CAMUNDA

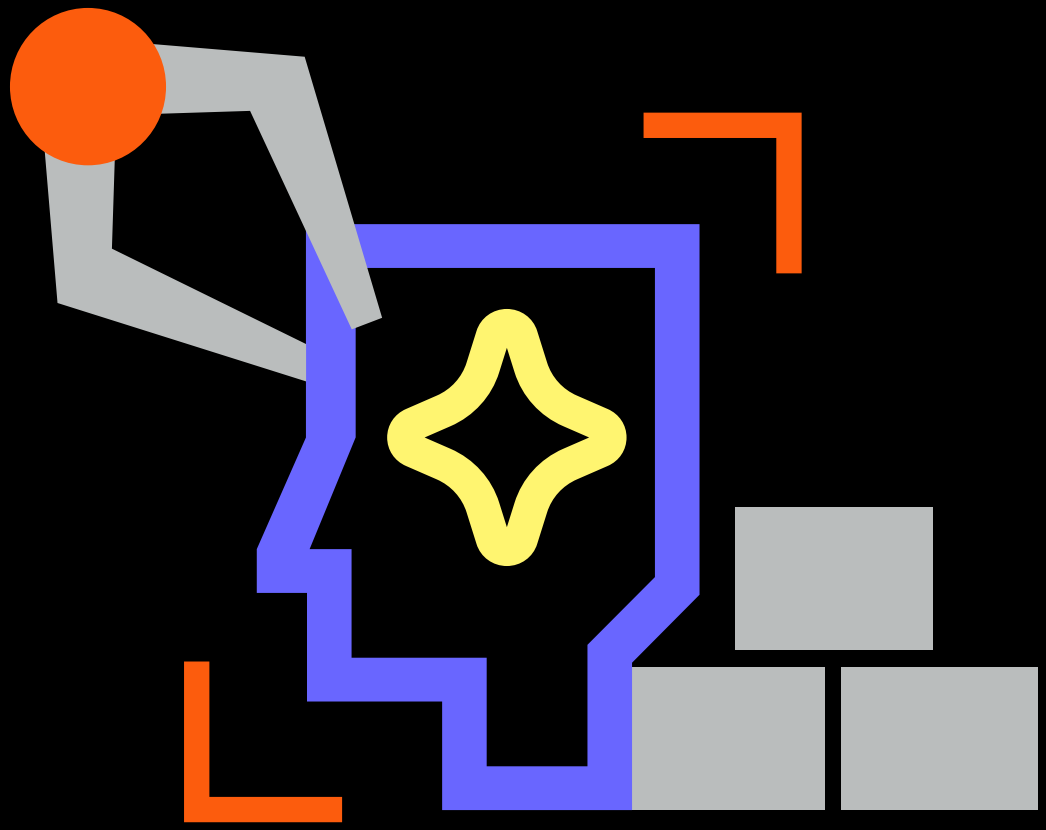

To Build or to Buy AI Agents?

How Agentic Orchestration Unlocks True Agent Value

February 2026

camunda.com

Contents

Introduction: Why the build-vs-buy question is exploding

AI agents have moved quickly from innovation labs into enterprise roadmaps. In regulated industries — such as banking, financial services, insurance, healthcare, and public sector — the pressure is not simply to "use AI," but to deliver outcomes that are trusted, auditable, and durable in real operations. That is a very different requirement than running an isolated pilot or deploying a chat-based assistant in a single department. It also explains why the build-versus-buy discussion has become so urgent when it comes to agents.

CIOs are juggling more AI work than ever. Most enterprises now run five or more AI projects, with nearly a third running ten or more[1]. The portfolio is expanding at the same time the automation tooling landscape is fragmenting across pro-code, low-code, and no-code options, which makes a single "right" decision harder to defend. It is also why many organizations experience pilot sprawl. The easiest path is to start, but the hardest part is to operationalize.

> The operationalization gap shows up clearly in agent adoption data. According to the State of Agentic Orchestration and Automation Report, 71% of organizations are using AI agents, yet only 11% of these agents have made it into production. Even when agents exist, they often run as isolated capabilities: 48% of organizations say their AI agents operate in silos and aren't woven into the end-to-end business process.

That is why "build or buy?" is often the wrong first question. A better question is whether the enterprise has a reliable way to put agents to work inside real business processes, with controls that match the risk profile of the work. In practice, the deciding factor is rarely the agent itself. Orchestration determines whether an agent is a durable business capability or another disconnected experiment.

