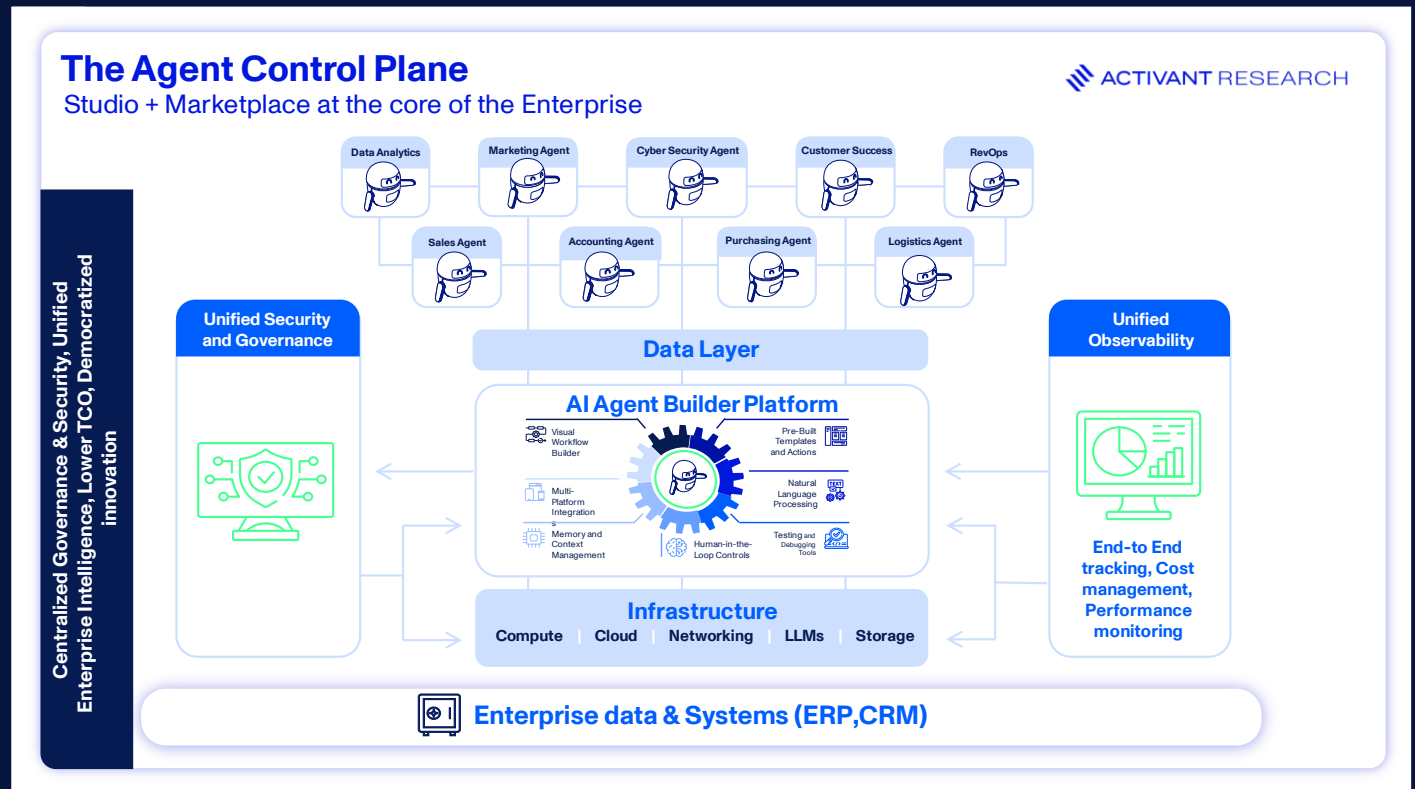




The Agent Control Plane

Build & Buy in One Governed Stack



Scott Watson, Emma Rowand

Q4 2025

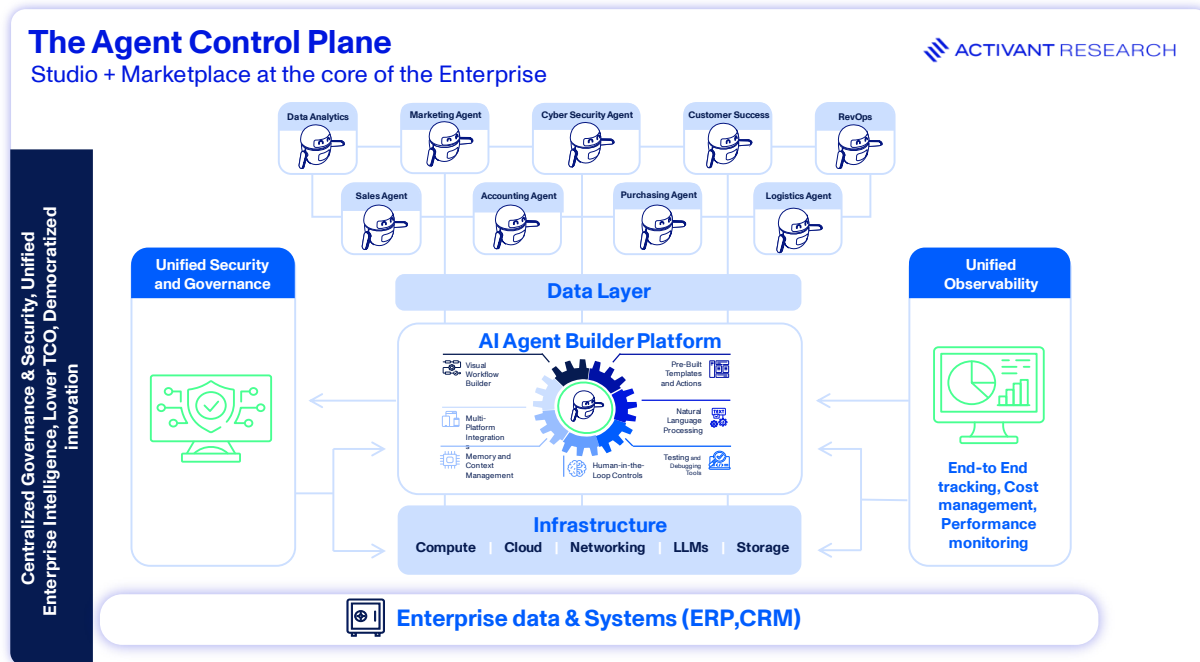
The first enterprise AI mandate was simple: **adopt**. In the ensuing rush, enterprises have inadvertently built a digital Tower of Babel. Twelve months later, leaders are staring at a familiar mess: overlapping copilots, shadow keys, brittle automations, and invoices nobody can explain. This "AI sprawl" is now a direct threat to enterprise security and efficiency. Security teams can't decipher what touched what. Finance can't tie spend to outcomes. Operators are drowning in tabs. This isn't a feature problem or a "pick a better LLM" problem. It's a **control-plane problem**.

