Futurum
RESEARCH

Salesforce's Real Moat Behind Koa Is Its CRM Training Data

NVIDIA's Open-Weight Nemotron Is Available to Every Rival but Replicating 27 Years of Encoded CRM Deployment History Is Not

Analyst(s): Keith Kirkpatrick
Publication Date: September 22, 2026
Document #: AINKK202609

What You Need to Know

- Koa's Hardest-to-Copy Asset Is Its CRM Training Data, Not the Reasoning Model: NVIDIA's Nemotron 3 Super is open-weight, so any rival foundation-model lab or platform vendor can post-train it the same way Salesforce did. Reproducing 27 years of encoded CRM deployment patterns across more than 14 industries is a different problem, one only Salesforce has the deployment history to solve today.
- The Interface Argument Behind Koa Is Not New: Salesforce's defensive moat is evolving from the application to the capability layer beneath it.
- Koa's Trust Boundary Is a Second, Less-Proven Pillar: Keeping inference inside Salesforce's own infrastructure via Guardian and Data 360 is an architectural choice, but the durability of that boundary is against a rival building the same capability, which is untested.
- The Benchmark Claims Are Unverified: Salesforce's three-times-fewer-errors figure comes entirely from its own CRM Bench.

The Futurum View

Salesforce announced its Koa CRM reasoning model at Dreamforce 2026 as part of a broader Headless 360 and Agentforce strategy. Based on comments from Salesforce President and Chief Platform Officer Rohan Kumar and on public reporting on Koa's technical design and early reception through September 2026, the move can be viewed as a way for Salesforce to leverage operational depth as a moat against the "build it yourself" narrative.

Indeed, ETR's September 2026 AI Tools Pulse data shows Claude experiencing broad enterprise deployment, jumping from 38% in June to 59%, while 78% of Claude users and 66% of ChatGPT users also increased consumption over the past month. Notably, agentic workflows are becoming a key driver of AI usage, with 41% of organizations saying at least 20% of their AI usage is agentic and 25% now running complex, multi-step agentic workflows.