1 Gartner, *How to Decide Whether to Build, Buy or Blend Your AI Projects, Whit Andrews, Jim Hare, 9 April, 2025*

The market forces behind the build-vs-buy debate

Customer experience is pushing agents into the core

Customer experience remains a top differentiator, especially in industries like financial services where user journeys cross both modern and legacy systems and customers expect personalization. AI agents promise the ability to adapt based on context, which is exactly what makes them attractive for core customer-facing operations. They also promise to automate knowledge work that traditional automation tools struggle to reach, such as correlating customer information across multiple platforms and using that context to determine next best actions.

The risk is that agents improve one interaction while the broader business process remains broken. A customer might have a great conversational experience with an agent, and still be forced into manual follow-ups because the agent cannot reliably trigger downstream verification steps, coordinate asynchronous tasks, or handle exceptions. When agents are not orchestrated, they tend to optimize single moments in time, instead of overall outcomes.

Composable architecture raised expectations, and exposed new gaps

Many enterprises are adopting composable automation architectures because they are flexible and let teams swap capabilities without rebuilding the whole tech stack. That flexibility enables faster adaptation to shifting business needs and changing technical requirements. It can also support cost management by making reusable components easier to share and maintain. Perhaps most importantly, it creates room to experiment with new technologies and incorporate them once they prove valuable.

At the same time, composability increases the need for coordination. Purchased agents (e.g. an agent within a proprietary tech platform) rarely integrate deeply across a heterogeneous stack, and built agents need a stable backbone to participate in enterprise processes that span people, systems, and devices. Gartner[2] notes that, “When choosing an AI agent development framework, software engineering leaders should direct their teams to assess the framework’s ability to support scalable, secure and efficient multiagent workflows while aligning with enterprise architecture and governance requirements.”

Agentic AI accelerated pilots, but also increased failure modes

Agent-based approaches promise dynamic, adaptive decision-making, but they often remain brittle or high-risk without governance, memory, or orchestration. That brittleness shows up in stalled programs and scrapped initiatives. MIT research shows that only 5% of AI pilot programs ever deliver measurable revenue impact. Another study from Gartner shows that 40% of agentic AI projects are expected to be scrapped by 2027 because they fail to create value. Organizations are also rightfully worried about the risks and inaccuracies of AI, with 44% of enterprises reporting incorrect AI outputs.

This is the context in which the build-vs-buy debate is happening. Enterprises are being pulled toward agents by competitive pressure, while being pushed away by risk and operational complexity.

2 Gartner, *How to Choose the Right Technology to Build LLM-Based AI Agents, Tigran Egiazarov, Jason Wong, Gary Olliffe, Jim Scheibmeir, Arun Batchu, Adrian Leow, 9 June 2025*

Blending a build and buy strategy

For most enterprises, the build-versus-buy question does not resolve cleanly in one direction. Instead, it evolves into a blended strategy shaped by risk tolerance, internal capability, and the nature of the business processes involved. Buying and building each introduce different tradeoffs, and neither approach consistently satisfies the full range of enterprise requirements on its own.

What “buying an agent” typically looks like

Buying typically means selecting pre-built copilots or domain-specific agents tied to a single platform. Often, this is not a conscious procurement decision. It is simply the path of least resistance: teams build agents inside platforms they already use because the data access and permissions are already there.

A common example is building an agent inside Salesforce if Salesforce is the system of record for a specific function. For non-technical users, this approach can be appealing because they do not have to solve data access, tool selection, or integration patterns from scratch. The tradeoff is that these agents are limited by the boundaries of the application. In other words, they work well inside Salesforce, but struggle when the process includes the other applications or human steps.

This is where agentic orchestration is most effective. Ideally, agentic orchestration lives outside the bounds of any single application. It coordinates agents’ work across systems, manages process state, and ensures agents have access to the right context no matter where data lives.