We believe the next mandate will be stricter: **govern**. The solution to this AI procurement mess is not another point solution, but a different architecture entirely. One emerging solution is the **AI Agent Builder Platform**, the potential end state for enterprise adoption: a centralized "studio" for an enterprise's entire AI workforce.

These platforms are not just tools; they are marketplaces. We don't expect a winner-takes-all market. Instead, enterprises will rely on a few core platforms to consolidate the number of vendors they engage with. Their power lies in offering a comprehensive set of solutions that caters to the entire organization:

- **For the time-constrained professional** who doesn't want to build: a marketplace to **buy** a vetted, pre-built agent for their specific workflow.
- **For the non-technical power-user** who wants to build: low-code and no-code frameworks to safely **build** and deploy agents that work precisely as desired.
- **For the technical/IT team** who must maintain a competitive edge: the full stack of tools to **build** unique, proprietary workflows, ensuring AI creates differentiation rather than parity.

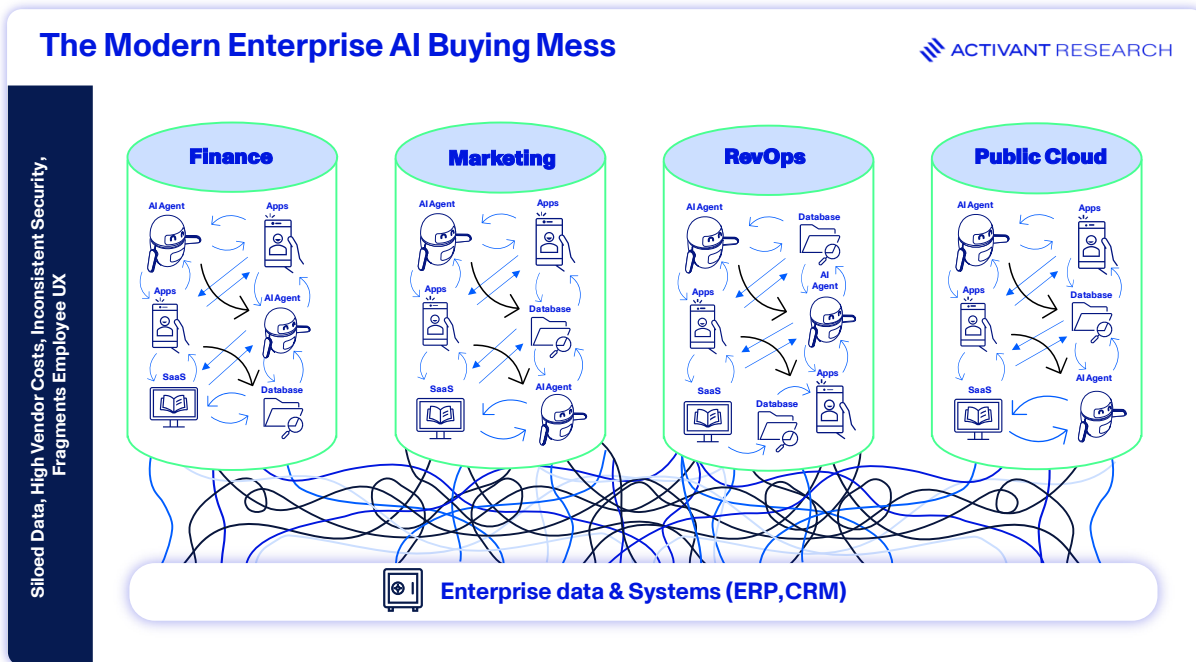
This architecture functions as a unified control plane, transforming AI from a siloed feature into a governed, enterprise-wide system. Visually, this platform sits at the center of the enterprise:



This architecture represents the end state: a rational, secure, and observable AI ecosystem. To understand why this shift is urgent, we must first analyze the chaotic reality most enterprises face today.

How We Got Here: The “AI Frankenstack”

The promise of enterprise software was to break down silos, yet the first wave of AI adoption achieved the exact opposite. Driven by departmental needs and a lack of central strategy, enterprises stumbled into a fragmented procurement process that created more complexity, not less. This chaotic approach results in what we call the "AI Frankenstack," a fragmented, department-led procurement process that creates more complexity than value.



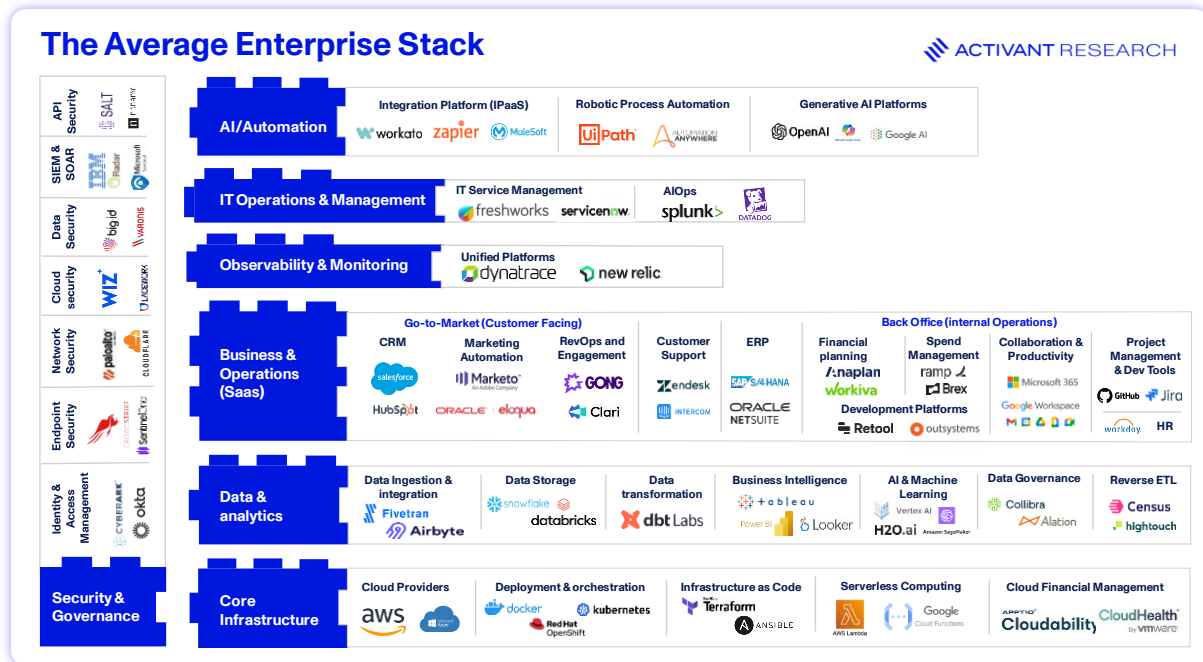
As the diagram illustrates, this well-intentioned collage of tools multiplies risk and hides cost. It's the core driver of what the 2025 MIT AI Report calls the "GenAI Divide." Despite \$30–\$40 billion in enterprise investment, 95% of organizations report zero measurable profit and loss (P&L) impact from their AI pilots.¹

This "tool bloat" has become a major hidden risk. While precise numbers vary, the complexity is undeniable, with numerous organizations managing dozens of overlapping security solutions. Attackers are actively exploiting the blind spots. Organizations with complex, fragmented security stacks face significantly longer times to identify and contain breaches, often taking months longer than those with consolidated platforms, according to IBM's 2025 Cost of a Data Breach Report.²

The pain extends beyond security. This proliferation of vendors imposes a hidden tax on every AI initiative. Each new tool introduces unique APIs, forcing technical teams to spend months on custom integration instead of strategic development. In regulated industries, every new vendor requires a separate, costly cycle of security audits.

