hat
- **Emily Chen** — Co-founder, Kaiyuanshe
- **Grace Francisco** — Technology Advisor, m2i
- **Thomas Garbay** — AI Project Manager, Savoir-faire Linux
- **Mazin Gilbert** — Executive Director, Agentic AI Foundation (AAIF)
- **Arnaud Giuliani** — Co-founder, Kotzilla
- **Diego Gosmar** — Head of AI, TESISQUARE
- **Angie Jones** — VP of Developer Relations, Agentic AI Foundation (AAIF)
- **Hao Li** — Founder & CEO, HAOEASY AI
- **Zach Lloyd** — Founder & CEO, Warp
- **Christian Maitre** — President, Caspera.lab
- **François Méro** — CEO, Probabl
- **Ania Musial** — Head of AI Platforms, Bloomberg
- **Rick Nucci** — Co-founder & CEO, Guru
- **Jérôme Oufella** — VP of Technologies, Savoir-faire Linux
- **Christine Rimer** — Independent Technology Advisor
- **Jonathan Rochelle** — Co-founder & CPO, Lutely
- **Xavier Shay** — CEO, Up Bank; former CDO, Bendigo Bank
- **Micaela Stump** — Team Lead PM, Adyen
- **Lin Sun** — VP of Open Source, solo.io
- **Brad Topol** — VP of Engineering, IBM
- **Andreas Traczyk** — Open Source Consultant, Savoir-faire Linux
- **Matt White** — Former Global CTO of AI, The Linux Foundation

About the authors

HILARY CARTER joined the Linux Foundation in 2021 to launch and lead LF Research, established to deliver empirical insights into open source trends, opportunities, and challenges. Prior to joining the Linux Foundation, Hilary was the managing director of the Toronto-based Blockchain Research Institute, a global, syndicated think tank focused on blockchain technology. She has contributed to nearly 200 research projects focused on open source innovation and its adoption across industries. Hilary holds a Master of Science in Management from the London School of Economics, and is a dual Canadian and Irish citizen.

HILLARY CURRAN is Head of Customer Innovation at Guru, where she leads AI enablement, customer education, and the innovation partner program through which enterprise teams put agentic AI into production. She has spent more than nine years at Guru, previously building and running its customer solutions engineering organization across the enterprise and digital segments. Her current work sits close to the questions this study examines: what an enterprise has to change about its knowledge, governance, and operating model before an agent can be trusted with real work. Hillary is a member of the Agentic AI Foundation and holds a Master of Science in Public Service from DePaul University.

ANNI LAI is a longtime open source community builder, advocate, and evangelist. With nearly a decade of active involvement in the global open source ecosystem, she has supported developer communities across more than 40 countries and five continents. She currently serves as Co-Chair of the Generative AI Commons and an elected board member of LF AI & Data, and she also advises LF Europe, GOSIM, and KAIYUANSHE and leads the Open Source Operations at Futurewei. Anni is a passionate believer that AI should serve the global population—rooted in ethics, trust, and openness. Anni holds a B.A. and an M.S. in Computer Science from Rutgers University and San José State University, respectively.

GURU

Founded in 2013, Guru is the governed knowledge layer for enterprise AI. It structures, governs, and continuously improves an organization's knowledge, turning scattered information into a trusted, AI-ready foundation. By enforcing verification, permissions, citations, and policy across every source it connects, Guru gives every person and every AI tool answers they can act on, from a knowledge layer the organization owns outright. Industry leaders like Spotify, Ro, DHL and NJM Insurance are a few of the thousands of companies using Guru as the trusted knowledge foundation driving their AI transformation. Guru has raised $70M in funding and is headquartered in Philadelphia, PA.

FUTUREWEI Technologies

Futurewei maintains ongoing, in-depth collaboration with forward-thinking companies worldwide. We pursue openness in research and development by embracing an open innovation model and striving to share ideas and knowledge with technology communities to create new business opportunities.

Our vision is Shaping the Future Toward a Fully Connected, Intelligent World.

Our mission is Developing Innovations to Benefit an Intelligent and Digital Society via Open Source, Standardization, and Collaboration within Ecosystems.

www.futurewei.com

THE LINUX FOUNDATION | Research

Founded in 2021, **Linux Foundation Research** explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. Through leveraging project databases and networks, and a commitment to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations the world over.

x.com/linuxfoundation

facebook.com/TheLinuxFoundation

linkedin.com/company/the-linux-foundation

youtube.com/user/TheLinuxFoundation

github.com/LF-Engineering

To reference this work, please cite as follows: Hillary Curran, Hilary Carter, and Anni Lai, Enterprise Transitions to the Agentic Era: A Qualitative Study of How Open Source AI Innovators and Enterprise Practitioners are Navigating the Shift from Cloud-first to Agent-first," foreword by Guillaume Blaquière, The Linux Foundation, September 2026.