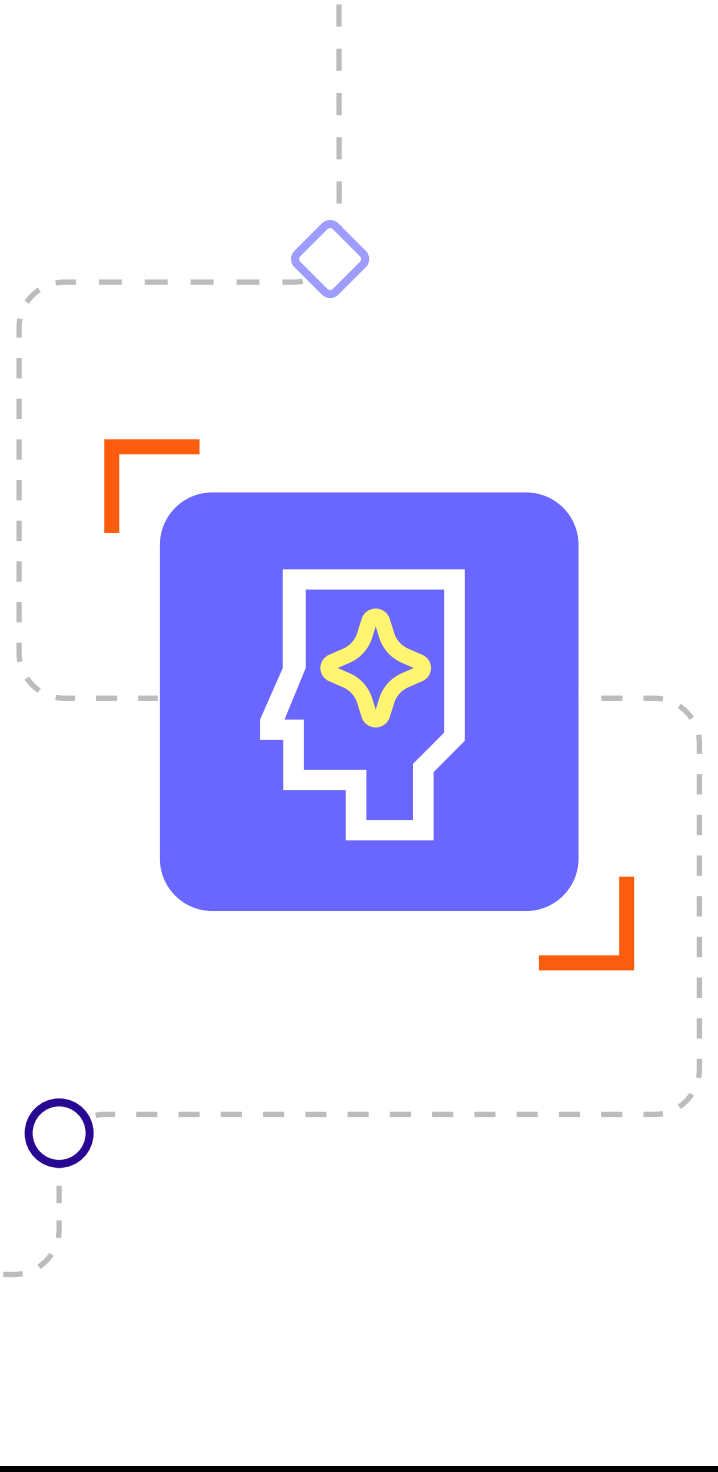

What it means to “build an enterprise-grade agent”

Building agents introduces a different set of dynamics. Custom-built agent often increase architectural and operational complexity, particularly when teams attempt to manage state, memory, reasoning loops, and integration logic themselves.

However, when building occurs within the right orchestration framework, that complexity becomes more manageable. Agents can be designed to align directly with existing business processes, policies, and data flows, making them more naturally compatible with how work already happens across the enterprise.

Why blended approaches are common

In practice, agentic orchestration is the key variable that determines whether the build or buy approach succeeds. Without orchestration, buying leads to siloed agents that improve isolated tasks but fail to influence end-to-end automation outcomes. Building, without orchestration, results in fragile experiments that struggle to scale or remain governable. With orchestration in place, both built and bought agents can participate in durable business processes that preserve context, accountability, and control.

Orchestration also plays a decisive role in how agents express autonomy. Purchased agents typically operate with lower autonomy. They excel at guided tasks, question-and-answer interactions, or channel-specific support, where decisions are constrained and the surface area of risk is limited. This can be entirely appropriate for certain use cases, particularly those focused on productivity or front-line assistance.

Built agents, by contrast, often enable higher levels of autonomy. They can reason over context, plan sequences of actions, and execute multi-step business processes that span systems and people. That autonomy can deliver significant value, but it also increases the need for governance, explainability, and human oversight.