The Anatomy of Chaos

This decentralized buying process has created "automation islands," a patchwork of AI tools managed in isolation. Instead of a unified system, enterprises end up with distinct, disconnected stacks for each function. The average enterprise stack has become a complex map of these automation islands.



The diagram provides a simplified illustration of this phenomenon. We consider this a bare minimum view; a mature organization typically has at least one vendor active in each of these subsectors, and often redundant tools within them. The true complexity far exceeds any simplified illustration. Real-world enterprise stacks encompass numerous additional subsectors and layers of technology, with significant variations in complexity from one organization to the next.

The functional siloing accelerates fragmentation:

- **Sales and revenue operations** adopt assistants from [Outreach](#), [Gong](#), and Salesforce to score leads and automate CRM updates.
- **Marketing teams** purchase a different set of tools, like [Jasper](#) and [Marketo](#), to generate content and personalize campaigns.
- **Finance and accounting** implement yet another collection of solutions from vendors such as [Anaplan](#) and [Stampli](#) for fraud detection and bookkeeping.
- **Human resources** uses its own AI-powered software like [Workable](#) or [HireVue](#) for screening resumes and managing onboarding.

This departmental free-for-all has led to adoption fatigue. Some reports indicate the average enterprise uses hundreds of SaaS applications (Okta's 2023 report cited over 300), overwhelming knowledge workers.³

This friction has created the so-called "shadow AI economy." MIT confirms this, finding that while only 40% of firms provide official AI subscriptions, a staggering 90% of employees report using personal AI tools at work. This shadow usage delivers real productivity but operates completely outside of enterprise governance, creating security and compliance risks. It is a clear signal that employees are demanding capable tools, and the official procurement process is failing to keep up.⁴



Enterprises don't have an automation shortage; they have an orchestration problem. Every team is forced to cobble systems together with brittle, outdated scripts and manual triage. At [Tines](#), we're building the connective tissue that transforms those ad-hoc efforts into secure, API-native, intelligent workflows. When security and IT teams can promote their best processes into repeatable, agentic workflows, the organization stops fighting fires and starts compounding intelligence."

Eoin Hinchy

Co-Founder & CEO, Tines

The Solution: The Studio and the Store

The Frankenstack is a temporary and unstable equilibrium. Operational pain, security risks, and spiraling costs are creating market pull for a new operating model.

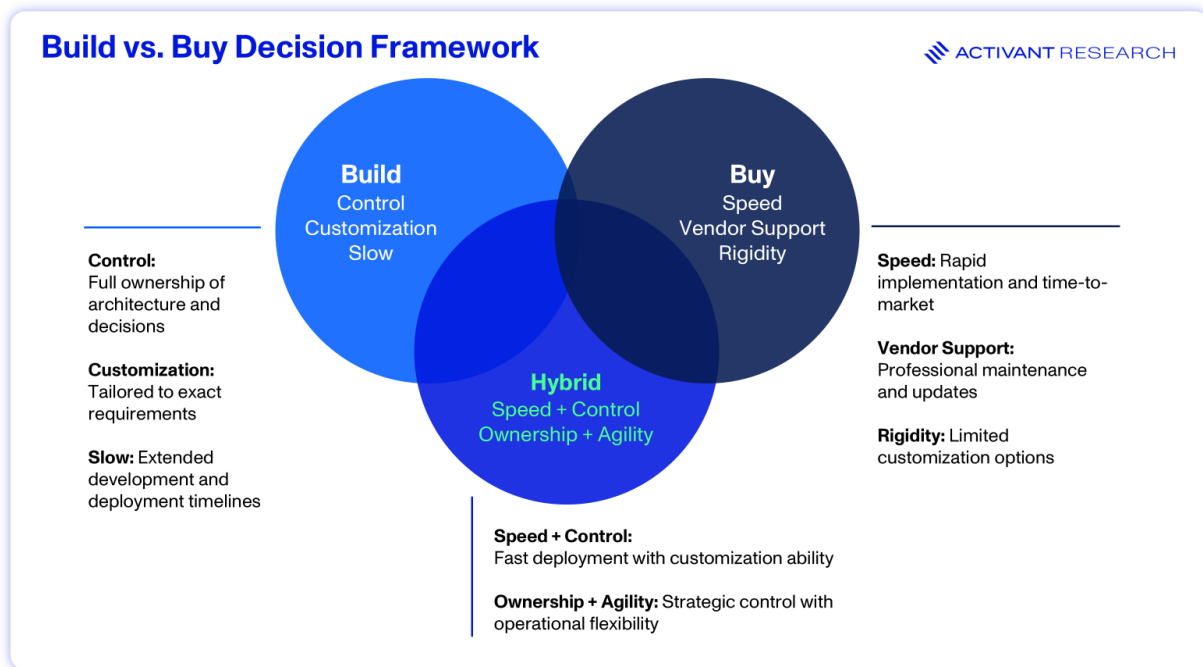
MIT's report identifies the core problem as a learning gap: tools fail because they "do not retain feedback, adapt to context, or improve over time."⁵ This is precisely why the solution is not another point solution, but a platform.

This is where the strategy pivots. The goal is no longer to simply *buy* siloed AI features, but to *build* a cohesive, intelligent workforce. This new platform ecosystem integrates two components: the Agent Studio (the "build" solution) and the AI Marketplace (the "buy" solution).

The clearest parallel is Apple's iOS ecosystem: Xcode is the "Agent Studio" where developers build custom applications, while the App Store is the "AI Marketplace" where users acquire pre-built, vetted solutions. You don't choose one or the other; the power comes from having both within a single, governed environment.

The point isn't to choose build or buy; it's to do both without fragmenting governance. This platform-based approach fundamentally resolves the traditional build vs. buy dilemma.

