




Aurora Mobile's GPTBots.ai Integrates Jev — Two Layers of AI, One Enterprise Platform

September 22, 2026

SINGAPORE, Sept. 22, 2026 (GLOBE NEWSWIRE) -- Aurora Mobile Limited (NASDAQ: JG), a leading global provider of customer engagement and marketing technology services, today announced that its enterprise-grade AI agent platform, [GPTBots.ai](#), has integrated Jev — the "System One" decision model from TypeSafe AI — to build what the team calls a "two-layer AI architecture": one layer that thinks, and one layer that judges.



(Screenshot of the GPTBots.ai Console)

The integration introduces a two-layer AI architecture inside GPTBots.ai: one layer that handles reasoning, and one layer that handles decisions — each optimized for what it does best. The move comes just one week after TypeSafe AI released Jev on September 15, a purpose-built decision model that has since been integrated by Vercel, Cloudflare, LangChain, and other major developer platforms.

The Problem: Enterprise AI Pays for Thinking When It Only Needs to Judge

Every AI agent workflow is full of decisions that don't require language generation. Is this a billing question or a technical one? Does this retrieved document actually answer the user's question? Which model should handle this task? These are binary or categorical judgments — yes/no, this/that, relevant/irrelevant — yet most platforms run them through full-scale LLMs that generate paragraphs of text just to arrive at a single classification.

The result: enterprises pay for words they don't need, wait for tokens that could have been a millisecond decision, and get no reliable measure of how confident the model actually is.

The Solution: A Dedicated Decision Layer

Jev was built for this exact gap. It does not generate text. It takes unstructured state as input and returns structured, probabilistic decisions — Choice, Score, or Yes/No — in a single parallel pass, with calibrated confidence scores attached to every answer.

According to TypeSafe's published benchmarks:

- **Speed:** 70–500ms end-to-end, vs. 3–329 seconds for frontier LLMs
- **Cost:** \$0.042 per million input tokens; output is free — up to 445× cheaper than comparable LLM decision tasks
- **Reliability:** 0% structured output error rate, vs. up to 45.5% for some frontier models
- **Adoption:** Integrated by Vercel, Cloudflare, LangChain, and Langfuse within days of launch

How GPTBots.ai Builds the Two-Layer Architecture

GPTBots.ai is not adding Jev as an afterthought. The platform already operated several decision-layer mechanisms. The Jev integration deepens and unifies these capabilities under a single, purpose-built model.

The architecture divides work by nature. The **Decision Layer**, powered by Jev, handles fast, high-volume judgments — routing, filtering, classification, relevance scoring — in under 500ms at a fraction of the cost of an LLM call. The **Reasoning Layer** handles complex reasoning, text generation, and open-ended dialogue using general-purpose LLMs such as GPT and Claude. Each task runs on the engine best suited for it.

Three existing GPTBots.ai capabilities now powered by Jev:

1. Model Auto-Router — Picking the right brain before the task starts

Before any task begins, GPTBots.ai evaluates the incoming query and routes it to the best-matched model. With Jev, this routing decision becomes faster and calibrated. Jev assesses query complexity, domain, and urgency, then returns a probabilistic recommendation that the router acts on in

milliseconds — replacing an expensive LLM call with a sub-500ms judgment at a fraction of a cent.

2. Dynamic Top-K — Filtering noise out of knowledge retrieval

When RAG retrieves dozens of document chunks, not all are relevant. GPTBots.ai's Dynamic Top-K capability discards irrelevant chunks before they reach the LLM. Jev strengthens this step by scoring each chunk's semantic relevance as a calibrated probability — so only the highest-confidence knowledge reaches the model, reducing hallucination risk and token waste.

3. Intent Classification in FlowAgent and Workflow — Routing to the right business branch

GPTBots.ai's FlowAgent and Workflow modules include a Classifier that identifies user intent and routes conversations to the correct business branch. With Jev, each routing decision carries a confidence score: high-confidence cases move forward automatically, uncertain cases escalate to a stronger model or a human agent.

"We already had the building blocks — model routing, dynamic retrieval, intent classification," said Chris Lo, Founder and CEO of GPTBots.ai. "What Jev gives us is a dedicated decision engine that handles these judgments at a speed and cost point that changes the economics of the entire pipeline. Instead of paying for a full LLM call on every routing decision, we now pay a fraction of a cent per judgment — and get calibrated confidence scores to boot."

