new relic

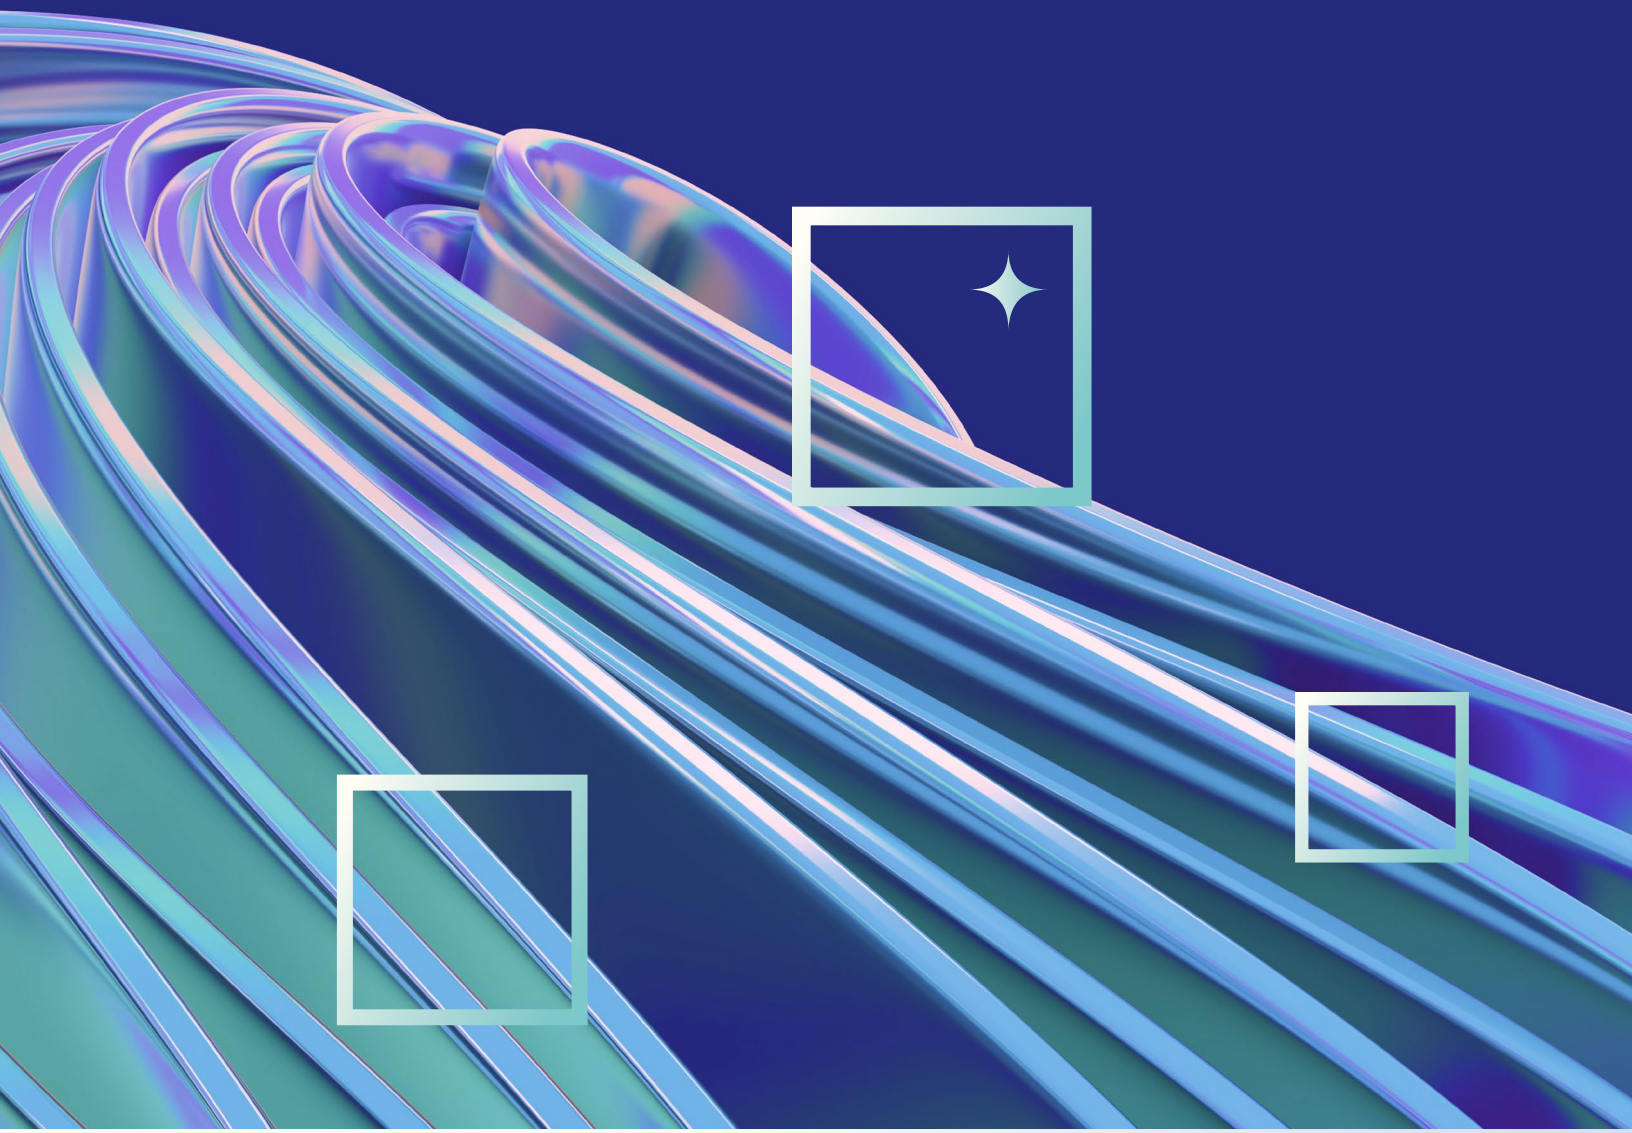

AI Unwrapped

2025 Impact Report

Gain insight from real-world AI implementations today.

Report / June 2025

Introduction

Shipping code is only half the battle. Without robust observability (the comprehensive logs, metrics, and traces that reveal precisely how systems are behaving in real-time), even the most brilliant applications become opaque, unreliable black boxes at scale. As artificial intelligence moves from speculative promise to foundational technology, we're seeing the same critical lesson emerge.

AI adoption is accelerating at an unprecedented pace, and the patterns in how organizations use it are evolving just as fast. For the first time, New Relic is sharing aggregated usage data that offers a front-row view into how developer behavior is continuously reshaping the AI landscape, revealing why AI monitoring has moved from optional to absolutely essential.

New Relic data shows that ChatGPT dominates usage across prompts observed as customers and developers rapidly shift to the latest versions like GPT-4o, even when lower-cost alternatives exist. At the same time, there's a sharp rise in developers creating custom models for use with GPT, a signal that more teams may start choosing niche tools and tailoring AI to their specific needs in the future.