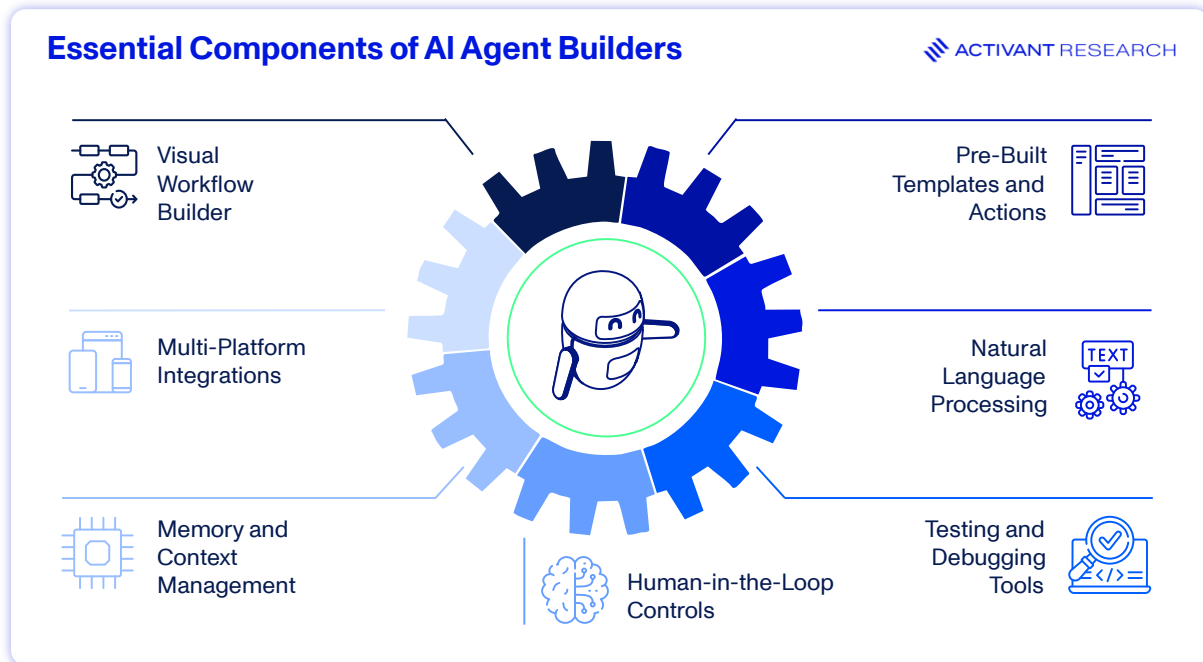
Instead of forcing a binary choice between slow, custom-built tools and rigid, off-the-shelf products, it provides a hybrid model that captures the benefits of both.



This combination of a "build" studio and a "buy" marketplace is what defines the new control plane. But to be effective, that studio component must be more than just a wrapper; it must be a production-grade factory.

What a Real Control Plane Feels Like

An Agent Studio is not another notebook or a clever wrapper. It is a production-grade, opinionated factory that enforces the non-negotiable components required for production AI:



In a studio:

1. Governance & Identity:

- ✓ **Per-agent identity.** Agents never borrow human credentials, running under least-privilege service accounts with clear boundaries.
- ✓ **Policy-as-code at runtime.** Geography rules, PII handling, model choices, and allow/deny lists are enforced automatically — not socially.
- ✓ **Built-in redaction.** Redaction is a capability in the runtime, not a policy document.

2. Context & Tooling

- ✓ **Scoped, mortal memory.** Agents remember only what they should, closing the “forgetful tools” gap.
- ✓ **Tools are contracts, not links.** Every tool is signed, rate-limited, and approval-gated, with explorable logs so you can replay the **why** as easily as the **what**.

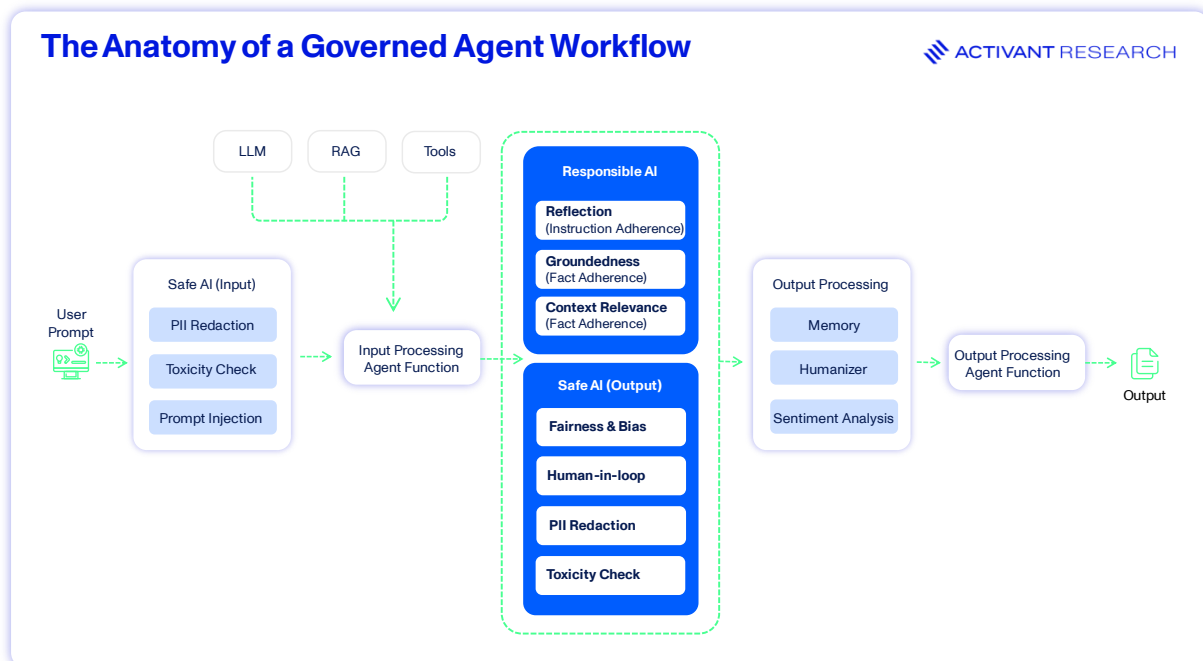
3. Studio & Orchestration:

- ✓ **Visual flows for real work.** Assemble multi-step, multi-agent workflows on a canvas.
- ✓ **Production-ready by default.** Flows are versioned, reproducible, and friendly to continuous integration (CI) and continuous delivery/deployment (CD).

4. Evals & Economics:

- ✓ **Evaluate before you ship.** First-class eval loops, sandboxes, and a reasoning paper trail measure accuracy, safety, and cost prior to GA.
- ✓ **Telemetry you can take to finance.** Every call is traceable, replayable, and priced. Tag spend to workflows so you can defend it in a budget review.

These principles are not just a checklist; they must be architected into the agent's core workflow. A production-grade agent's lifecycle is a sophisticated, multi-stage process of security, policy, and governance checks from the moment a prompt is received to the moment an output is delivered.



