

Orq.ai Launches Standalone AI Router, Addressing Enterprise Sovereignty and Cost Control Demands

European AI development platform responds to rising AI costs, infrastructure lock-in and sovereignty demands as routing becomes a critical control layer.

AMSTERDAM, NETHERLANDS, February 17, 2026 /EINPresswire.com/ -- [Orq.ai](https://www.orq.ai), the generative AI collaboration platform for building, deploying, and managing production-grade AI agents, today announced the launch of the Orq.ai Router as a standalone product. The move gives engineering teams a dedicated entry point for managing, routing, and optimizing requests across multiple large language models (LLMs) through a single gateway.



New LLMs are entering the market at a rapid pace, forcing teams to upgrade and switch models more frequently as AI systems move into production. To balance cost, performance, and reliability, most organizations now rely on multiple models, a setup that introduces operational complexity and unpredictable costs that traditional application stacks were never designed to manage. This challenge is especially pronounced for Orq.ai's early customers in regulated industries such as financial services, fintech, and healthcare, where reliability, control, and compliance are non-negotiable.

For bunq, one of Europe's leading fintech banks, these challenges had already surfaced while operating an internally built routing solution.

"We built our own LLM routing infrastructure, but maintaining it became increasingly expensive and time-consuming, while still leaving gaps in observability and performance," said Benjamin Kleppe, GenAI Lead at bunq. "We chose to work with Orq.ai to replace that internal setup with a production-ready AI Router that meets our governance, scalability, and cost-monitoring

requirements.”

The Orq.ai Router addresses these challenges at their point of origin: the routing layer.

“As soon as AI systems move beyond a single model, routing turns from plumbing into a production bottleneck. Making the router standalone lets teams regain control early, with a single line of code,” said Sohrab Hosseini, Co-founder of Orq.ai.

Previously available only as part of Orq.ai’s broader agent lifecycle platform, the Orq.ai Router can now be deployed as a standalone gateway, enabling teams to start with routing and cost control and expand into broader agent lifecycle capabilities over time.

“In Europe, AI sovereignty is no longer an abstract policy debate; it’s a direct consequence of today’s geopolitical reality,” said Sohrab Hosseini, Co-founder of Orq.ai. “Enterprises need to know where AI inference runs, who controls the infrastructure, and how quickly they can adapt as conditions change. Those decisions are enforced at the routing layer, which is why we made it available as a standalone product.”

To support these requirements, the Orq.ai Router gives teams explicit control over how requests are routed across providers and regions. Routing policies can be defined around factors such as geography, latency, cost, or custom constraints, allowing organizations to adapt infrastructure decisions without rewriting applications or re-architecting systems. Unlike other AI routers, Orq.ai can run entirely within a customer’s own infrastructure, supporting deployments across both public and private models

This approach also shapes how the Orq.ai Router is priced. Rather than applying a fixed platform fee on every request, the Orq.ai Router separates routing from usage-based services. Many AI gateways focus primarily on routing and apply percentage-based markups on top of underlying model costs, often in the range of 5–6%, which can compound quickly as usage scales. The Orq.ai Router takes a different approach: routing itself is available without a platform markup, while teams pay only for tracing and logging of processed data, based on volume.

The standalone release of the Orq.ai Router gives teams a clear entry point into production AI infrastructure, starting with routing and cost control while preserving flexibility as systems evolve. Organizations can deploy the Router independently today and expand into Orq.ai’s broader agent lifecycle capabilities over time as operational requirements grow.

The Orq.ai Router is available immediately as a standalone product, with further details available at <https://router.orq.ai/>

About Orq.ai

Founded in 2022 and headquartered in Amsterdam, Orq.ai is the enterprise control platform for the full AI agent lifecycle. The company's Agent Platform 4.0 enables organizations to

experiment, evaluate, deploy, observe, and optimize AI agents through a unified interface that prioritizes safety, collaboration, and scalability. Orq.ai supports over 300+ LLM models and offers flexible deployment options, including cloud, hybrid, and on-premises configurations to meet enterprise data sovereignty requirements. The company serves forward-thinking enterprises and AI-native startups across Europe and the United States. For more information, visit orq.ai.

Kuldeep Yadav, Head of Growth

Orq.ai

[email us here](#)

Visit us on social media:

[LinkedIn](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/893046649>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.