These trends point to a fundamental inflection point: AI is moving from experimentation to operationalization, rapidly. The demands of this dynamic environment require more than basic, task-specific monitoring tools, underscoring the rising adoption of New Relic's AI Monitoring solution. And as these trends continue, it's becoming clear that AI monitoring is now a core enabler of technical agility and operational resilience for developers, infrastructure teams, and CTOs alike.

ChatGPT dominates usage
across prompts observed
as customer adoption.

ChatGPT dominates

Developers are driving exceptionally strong growth in AI model experimentation and specialization. New Relic saw a 92% increase in the number of unique large language models used across AI apps in the most recent quarter.

More specifically, developers are embracing the largest LLMs, led by OpenAI's ChatGPT, which accounted for **more than 86%** of all LLM tokens processed by New Relic customers, with Meta's Llama coming in second position. That lion's share strongly suggests that developers are prioritizing ease of use, reliability, and top-tier performance with minimal setup or tuning. As of now, it's fair to conclude that organizations want to deploy rapidly using enterprise-ready models rather than invest time in building or fine-tuning alternatives.

>86%

of all LLM tokens processed by New Relic customers used ChatGPT

Even with ChatGPT's dominant model usage, observability is not optional—teams still need to track usage, benchmark performance, and detect anomalies to ensure cost-efficiency, optimal performance, and reliability.