Most enterprises need both. They need mixed autonomy, where deterministic process logic provides predictability and compliance, and AI agents enable dynamic reasoning where variability and judgment are unavoidable. In this model, autonomy is not an all-or-nothing decision. It becomes a dial organizations can adjust based on process risk, regulatory exposure, competitive changes or market forces, and business criticality. Agentic orchestration is what makes that adjustment possible, allowing organizations to scale agentic capabilities without losing control.

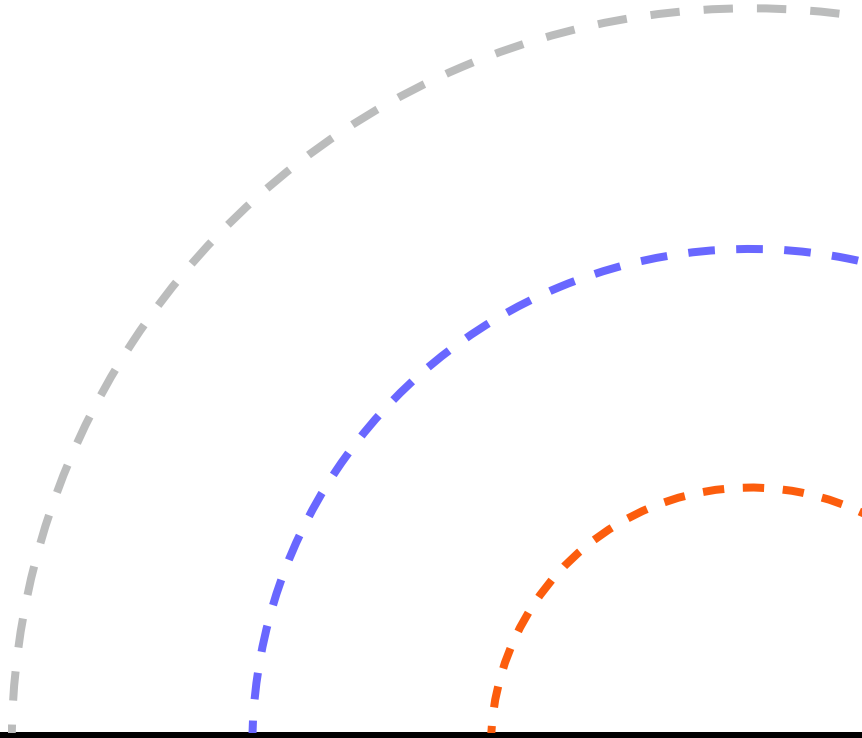

Decision framework: When to buy vs. build

The strongest build-vs-buy decisions start with the business process, not the agent.

When buying makes sense

Buying tends to make sense when a process is standardized, complexity is low, and time-to-deploy matters most. This aligns with Gartner's "Defend" category[3] for AI — augmenting productivity through capabilities that are easier to adopt and easier to constrain.

In these cases, the goal is usually to improve speed and consistency within a defined surface area. If the outcome can be achieved without deep cross-system orchestration and without high-risk autonomous decisions, buying can be a pragmatic choice.

When building makes sense

Building makes sense when a process is high value, heavily regulated, or requires extensive transparency and oversight. It also makes sense when the enterprise needs reuse across multiple business processes, or when composability is essential because no single vendor tool can own the whole journey.

Building becomes especially compelling when organizations need to shape agent autonomy precisely — dialed up in low-risk segments of the process, and dialed down where human judgment or compliance review is required.

When hybrid models are best

Hybrid models often win when organizations face contextual constraints, such as being restricted to one or more approved LLMs or needing to host an LLM locally. These constraints introduce architectural and operational tradeoffs that make a single-vendor approach risky, while also making a fully custom approach expensive and difficult to future-proof.

A hybrid approach can preserve flexibility while still meeting constraints, but only if orchestration provides a stable control plane. Without that, "hybrid" becomes another word for fragmented.

3 Gartner, *How to Decide Whether to Build, Buy or Blend Your AI Projects, Whit Andrews, Jim Hare, 9 April, 2025*

Why agentic orchestration is the deciding factor

Composable technology stacks make it easy to add agentic capabilities, but that same ease can easily create sprawl. Purchased agents can quickly become isolated inside one application, while built agents can turn into disconnected experiments owned by individual teams. Enterprises need a way to connect deterministic business logic, dynamic agent reasoning, and human oversight into one durable end-to-end process.