This level of built-in governance is what makes an AI agent trustworthy enough for the enterprise. However, even the most secure and well-designed AI will fail if employees refuse to use it. This brings us to the next barrier to adoption: the user experience.

The Mandate for Invisibility

The barrier to successful AI adoption is simple but lethal: enterprises don't need one more tool; they need AI that lives invisibly within the tools employees already use. "Tool fatigue" means any new standalone platform is often perceived as an extra burden, leading to low adoption.

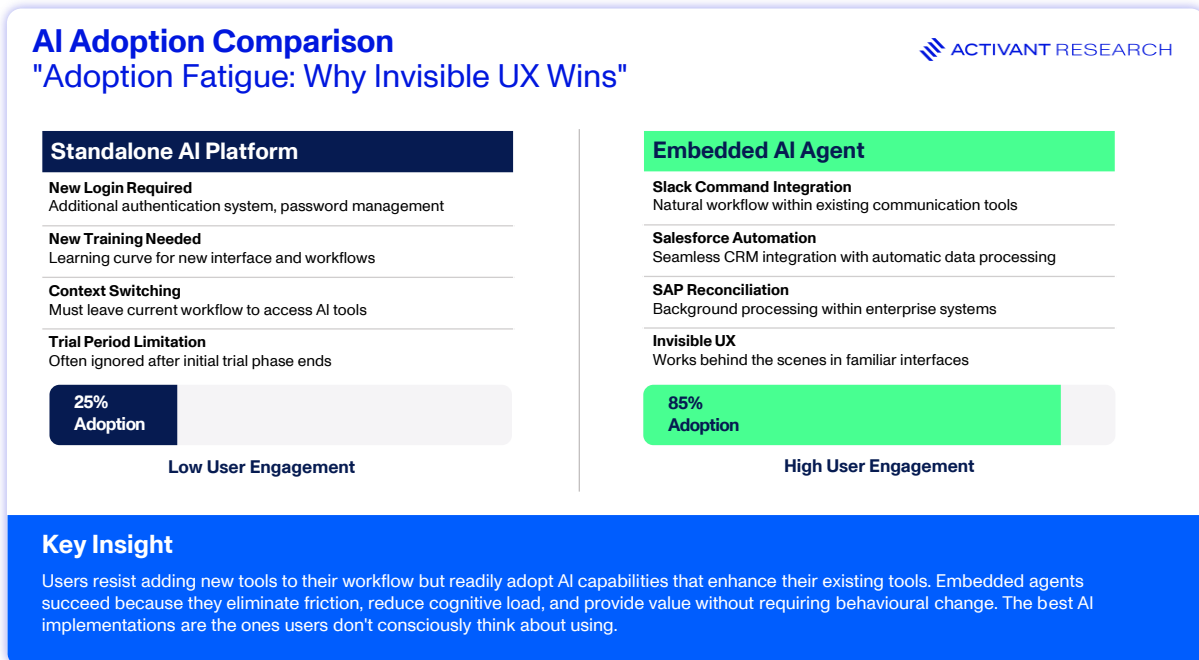
This fatigue is measurable. MIT found that enterprise-grade systems are being "quietly rejected": while 60% of organizations evaluated such tools, only 20% reached pilots and a mere 5% made it to production, most failing due to brittle workflows and poor integration with daily operations. The lesson is simple: if AI isn't embedded where work already happens, it won't be adopted.⁶

Work happens in collaboration suites. Microsoft Teams reported 360 million monthly active users (MAU) as of June 2025, while Slack boasts an estimated 47 million daily active users (DAU) in 2025.^{7,8} Embedding into existing surfaces outperforms launching a new tab.

Successful platforms will prioritize deep, native integrations:

- **Slack- and Teams-Native Agents:** Employees trigger agents directly from chat, eliminating context switching.
- **CRM and ERP Embedding:** A sales agent operates directly inside Salesforce. A finance agent reconciles invoices inside SAP.
- **Invisible UX by Design:** An HR agent appears as an "approve/reject suggestion" inside Workday, not as a separate portal.

The new generation of agent builders is taking this principle to heart. [Gumloop](#), for example, recently launched *Gumloop Agents* with Model Context Protocol (MCP) support, allowing teams to promote no-code flows into tool-using agents that trigger from Slack. It's a subtle but important shift: automation becomes ambient, living inside existing tools rather than beside them.



Making AI invisible removes the primary barrier to adoption and unlocks the productivity gains missing from most pilot-stage projects. The enterprise systems that succeed won't feel like software upgrades at all; they'll feel like work simply got easier.

The Evidence: Why Governance Wins

While the move to governed Agent Studios is not yet widespread, the **catalysts for it** are clear. We see this as an **inevitable shift**, driven by market forces and tangible economic proof points:

- **Adoption is real but uneven:** MIT notes that while over 80% of organizations have piloted tools like ChatGPT, these primarily boost individual productivity, not P&L performance.⁹ The biggest EBIT gains only show up when core workflows are rewired, precisely what Agent Studios are designed to enable.
- **Shadow AI is widespread:** 90% of employees use unapproved AI tools, proving that "ban and block" doesn't work. The only answer is a governed path that's fast enough to win behavior.¹⁰
- **Risk and cost pressure are measurable:** Data breach costs remain stubbornly high, with the global average declining slightly in 2025 but the US average hitting a record \$10 million per incident. AI adds new supply-chain edges and attack surfaces (e.g., Prompt Injection, Model Poisoning) requiring first-class governance.^{11,12}

- **Regulatory gravity is set:** The EU AI Act is in force, with general application starting August 2026 and full effect by 2027. Improvisation accrues liability; a single control plane accrues value.¹³
- **Operator proof points exist:** At production scale, governed agents have collapsed cycle times. Klarna's customer-service agent handled roughly two-thirds of chats, cut resolution from 11 minutes to less than 2 minutes, and was projected to add \$40M in annual profit, though this figure stems from earlier 2024 reports. This is a rare example of the 5% of successful pilots cited by MIT that deliver tangible, multi-million-dollar value.^{14,15}
- **Spending is concentrating:** Forecasts show global AI spend potentially doubling by 2028. As compliance costs rise, buyers will trade dozens of experiments for a few governed platforms that pass audits and scale.¹⁶

Crossing the Production Bridge

Most AI efforts stall in the gap between demo and deployment. MIT calls this the "pilot-to-production chasm." The 5% of organizations that successfully cross it do so by demanding deep customization and holding vendors accountable to business metrics, not just buying "science projects."¹⁷

The bridge between the two is built from a few straightforward ideas:

1. **Define measurable outcomes:** Not "help sales," but "produce a client-ready proposal in three minutes with <2% hallucination rate." If you can't measure it, you can't improve it.
2. **Pin a service-level objective (SLO):** Choose your p50 and p95 latency budgets and decide what happens when they aren't met. No SLO, no promotion.
3. **Evaluate relentlessly:** Treat evals like CI. Every model or prompt change runs the full suite. Failing agents don't ship.
4. **Limit blast radius:** Start with tightly scoped identity and tool access. Expand only as reliability and confidence grows.
5. **Track cost-per-outcome:** Tokens are trivia. Time, quality, and dollars per accepted result are the scoreboard.