Confidence-Driven Execution: Act, Review, Escalate

The two-layer architecture introduces a configurable confidence threshold system. Every Jev decision comes with a probability score. Enterprises can set their own thresholds per workflow — deciding which decisions execute automatically, which require review, and which escalate to human agents. A billing classification might auto-execute at high confidence; a compliance decision might require near-certainty. The platform adapts to the cost of being wrong, not just the speed of being right.

What This Means for Enterprise Customers

- **Cost reduction:** Decision tasks that previously consumed full LLM tokens now cost fractions of a cent. For platforms processing millions of judgments daily, the savings compound quickly.
- **Lower latency:** Sub-500ms decision responses keep real-time conversations and automated workflows moving without waiting for LLM round-trips.
- **Calibrated confidence:** Every decision comes with a probability score — replacing "the model seems sure" with a number you can build automation rules around.
- **Zero type errors:** Jev outputs are schema-guaranteed. No JSON parsing failures, no malformed responses, no pipeline breaks from unexpected output formats.

About GPTBots.ai

[GPTBots.ai](#) is an enterprise-grade AI agent platform by Aurora Mobile (NASDAQ: JG), offering no-code/low-code AI agent development, deployment, and management for enterprise clients. Through multi-model LLM integration and intelligent workflow automation, GPTBots.ai empowers organizations to enhance customer service, optimize knowledge management, and automate business processes, driving digital innovation and transformation worldwide.

GPTBots.ai's global customers include Axios Management, GP Batteries, Ruko and many more.

For more information, please contact: marketing@gptbots.ai

About Aurora Mobile

Founded in 2011, Aurora Mobile (NASDAQ: JG) is a leading provider of customer engagement and marketing technology services. The Company is dedicated to empowering global enterprises with stable, efficient, and intelligent customer interaction solutions. Leveraging its first-mover advantage in mobile messaging, Aurora Mobile has evolved into a comprehensive platform that integrates Omnichannel Engagement, AI-Driven Marketing, Advanced AI Customer Support, and Frictionless Identity Security. Through its flagship brand EngageLab and its robust AI infrastructure GPTBots.ai, the Company helps businesses achieve seamless customer reach, automate complex marketing journeys, and optimize service efficiency with AI agents, accelerating digital transformation for clients worldwide.

For more information, please visit: <https://ir.aurora-mobile.com/>

Media contact: marketing@aurora-mobile.com

Safe Harbor Statement

This announcement contains forward-looking statements. These statements are made under the "safe harbor" provisions of the U.S. Private Securities Litigation Reform Act of 1995. These forward-looking statements can be identified by terminology such as "will," "expects," "anticipates," "future," "intends," "plans," "believes," "estimates," "confident" and similar statements. Among other things, the Business Outlook and quotations from management in this announcement, as well as Aurora Mobile's strategic and operational plans, contain forward-looking statements. Aurora Mobile may also make written or oral forward-looking statements in its reports to the U.S. Securities and Exchange Commission, in its annual report to shareholders, in press releases and other written materials and in oral statements made by its officers, directors or employees to third parties. Statements that are not historical facts, including but not limited to statements about Aurora Mobile's beliefs and expectations, are forward-looking statements. Forward-looking statements involve inherent risks and uncertainties. A number of factors could cause actual results to differ materially from those contained in any forward-looking statement, including but not limited to the following: Aurora Mobile's strategies; Aurora Mobile's future business development, financial condition and results of operations; Aurora Mobile's ability to attract and retain customers; its ability to develop and effectively market data solutions, and penetrate the existing market for developer services; its ability to transition to the new advertising-driven SAAS business model; its ability to maintain or enhance its brand; the competition with current or future competitors; its ability to continue to gain access to mobile data in the future; the laws and regulations relating to data privacy and protection; general economic and business conditions globally and in

China and assumptions underlying or related to any of the foregoing. Further information regarding these and other risks is included in the Company's filings with the Securities and Exchange Commission. All information provided in this press release and in the attachments is as of the date of the press release, and Aurora Mobile undertakes no duty to update such information, except as required under applicable law.

For more information, please contact:

Aurora Mobile Limited

E-mail: ir@aurora-mobile.com

Christensen Advisory

Ms. Xiaoyan Su

E-mail: Xiaoyan.Su@christensencomms.com

A photo accompanying this announcement is available at <https://www.globenewswire.com/NewsRoom/AttachmentNg/8e3b368a-2f58-45c8-b87b-070b97cf91fd>