In general, many organizations are still in earlier stages of AI deployment, so it makes sense that many rely on generalist models. However, the data shows that developers are also testing a wide variety of models, including domain-specific and task-specific, although at a smaller scale. As teams and organizations become more sophisticated, the need for advanced AI monitoring to maximize performance, manage costs, ensure governance compliance, and navigate multi-model environments is likely to increase substantially in the following years.

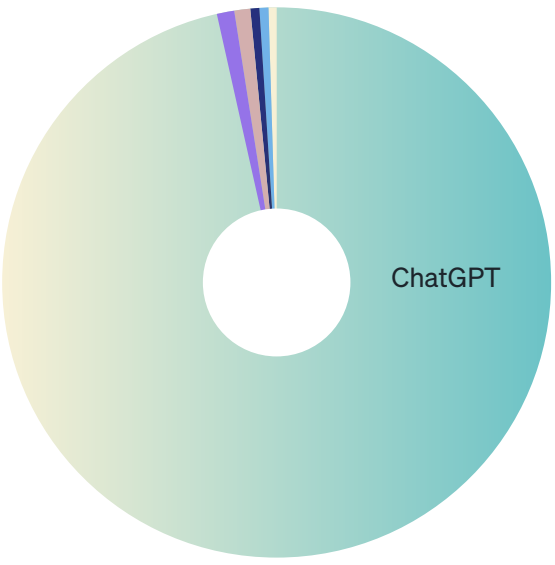

Most popular models by observed LLM token consumption

- ChatGPT
- Meta Llama
- Ada
- TheDrummer/Unslop
- Anthrocite Magnum
- Anthropic

Developers shift rapidly

As the pace of AI innovation accelerates, developer behavior is proving fluid: model version switching happens rapidly. ChatGPT-4o has been dominating more recently followed by ChatGPT-4o mini. Given that ChatGPT-4o mini is cheaper than ChatGPT-4o, it possibly suggests that developers value cutting-edge performance and features (e.g. better multimodal capabilities, lower latency) more than savings, at least right now. It will be worth watching to see if users will start optimizing cost over time.

In some cases, pivoting from ChatGPT version-to-version is occurring seemingly overnight. For example, New Relic users have been rapidly shifting from ChatGPT-3.5 Turbo to ChatGPT-4.1 mini since it was announced in April. It seems clear that enterprises are closely monitoring the latest innovations from OpenAI to ensure they stay competitive.

In this fast-moving environment, organizations need AI monitoring that keeps pace with the ever-evolving AI stack, especially as OpenAI releases incremental versions and updates to each model over time. Developer teams will need to evaluate new models for greater visibility without losing control over cost, reliability, or compliance.

Pivoting from ChatGPT version-to-version is occurring seemingly overnight.

Organizations need AI monitoring that keeps pace with the ever-evolving AI stack.