Figure 1: AI Model Deployment Among Enterprise Buyers

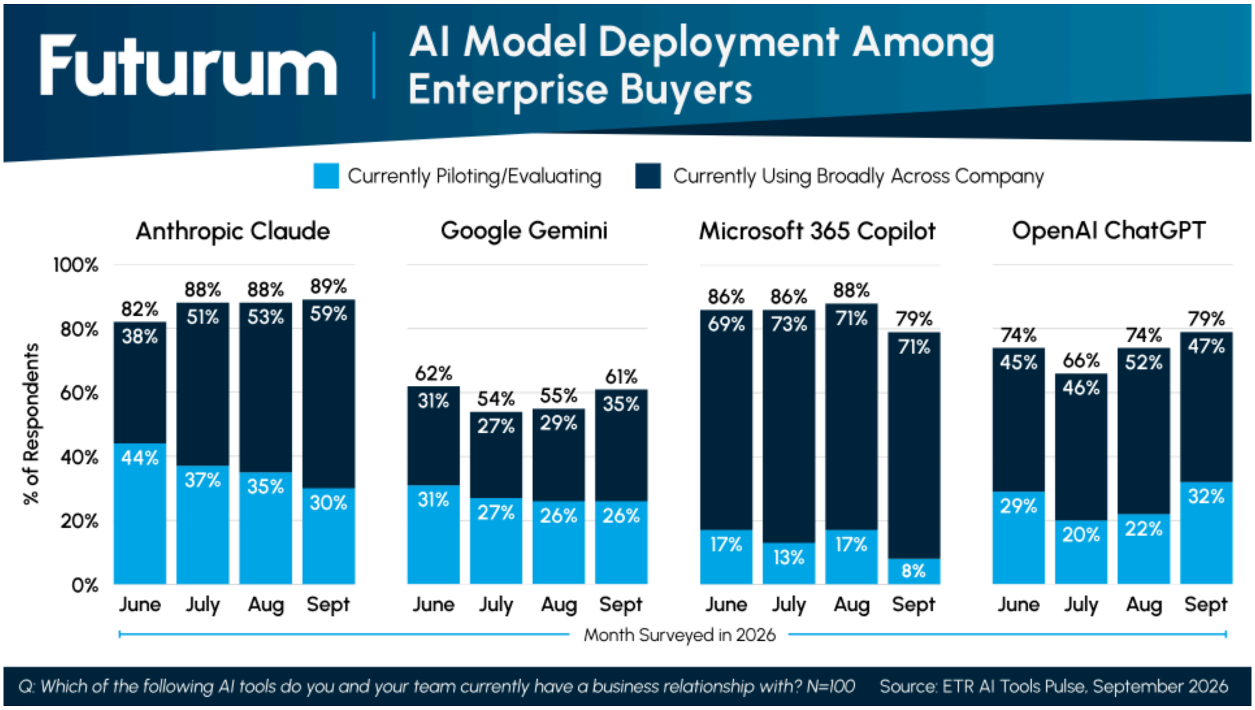

Source: ETR AI Tools Pulse, September 2026

Koa's CRM Training Data Is Harder to Copy Than the Model

NVIDIA's Nemotron 3 Super is an open-weight model. Salesforce post-trained it into Koa using supervised fine-tuning and reinforcement learning, and nothing in that process is proprietary to Salesforce. Any foundation-model lab, systems integrator, or rival platform vendor, Microsoft, SAP, ServiceNow, or Workday, among them, can download the same base model and run the same post-training recipe. If the model architecture were the moat, it would not survive a product cycle.

What is harder to reach is the training data. Salesforce says it generated synthetic scenarios simulating real personas, action sequences, and tool calls across more than 14 industries, drawing on 27 years of CRM deployment patterns, though none of that was proprietary customer data. That data set encodes a specific kind of knowledge: what a deal looks like as it moves through a pipeline, what a service case escalation looks like, and how both vary by industry, learned from three decades of being the system where those events happened. A rival with a shorter or thinner deployment history has more limited pattern depth from which to synthesize data, and new competitors entering the market today have none.

Microsoft, SAP, ServiceNow, and Workday all have their own multi-decade deployment histories in adjacent systems of record, and any of them could synthesize a comparable training set from their own history rather than Salesforce's. However, Salesforce's specific combination of tenure and CRM-specific breadth stays ahead of what most of these rivals can synthesize from their own systems.

The New Moat: The Metadata-Aware Capabilities Layer

The broader case, that as AI agents reach enterprise data through APIs instead of screens, defensible value moves from the interface to whatever sits beneath it, predates Koa. Indeed, the next competitive moat in enterprise software is not the application itself, but the governed, metadata-aware capability layer sitting beneath it. Koa is a concrete answer to this question.

Reasoning and Data Synthesis Remain in Salesforce

Kumar drew a second line in the briefing with analysts at Dreamforce, noting that "a lot of the value Salesforce gets is the context it understands about the enterprise," meaning the customer segmentation, identity resolution, and operational metadata assembled through Data 360's zero-copy federation. Salesforce Guardian's agent-identity and data-classification controls, paired with that zero-copy link, keep the actual reasoning and data synthesis running inside Salesforce's infrastructure even when the front end is Slack or ChatGPT, so Salesforce remains the system doing the work even when it is not the system the user sees.

Kumar extended the argument when asked whether model providers could simply pull enterprise context out of Salesforce and build it themselves: context Salesforce updates within seconds of a customer action is a different asset than context reconstructed later from a data export. This extends the moat, particularly in cases where real-time context matters in decisioning. The advantage only pays off if Salesforce's own harness keeps routing to Koa over Claude, ChatGPT, and Gemini, and no party outside Salesforce can yet confirm it does.

