



NUCLEUS
RESEARCH

Salesforce unveils its AI strategy with Data 360 and Koa

ANALYST

Alexander Wurm

The Bottom Line

In 2026, enterprise software is turning headless and ambient, and the agents taking on real customer-facing work are only as good as the context they run on. A general model bolted onto that data misses the records it needs, and the agents built on it inherit the gap. At Dreamforce 2026, Salesforce announced Data 360 as the data layer of its agent-first Salesforce platform, with Koa, its proprietary CRM reasoning model built with NVIDIA, tuned to run on it. The data process and metadata are the key to the unlock, letting an agent act reliably on the system of record. Nucleus expects adopters to run agents on work they previously kept human and to cut the cost of moving data out of the CRM. The owned, governed data layer gives adopters the edge in agentic CRM, and Nucleus expects that layer, not the model, to become the reference architecture for agentic CRM in 2026 and forward.

Overview

Enterprise agents have changed what customers ask of the CRM. An agent runs on the context behind it, and that context has to be curated, governed, and refreshed at the pace the use case demands. Customers are no longer evaluating the CRM as a system of record that humans query. They are evaluating it as the layer that determines whether their agents work, and they want that layer to serve the agents alongside the analysts and application developers. A general-purpose model bolted onto that data has to pull records out of the CRM to reason across them, so every action the agent takes starts with moving data.

Data 360 is the data layer of the Salesforce Platform, bringing consistency to agentic and human users.

Salesforce addressed this with Data 360, the data layer under its new agent-first Salesforce platform. Data 360 brings the data process and the metadata that describe the customer relationship into one governed layer, so an agent works from the system of record rather than on a copy moved elsewhere. On top of that layer, Salesforce announced Koa, the company's first CRM reasoning model built with NVIDIA by post-training Nemotron. Koa is in pilot now.

Salesforce Dreamforce 2026

Headless Toolkit and Data 360

Salesforce framed the shift around the platform, not the model. The company describes enterprise software as becoming more headless and more ambient, with autonomous agents operating in the background. It positions itself as the system of record and system of engagement, the trusted foundation for customer-facing work in the age of AI. The question the company poses is which platforms agents will trust when they need to get real work done, and Data 360, the data layer of the Headless Toolkit, is built to be that platform. It gathers the customer data and the metadata that describe it into a single governed layer, so agents act on the system of record. Because the context is owned, governed, and refreshed in place, an agent can act on real CRM work without leaving the CRM to reach an external model. The business keeps the controls its compliance and data teams require.

Koa runs inside Salesforce's trust boundary, allowing organizations to better trust the work that is done.

Koa, the model tuned to run on Data 360

Koa is the second-order piece, the model that puts the governed Data 360 layer to work. The model is purpose-built to help Agentforce agents reason through complex, multistep CRM workflows and use the

right tools. It runs inside Salesforce's trust boundary, so the reasoning stays on the governed data layer, not in an external model. The vendor reports that Koa matches or exceeds leading model performance on CRM actions with 3x fewer errors on Salesforce's own CRM benchmark. Independent press notes that the benchmark is self-defined. Koa is in pilot now with select customers, and general availability is expected in U.S. regions in winter 2026.

Koa is a post-training of NVIDIA Nemotron 3 Super on a synthetic dataset modeled on 27 years of CRM intelligence.

A national virtual accounting and tax services firm, an established Agentforce user, is among the Koa pilot customers. Salesforce reports that the firm resolves 50 percent of its support inquiries with digital labor. The firm also resolved more than 1,000 client engagements in the first 24 hours. Both figures come from its 2025 Agentforce deployment, which predates Koa. The firm runs an established agent workload on a CRM-native model on the governed data layer.

Expected Benefits

Data 360 and Koa are expected to deliver gains across agent reliability, the cost of running agents on the system of record, and the controls that let agents take on higher-value work.

A data layer agents can trust

Curated, governed Data 360 context gives agents the records they need to act, so they run on real CRM work with fewer corrections and more throughput. Agents take on work the organization previously kept human because they reason on governed context, not on a copy they have to trust blindly. The organization can expand the share of customer-facing work it hands to agents without adding a human review step.

Customers can expect up to 3x fewer errors CRM tasks.

Lower cost to run agents on the system of record

Because Koa runs on the owned data layer, an agent reaches the CRM without moving data out of it. The organization avoids the infrastructure and re-wiring that a bolted-on external model requires, and it keeps the data where it is governed. The business redirects the staff time it no longer spends moving records between systems to new agent use cases.

Efficient action in parallel to headless

Koa is tuned to Data 360, so it works in parallel to the headless platform and takes on work a frontier model would otherwise run. The organization can point agents at higher-value CRM processes under

the controls its compliance and data teams set, because the reasoning stays inside the trust boundary on governed context.

Looking Ahead

Nucleus expects the owned, governed data layer, with a CRM-native model tuned to run on it, to set the direction for agentic CRM in 2026 and forward. As enterprise software turns headless and ambient and agents take on real customer-facing work, organizations gain the advantage on the platform that curates and governs its customer data at the pace the use case demands. The model is the second-order piece, and it makes the agent efficient on the data the organization owns, so the platform holds the workloads a general model would otherwise take. For buyers, organizations that already run Agentforce on the Salesforce platform should evaluate Data 360 now, because the data layer is the foundation the model and the agents build on. Buyers that are not yet on the platform should wait for general availability, verify the governed data layer, and review the first measured outcomes before committing.

A national virtual accounting firm resolves 50 percent of its support inquiries with digital labor, using Koa on Data 360.