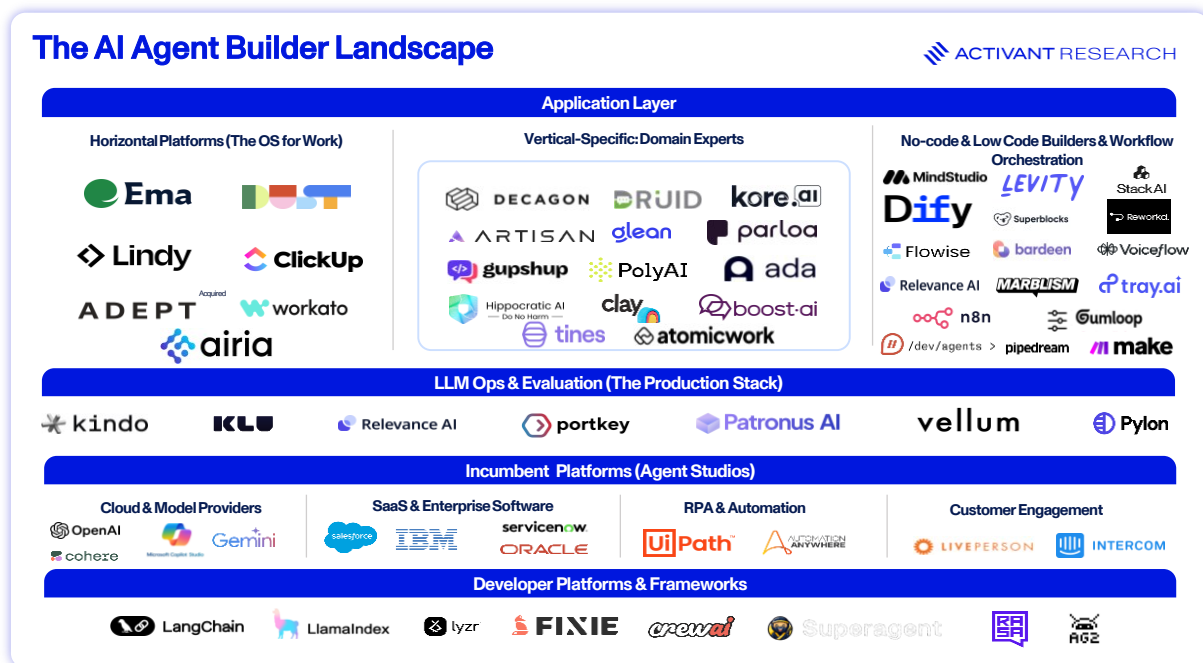
Market Analysis: Sizing the Prize

The opportunity presented by AI agent platforms is not a niche play but a chance to capture a significant share of one of the fastest-growing technology markets in history.

- **Total Addressable Market (TAM):** The broadest view is the global AI market. Grand View Research projects this market to grow from **\$279 billion in 2024** to **\$3.5 trillion by 2033**.¹⁸
- **Serviceable Addressable Market (SAM):** A more focused measure is the Enterprise AI market, estimated by Mordor Intelligence at **\$97.2 billion in 2025** and forecast to reach **\$229.3 billion by 2030** (CAGR of 18.9%).¹⁹
- **The Real Prize (AI Agents):** Within this, the market specifically for AI Agents is projected to grow from **\$5.1 billion in 2024** to **\$47.1 billion by 2030**, expanding at a CAGR of **44.8%**.²⁰
- **Serviceable Obtainable Market (SOM):** This validates our thesis: the agentic paradigm is where new enterprise spending is concentrating. Assuming a winning platform captures 5-10% of the AI Agents market, this translates to a serviceable obtainable market of approximately **\$2.3 billion to \$4.7 billion annually by 2030**.

The Arena: A Field Guide to the Platform Wars

This is an ecosystem race, not a single-vendor sweep. The race to become the dominant platform for the agentic enterprise is intensifying, pitting the scale of tech giants against the focus of new startups. This market is not a single category but a complex, multi-layered stack. To navigate it, we first need a map of the key domains, from the foundational developer frameworks to the application layer where users and agents will interact.



As the map illustrates, this is a complex, multi-layered ecosystem. The race to own it is being fought between two primary forces: the established incumbents, who are leveraging their massive scale and the new, focused challengers.

Legacy Players / Incumbents

The incumbents are not standing still; they are aggressively defending their turf by attacking from different angles:

- **Data Gravity:** If your crown-jewel data sits in **Snowflake** or **Databricks**, their agent runtime (Cortex Agents, Mosaic AI Agent Framework) has a first look at your workloads.²¹
- **Tooling Surface:** Suites with deep hooks into collaboration (**Microsoft's** Graph and M365) or API management (**Google's** Apigee) can make agents feel native (Azure AI Foundry, Vertex AI Agent Builder).
- **Operational Hardening:** Some platforms lead on runtime security and non-human identity (**AWS** Bedrock AgentCore).
- **Workflow Integration:** For enterprises where **ServiceNow** is the system of record, its AI Agent Studio's time-to-value is hard to beat. **Salesforce** is embedding agents via Einstein and Agentforce – a powerful but potentially walled garden.
- **Model-Forward:** **OpenAI's** tight feedback loop between its models and its new **AgentKit** and **Agent Builder** will likely lead the frontier of what feels like magic.

Each incumbent is leveraging its unique, embedded strength, whether it's **data gravity**, **workflow integration**, or **model leadership**, to win the horizontal platform war. They aren't just building agents; they are extending their existing moats to *own* the agentic layer. This is the core reason why competing with them head-on as a horizontal startup is so difficult.

The Activant Thesis: The Three Fronts of the Platform War

The M&A activity in 2024-2025 (ServiceNow acquiring Moveworks, SAP acquiring WalkMe, and Workday acquiring Flowise) isn't the end of the war; it's the sound of the incumbents waking up.^{22,23,24}

They're scrambling to defend their dominance against a new generation of AI-native challengers pioneering the **AI Agent Builder Platform**.

This is a market shaped not by a single winner but by three simultaneous battles:

1. **The Horizontal Platforms** — the fight to build the “OS for Work,” where agents span every department and data system.
2. **The Vertical Platforms** — deep, domain-specific studios that own a single workflow end-to-end.
3. **The Picks & Shovels** — the foundational infrastructure and frameworks that power them all.

Each front demands a different playbook, but together, they define the architecture of the agentic enterprise.

1. The Horizontal Platform (The "OS for Work")

Horizontal is arguably the biggest prize: the race to build the true “OS for Work”, a unified studio where anyone can design, deploy, and manage agents across the enterprise. These platforms go broad, aiming to connect every workflow, model, and dataset under one governed surface.