The Benchmark Claims Are Unverified

Koa's three-times-fewer-errors claim comes entirely from Salesforce's internal CRM Bench, with no independent auditor's results published. Until the company publishes third-party verified data around accuracy, speed, and efficiency, prospects and customers will view Koa as an untested model.

Salesforce is not betting everything on Koa. Kumar said in the briefing that the company is pursuing model choice and its own reasoning model "both" at once, and Headless 360's architecture is explicitly open to Claude, ChatGPT, and Gemini running inside the same harness. Furthermore, Salesforce announced Claudeforce, a continued partnership with Anthropic, which essentially provides more optionality for customers. If a customer's harness can route a task to whichever model is cheapest that week, Koa has to keep winning that decision on its own merits rather than by default.

Two Moats, Not One

Futurum's June 2026 assessment of Salesforce's "hidden moat" assessed the company on a different basis entirely: operational depth, meaning data governance, compliance frameworks, identity resolution, and the accumulated integrations connecting one enterprise system to another. That moat and the one this analysis describes are not competing explanations.

Operational depth protects Salesforce against a customer trying to build a replacement in-house; the training-data pipeline behind Koa protects Salesforce against a rival vendor trying to build a competing reasoning model. A buyer or a competitor who clears one barrier still has to clear the other, providing Salesforce with an advantage in the market, at least for now.

You can read the full press release on Salesforce's Koa at the company's website.

What to Watch

The following factors, tied back to the takeaways above, will determine whether Koa's moat argument holds over the next several quarters.

- Whether a Rival Can Synthesize Comparable Training Data: Microsoft, SAP, ServiceNow, and Workday each have multi-decade deployment histories of their own; whether any of them can turn that history into a training set as CRM-specific as Koa's is the real test of this analysis's central claim.
- Independent Verification of the CRM Bench Results: Until a party outside Salesforce audits Koa's error-rate and cost claims, buyers should treat the three-times-fewer-errors figure as a vendor benchmark, not a proven outcome.
- Whether Salesforce's Own Harness Keeps Routing to Koa: Headless 360 is explicitly open to Claude, ChatGPT, and Gemini, so the real test is whether Salesforce's own AI gateway keeps choosing Koa once customers can compare it against rival models on cost and accuracy, case by case.
- Agentforce's Broader ROI Track Record: Reports of implementation costs and timelines missing marketed expectations for Agentforce raise the bar Koa has to clear before enterprises treat Salesforce's reasoning-model strategy as proven rather than promised.

Other Insights from Futurum

Salesforce Bets the Platform on Headless 360

The Hidden Moat: Why Operational Depth Defeats the 'Build It Yourself' Narrative

AIforce Turns Salesforce Into an Everywhere Intelligence Layer

About Us

About the Authors

Keith Kirkpatrick is VP, Research, Enterprise Software & Digital Workflows at Futurum. His coverage area encompasses enterprise applications, platforms, and tools, as well as the technologies and tools that underpin them, including artificial intelligence, automation, and integration. He has more than 25 years of experience in research, marketing, and consulting-based fields, and is frequently sought out by the vendor community and trade media for his insights and commentary. He can be reached at kkirkpatrick@futurumgroup.com.

About The Futurum Group

Every day, The Futurum Group's analysts, researchers, and advisors help business leaders worldwide anticipate tectonic shifts in their industries and leverage disruptive innovation. Unlike traditional analysts, The Futurum Group works not only in analysis and research but also takes that insight and knowledge even further, engaging all the way through the go-to-market process.

Futurum Research provides in-depth research and insights on global technology markets using advisory services, custom research reports, strategic consulting engagements, digital events, go-to-market planning, and message testing. It also creates, distributes, and amplifies rich media content that all stakeholders read, watch, and listen to.

See more details on The Futurum Group at futurumgroup.com.

Copyright & Use License

Copyright Notice

License Notice

Limitation of Liability Notice