Agents need a conductor

In an enterprise setting, deterministic logic is essential for compliance and repeatability. Dynamic reasoning is essential for adaptability and contextual decisions. Agentic Business Process Model and Notation (BPMN) blends both inside a single executable model, allowing teams to design processes where agents can plan and act within auditable boundaries.

This is why the build-vs-buy debate is often a false binary. Both built and bought agents can be governed, monitored, and made durable when they operate inside a consistent agentic orchestration framework.

Example: An orchestrated mortgage application journey

A mortgage application is a useful lens for understanding how build-and-buy strategies converge in practice, because it combines regulated decision-making, high customer expectations, and complex, asynchronous business processes. No single agent or tool can own the journey end to end. What matters is how specialized capabilities work together inside a governed process.

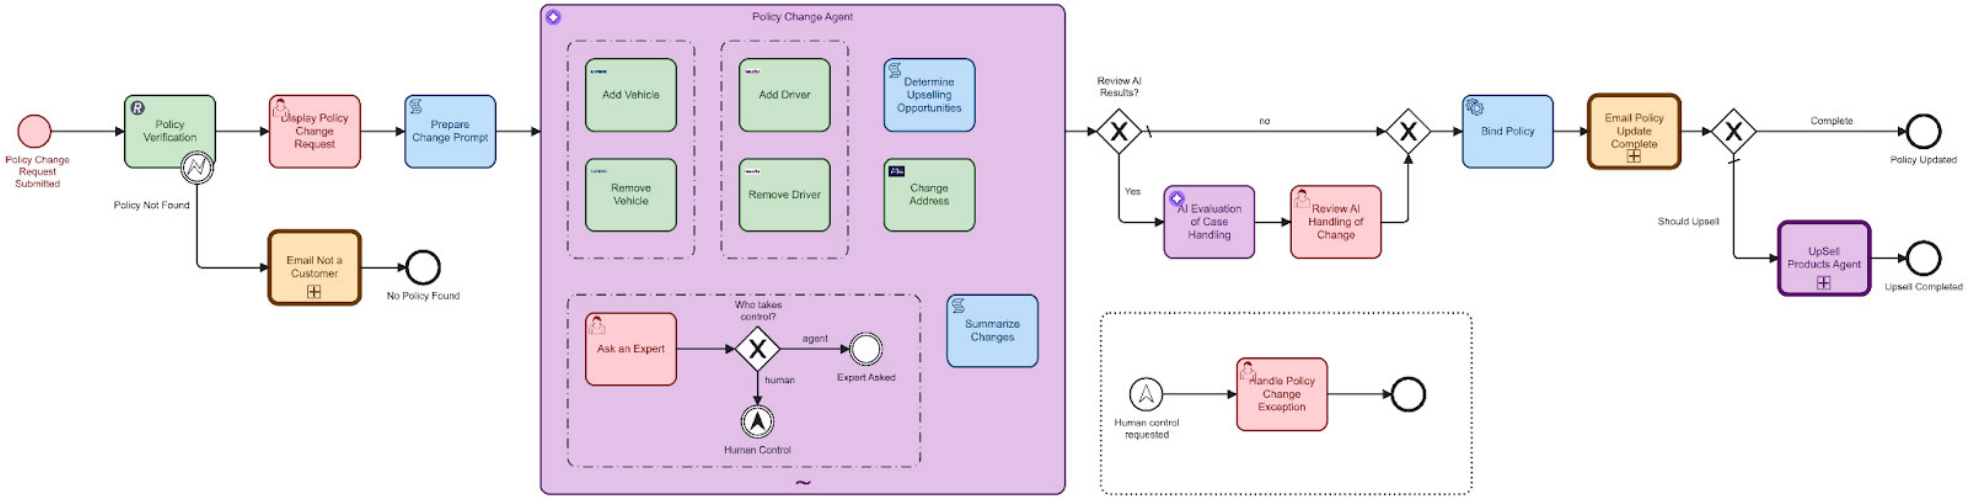

At CamundaCon New York 2025, Camunda demonstrated how a composable digital experience architecture can support this kind of journey when it is orchestrated end to end. In this example, Camunda integrates with a variety of third-party technologies, which are acting as an intelligent agent or interface to a mortgage application.

- **Liferay provides the experience layer.** It serves as the unified entry point where customers interact with services, content, and agents across channels. By managing identity, access controls, and continuity, Liferay ensures that the customer experience remains consistent and governed from the first interaction through completion.