Two camps are emerging:

The incumbents — **Microsoft**, **Google**, **Salesforce**, and **ServiceNow** — are leveraging their user bases, data gravity, and deep integrations. Their AI layers (Graph, Vertex Agent Builder, Agentforce, and AI Agent Studio) are transforming existing suites into governed agent platforms.

The challengers are building from first principles, often attacking the Frankenstack problem by decoupling the control plane from the underlying models:

- [Ema](#) is crafting an “AI-native workplace” that fuses chat, knowledge, and automation.
- [Airia](#) tackles the governance gap directly, positioning itself as a "model-agnostic" orchestration layer. It provides a centralized control plane where IT can apply unified security policies (like "Agent Constraints") across any model or agent, effectively preventing "shadow AI" sprawl before it starts.
- [Lindy](#) offers programmable personal agents that coordinate tasks autonomously, while [ClickUp](#) is evolving from task management into an agent-orchestrated productivity hub.

To win horizontally, a platform must do more than simply add agents. It must *become* the work surface itself: the single, multimodal interface where every function and file speaks agent.

2. The Vertical Platforms (The "Deep" Play)

This is where the most defensible, immediate value is being created. Unlike the horizontal players, these startups aren't building generic tools; they are constructing deep, studio-like platforms tailored to specific, high-value workflows. Their edge lies in proprietary data, domain-specific evaluation models, and the ability to deliver measurable outcomes that horizontal platforms struggle to match.

Early leaders are already carving out territory: [Cognition](#) is redefining code generation and pair programming; [Sierra](#) is reinventing customer service with autonomous resolution loops; [Artisan AI](#) is attacking the sales-productivity stack; and [Atomicwork](#) is transforming ITSM through tightly integrated agent workflows.

Around them, a second wave of vertical innovators is proving that agentic systems can outperform traditional SaaS in production. [Decagon](#) powers customer-experience agents across chat, email, and voice through an "AI agent engine" already trusted by brands like Hertz, Chime, and Duolingo. In cybersecurity, [Tines](#) (Activant portfolio company) embeds a private "AI Agent" action directly inside SecOps workflows to automate enrichment and case closure without exposing data to external models. [Pylon](#) takes a similar approach in support operations, turning internal runbooks into executable agents that deflect tickets and act across channels while feeding analytics back to managers.

In go-to-market, [Clay](#) uses its *Claygent* to browse and enrich millions of leads in real time— a modern web-scale research layer for revenue teams— while [Dust](#) enables enterprises to connect agents to internal knowledge and deploy them directly within Microsoft Teams. Together, these platforms show that the "vertical" play is less about niche focus and more about depth: owning the workflow, the data, and the outcome.

3. The Picks & Shovels (The "Wide" Play)

The third opportunity is to build the critical, often unglamorous plumbing that every platform—incumbent and challenger alike—will rely on. This wide play underpins the ecosystem and falls into three layers.

Developer Frameworks: This foundational layer provides the building blocks for agent creation and orchestration, long defined by [LangChain](#), [LlamaIndex](#), and [Lyzr](#). These stacks are maturing fast: [LlamaIndex](#) now pairs pre-built agents with a production-grade *AgentWorkflow* orchestrator; [CrewAI](#) blends an open-source multi-agent framework with a Studio/AMP interface for guardrails and observability; and [Lyzr](#)

delivers a low-code *Agent Studio* that moves enterprises from prototype to governed deployment. [LangChain](#) still dominates developer mindshare, anchoring much of the early agentic experimentation, while [Vellum](#) extends the toolchain further, enabling teams to design, version, and test LLM applications with built-in monitoring and prompt management. Together, these platforms form the foundation of the emerging agent development stack.

Automation Infrastructure: Above this sits the connective tissue that links agents to enterprise systems. [n8n](#) leads the open-source camp, expanding its automation fabric to support agentic workflows and LLM nodes across thousands of integrations. [Superblocks](#) takes the enterprise route: its “Clark” agent builds and operates internal apps from natural-language prompts, while on-prem options keep data inside the VPC.

LLM Ops & Evaluation: As organizations deploy agents at scale, evaluation, observability, and governance become non-negotiable. Tools like [Patronus](#) handle testing and evaluation, [Kindo](#) and [Portkey](#) provide observability, and emerging security layers close the loop between performance, safety, and cost. The winners will be the ones that turn experimentation into reliable, auditable production systems.



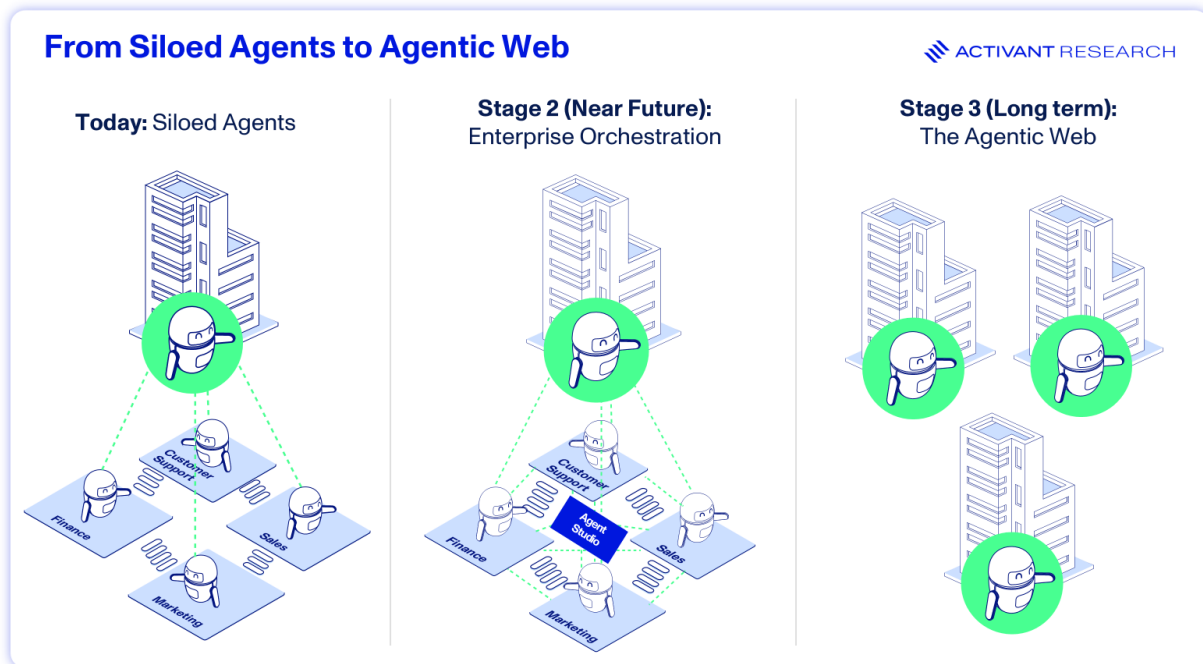
Enterprises will inevitably have 10,000+ agents across platforms like ServiceNow, Microsoft, and Google. The challenge isn't just building them; it's managing the fragmentation. Lyzr provides the complete platform to govern these agents and prevent data from hopping between ecosystems. Our larger vision is OGI (Organizational General Intelligence), which unifies these heterogeneous agents into one central brain to proactively identify insights and continuously enrich the organization's intelligence."