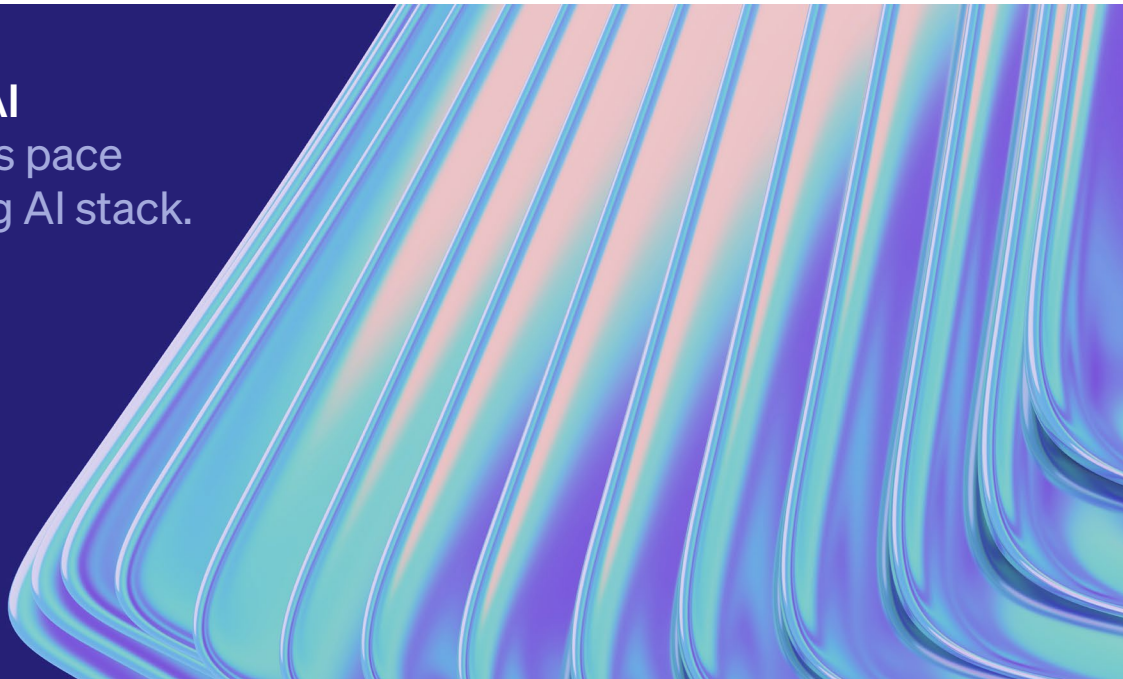

AI monitoring surges

Since its launch last year, enterprises have been adopting New Relic AI Monitoring to ensure AI model reliability, accuracy, compliance, and cost efficiency. New Relic saw a steady 30% growth in usage quarter-over-quarter in the last 12 months.

New Relic's *2024 Observability Forecast* revealed that the rapid adoption of AI is the primary driver of organizations deploying AI monitoring solutions. As AI transitions from experimental to operational, companies face new monitoring challenges. This evolution requires comprehensive AI observability capabilities that extend beyond large language models. Observability systems need to encompass the complete ecosystem of AI-related technologies, seamlessly integrated within existing operational infrastructure.

Meanwhile, some enterprises use AI models without a robust AI monitoring solution, placing them at substantial risk, including reliability, response accuracy, security, and cost spikes. AI without proper monitoring is like deploying software without logs—organizations are blind to performance issues, vulnerable to failure, and not fully prepared for AI governance. That could slow the pace of AI innovation and can undermine confidence in scaling AI. Organizations need a unified AI monitoring solution that is easy to set up, configure over time, and provides an intuitive experience for any user—from devops to executives.

> **Our new AI tool connects to New Relic as well, where we monitor its uptime, performance, and response times.** It's keeping costs down by allowing our developers to focus on the user-facing product features that matter."
>
> **Jared Chung**
> Executive Director at CareerVillage.org