- **GetVocal operates as a conversational agent embedded within that experience.** It blends deterministic AI with human oversight to guide customers through the application process via voice, chat, or messaging. The agent collects information, answers questions, and takes predefined actions, escalating to a human when confidence thresholds are not met. Each interaction is protocol-driven, explainable, and fully visible, which is essential in regulated environments.

- **Camunda functions as the orchestration backbone.** Rather than acting as an agent itself, it coordinates the entire journey, managing state, sequencing tasks, and enforcing process-level governance. As the application progresses, Camunda triggers backend processes such as identity verification, document handling, and asynchronous credit checks. It ensures that each step executes in the correct order, with full auditability and resilience.

- **Taktile serves as the agentic decisioning layer.** It autonomously retrieves relevant data, applies expert-defined policies, and executes the business logic required to assess risk, determine eligibility, and personalize offer terms. While the underwriting logic is agent-driven, guardrails and decision criteria are defined by human experts, ensuring transparency, consistency, and regulatory alignment.

When a process encounters uncertainty or risk, Camunda can route the case to a different agent, deterministic logic, or a human reviewer with full context. In the case of a human review, the reviewer does not start from scratch. They see the data the agents used, the decisions that were made, and the reasons behind them. Once resolved, the process continues without losing continuity.

This is agentic orchestration in action. Bought capabilities such as conversational interfaces and decisioning services are combined with built process logic and governance. Deterministic orchestration provides predictability and compliance. Agentic orchestration delivers contextual reasoning and adaptability. Together, they enable end-to-end business processes to work seamlessly across a modular technology stack, without sacrificing control.

The coming wave of “buyable” agents as business solutions

The market is heading toward horizontal, agent-powered solutions that look more like packaged business capabilities than chat assistants. These solutions will proliferate, but they will still require orchestration to meet enterprise needs around compliance, durability, and adaptability.

This is where agentic orchestration has a clear job. It must let teams buy accelerators when speed matters, and build differentiation where it creates competitive advantage. Most importantly, it must orchestrate both built and purchased solutions, so that agent adoption does not fragment into tool-specific silos.

Pros and cons of buying vs. building

Buying agents

Pros

- Faster deployment and time to value
- Lower upfront investment
- Predictable early performance within a constrained scope

Cons

- Limited customization and flexibility
- Decision logic often remains siloed within a single platform
- Harder to scale across multiple business processes
- Governance and accountability may be weaker without orchestration
- Long-term cost risk if pricing changes or vendor dependency increases

Building agents

Pros

- Tailored to enterprise-specific business processes
- Greater control over autonomy and decision boundaries
- Integrated governance and human-in-the-loop design
- Reusable components across business processes
- Improved visibility into agent behavior and decisioning

Cons

- Requires specialized skills and engineering effort
- Higher upfront investment
- Slower rollout without orchestration to ensure durability and reuse

Why orchestration matters

- Reduces the rigidity of buying by integrating purchased agents into end-to-end processes
- Mitigates the risk of building by providing governance, durability, and reuse
- Enables both built and bought agents to operate as enterprise-grade capabilities

Conclusion: Build vs. buy is less decisive when you orchestrate

Enterprises ask “build or buy?” because they want to reduce risk, manage cost, and move quickly. Those are rational goals. The problem is that the choice becomes much less predictive of success when agents are not operationalized inside end-to-end business processes.

Success comes from blending autonomy, governance, and composability. Agentic orchestration turns agents (whether built or bought) into trusted business capabilities that can run at scale with the same rigor as any core business process.

When orchestration is treated as the control plane, enterprises can:

- Adopt agents without losing accountability
- Dial autonomy up or down based on process risk
- Measure outcomes instead of counting pilots.

Over time, agents will stop being experiments and become part of the way the organization runs its core processes.

Want to learn more about unlocking value from AI agents in your automated business processes?

Contact our team today

CAMUNDA

About Camunda

Camunda is the leader in enterprise agentic automation, orchestrating complex business processes, including high-value knowledge work, across agents, people, and systems. By creating production-ready, enterprise-grade agents with built-in governance, Camunda uniquely delivers trusted AI agents for business-critical processes. Over 700 leading innovators like Atlassian, ING, and Vodafone, rely on Camunda to slash time-to-value from months to days, boost operational efficiency, and elevate customer experiences. Ready to become an AI-first enterprise?

Visit camunda.com.