Siva Surendira

Co-Founder & CEO, Lyzr

Where It's Going: From Enterprise Agents to the Agentic Web

Gartner predicts that by 2029, the world will move from today's task-specific agents to collaborative, multi-agent ecosystems. In other words, the future is not dozens of isolated copilots, but an *Agentic Web* — a mesh of interoperable agents that can coordinate across vendors, companies, and domains as easily as web services do today.²⁵



This evolution will feel less like science fiction and more like enterprise plumbing: a collections agent handing a dispute case to a policy agent; a procurement agent negotiating with a supplier bot; or a lineage agent tracing where a number originated, all without a meeting, email, or API call. The winners of the next decade will be those who master interoperability, standardization, and security across this emerging agent-to-agent economy.

A glimpse of that future is already visible in companies building the early infrastructure. One is [/dev/agents](#) (Activant portfolio company), led by ex-Google, Meta, and Stripe engineers. They are developing an “operating system for AI agents,” a runtime designed for cross-agent collaboration, permissions, and shared state. It represents the earliest scaffolding of a truly open Agentic Web, where agents are no longer confined to silos but can communicate safely across organizational boundaries.

The implication for enterprises is profound: the control plane you build today must not only govern internal agents but be ready to interoperate externally. Governance and interoperability will converge, and the companies that design for that intersection early will define the standards everyone else builds on.

Economics a CFO Will Sign

AI will not be cheap in absolute dollars. It will be valuable if it collapses a real cost curve or opens new revenue. The easiest way to make this objective is to price the outcome, not the inference.

Today's Frankenstack carries a massive, hidden Total Cost of Ownership (TCO) from redundant licensing, custom integration, and fragmented security reviews. A unified platform fundamentally changes this equation.

The discipline is simple:

1. **Consolidation Savings:** Retire overlapping licenses, kill custom glue, and cut audit overhead by centralizing controls.
2. **Outcome Lift:** Cycle time down, win rate up, error rate down. Tie gains to an owner and a ledger.
3. **Platform Cost:** Fees for the studio and marketplace usage.

ROI = (Consolidation + Outcome Lift) – Platform Cost.

Once telemetry ties spend to workflows, finance can ask the most important question: "What is our cost per accepted proposal? Per cleared invoice? Per vendor review?" When that exists, budgets shift from tool line-items to conversations about efficiency.

The Inevitable Consolidation

The first phase of enterprise AI was a chaotic land grab, leaving enterprises with a fragmented Frankenstack that multiplies risk, hides cost, and limits real P&L impact.

The next phase will be one of deliberate consolidation. The signals are clear: operational pain, the massive security risks, and accelerating M&A are pushing enterprises toward a more rational operating model.

This is the **AI Agent Builder Platform**.

We believe a centralized "Studio" to **build** and "Marketplace" to **buy** is the inevitable end state, because it directly solves the core failures of the Frankenstack.

It is the only architecture that allows an enterprise to:

- **consolidate** vendor sprawl,

- **govern** all AI activity from a single, auditable control plane, and
- **end** the "build vs. buy" stalemate by combining pre-built agent speed with custom-built workflow differentiation.

For enterprise operators, the mandate is simple: stop funding the chaos. Stop buying point solutions and start architecting the control plane.

For founders, the opportunity is to build the fabric of this new era. The winners will be those who compete to build the new **Horizontal Platform** "OS for Work," the deep **Vertical Platforms** for specific industries, or the critical **Picks and Shovels** that the entire ecosystem relies on.

The opportunity is not just to build a better tool, but to build the new, cohesive, and intelligent operational layer of the modern enterprise.

If you are a founder building that platform, Activant Research wants to hear from you.

Disclaimer: The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice, and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at <https://activantcapital.com/companies/> for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.

-
- ¹ mind the PRODUCT, [Why most AI products fail: Key findings from MIT's 2025 AI Report](#), 2025
- ² BAKER DONELSON, [Ten Key Insights from IBM's Cost of a Data Breach Report 2025](#), 2025
- ³ Avatier, [Password Synchronization: Why Avatier Handles Passwords Better Than Okta](#), 2025
- ⁴ mind the PRODUCT, [Why most AI products fail: Key findings from MIT's 2025 AI Report](#), 2025
- ⁵ Ibid
- ⁶ Ibid
- ⁷ SQ Magazine, [Microsoft Teams Statistics 2025: How Teams Dominates Collaboration](#), 2025
- ⁸ demandsage, [Slack Statistics \(2025\): Revenue & Number of Users](#), 2025
- ⁹ MIT, [The GenAI Divided: State of AI Business 2025](#), 2025
- ¹⁰ Ibid
- ¹¹ BAKER DONELSON, [Ten Key Insights from IBM's Cost of a Data Breach Report 2025](#), 2025
- ¹² INDUSFACE, [OWASP Top 10 LLM Applications 2025 – Critical Vulnerabilities & Risk Mitigation](#), 2025
- ¹³ European Parliament, [AI Act implementation timeline](#), 2025
- ¹⁴ PYMNTS, [Klarna CEO Says AI Will Drive 'Massive Shift' in Employment](#), 2025
- ¹⁵ mind the PRODUCT, [Why most AI products fail: Key findings from MIT's 2025 AI Report](#), 2025
- ¹⁶ Stanford University Human-Centered Artificial Intelligence, [The 2025 AI Index Report](#), 2025
- ¹⁷ MIT, [The GenAI Divided: State of AI Business 2025](#), 2025
- ¹⁸ Nasdaq, [Better Artificial Intelligence Stock: Nebius Group vs. ASML](#), 2025
- ¹⁹ Mordor Intelligence, [Enterprise AI Market Size & Share Analysis - Growth Trends And Forecast \(2025 – 2030\)](#), 2025
- ²⁰ Datagrid, [The Evolution of AI Agents in Sales: 30 Key Statistics & Adoption Trends](#), 2025
- ²¹ Databricks, [Mosaic AI Agent Framework](#), 2025
- ²² ServiceNow, [ServiceNow to extend leading agentic AI to every employee for every corner of the business with acquisition of Moveworks](#), 2025
- ²³ Ibid
- ²⁴ Workday, [Workday Acquires Flowise, Bringing Powerful AI Agent Builder Capabilities to the Workday Platform](#), 2025
- ²⁵ IntuitionLabs, [Agentic AI Workflows: Why Orchestration with Temporal is Key](#), 2025