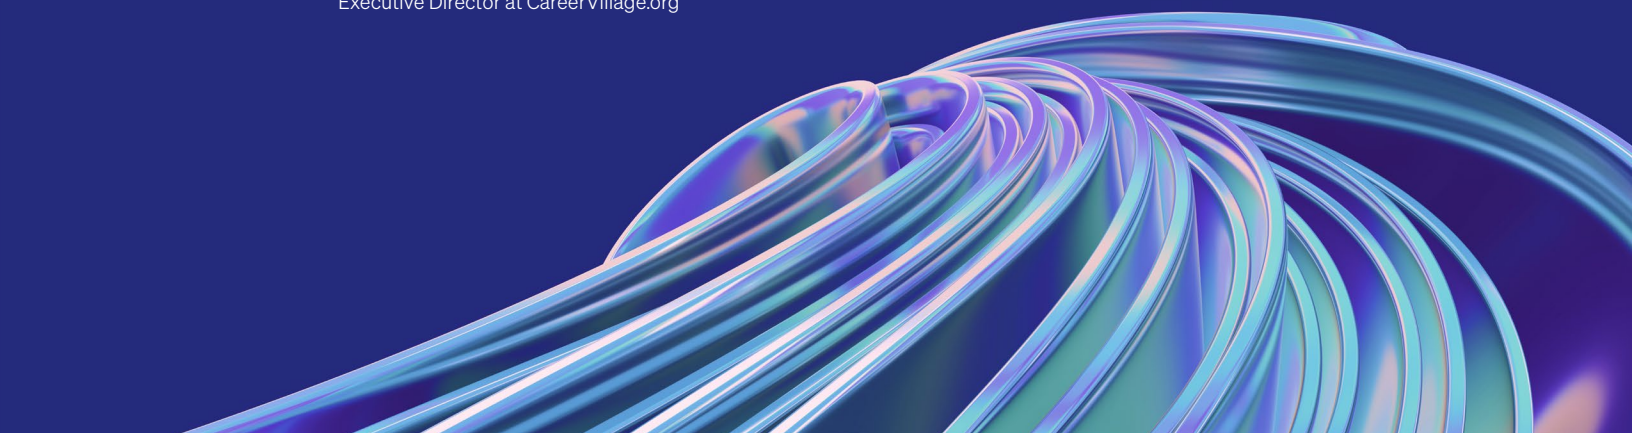

Python dominates

Customers use New Relic to observe AI applications. From those applications, Python continues to dominate in AI applications followed by Node.js in terms of both scale of requests as well as customer adoption. Python customer adoption grew nearly 45% since last quarter.

As AI becomes deeply embedded in enterprise operations, the range of programming languages for AI applications will necessarily expand. This diversification is a natural consequence of enterprises needing to integrate AI with legacy systems (often built in languages like Java or C#), address specific performance and scalability demands (driving adoption of C++, Go, or Rust), and deploy AI in varied environments, from cloud to edge devices.

Java usage growing

In fact, New Relic recently saw Java usage growing rapidly, 34% since last quarter. New Relic believes that this trend is worth watching and could be a sign of more production-grade Java-based LLM applications to come from large enterprises.

It is no surprise to see Python leading in adoption and usage, given that it has the most momentum, support, and tooling. However, New Relic believes tooling of other languages will continue to mature, leading to greater proliferation over time, which explains why New Relic supports the largest number of languages, providing comprehensive observability for AI applications built with Python, including those using LLMs.

34%

growth of Java usage since last quarter

New Relic supports the largest number of languages, providing comprehensive observability for AI applications built with Python, including those using LLMs.

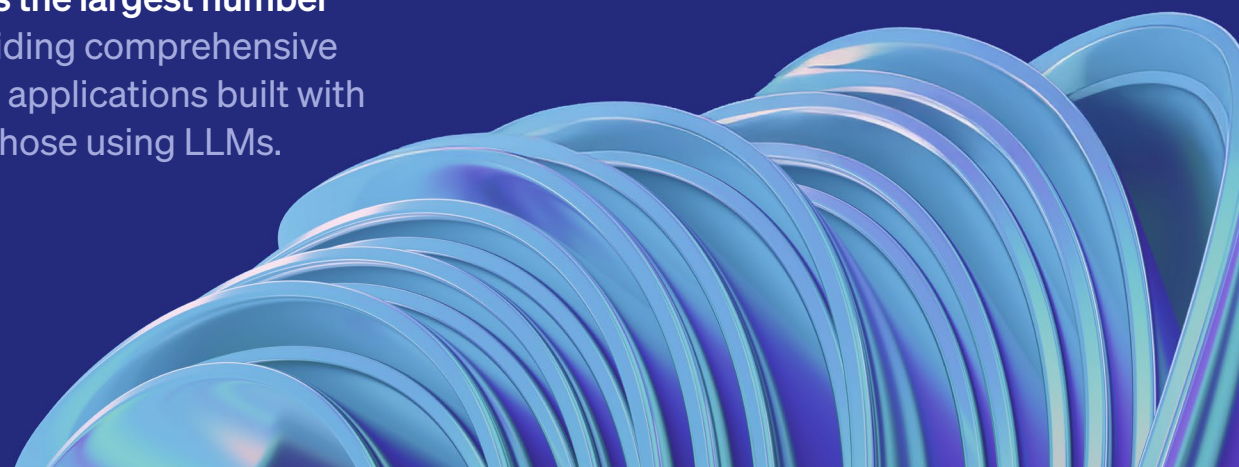

Methodology

This report is based on aggregated and de-identified usage data gathered from 85,000 active New Relic customers collected between April 30, 2024 and April 30, 2025.

About New Relic

The New Relic Intelligent Observability Platform helps businesses eliminate interruptions in digital experiences. New Relic is the only AI-strengthened platform to unify and pair telemetry data to provide clarity over your entire digital estate. We move your problem solving past proactive to predictive by processing the right data at the right time to maximize value and control costs. That's why businesses around the world—including Adidas Runtastic, Domino's, GoTo Group, Ryanair, Topgolf, and William Hill—run on New Relic to drive innovation, improve reliability, and deliver exceptional customer experiences to fuel growth.

Explore New Relic AI Monitoring

new relic.