




AI AGENTS

When Your CRM Starts Building Itself: The CIO's Role in the Agentic SaaS Era

By Sakshit Teotia, [Rajeev Singh](#), Matt Emara, Karan Gupta, and Srikanth Srinivasan

ARTICLE SEPTEMBER 15, 2026 12 MIN READ

CRM development is becoming a lot more democratic, and CIOs need to quickly adapt to this change. Doing so may seem challenging, but it represents a genuine opportunity.

With agentic AI, business users can build and automate more of their own workflows directly in the CRM, without deep technical expertise or day-to-day reliance on IT. Every major enterprise

SaaS platform now ships with capabilities that put more control in the hands of business users. The question is when they will start to use them.

CIOs have seen citizen development before, from Access databases and SharePoint to low-code apps, and they learned to live with the sprawl. But AI agents change the equation because they don't just present information. They make decisions, act on them, and run up costs with every action they take.

For CRM—the system of record for the customer and the engine for revenue-critical workflows—the risks are particularly acute. They include fragmentation of the customer record, ungoverned agent decisions, vendor lock-in, and runaway costs as SaaS vendors move from seat-based to consumption-based pricing.

Already, 70% of technology leaders say that their business teams are deploying technology faster than IT can track. The traditional model had clear advantages for CIOs: IT had full visibility into and control of every change. But it came with a real cost: bottlenecks were common, as IT's limited capacity to build what the business needed constrained innovation.

Herein lies the opportunity for CIOs to capture the best of both worlds without inheriting the downside. The goal is an environment in which development moves closer to the people who understand the business best, yet remains safe, consistent, and economically viable.

So what should CIOs do? Stop owning delivery on the platform and start owning the foundations: the data architecture, the agent governance model, the orchestration layer, and the cost model underpinning the commercial strategy.

CIOs that get this right can turn an operational shift that initially raises flags into a force multiplier for the entire organization. They can expand their own influence and impact while fostering innovation and growth.

Why CRM Raises the Stakes

The pressure to open up development will only increase as more business users rely on AI agents. More than half of sellers have used agents, according to Salesforce's 2026 *State of Sales* report, and nearly nine in ten expect to do so by 2027. Among sales leaders whose organizations already use them, 94% say that agents are critical to meeting business demands.

As reliance on agents grows, so will demand for more of them. Within a couple of years, a large enterprise can reasonably expect to be running hundreds of agents, and in some cases thousands. Few technology leaders are prepared to manage at that scale.

We are starting to see how this will play out, as enterprises discover dozens of ungoverned AI agents running across sales and service functions, making decisions about pricing, lead prioritization, and customer outreach before any framework exists to track them. The gap between what agents are doing and what IT knows about them is a reality for many organizations today.

That gap matters across the technology landscape. But nowhere are the stakes higher than in CRM, which sits at the intersection of three things that enterprises can least afford to get wrong: the customer record, revenue-critical workflows, and customer trust. The CRM holds the data defining who the customers are and what the enterprise knows about them. It powers processes from lead management and pricing to renewals and service. And it shapes interactions that customers experience directly, from outreach and offers to service resolution.

This makes CRM particularly well suited to distributed building: sales, service, and marketing teams can create agents around the customer processes they know best. But it also raises the stakes when things go wrong. When an agent misfires on CRM, the impact can hit data, revenue, and customer trust at the same time, often before IT even knows that the agent exists.

In practice, the risks tend to surface in four ways:

- **The customer record fragments.** When teams build agents independently, they may create different views of the customer. Sales may define an “active customer” one way; service, another. Over time, those differences can undermine the shared customer record that CRM is supposed to provide. The effects ripple into billing, forecasting, analytics, and compliance, as different parts of the business begin operating from different versions of the truth.
- **Agents make revenue decisions without adequate controls.** Agents can prioritize leads, recommend discounts, and generate offers. They can make decisions that once required human judgment, and they can execute them at machine speed. Without clear limits on what agents can decide and do autonomously, a bad decision can quickly become a bad decision at scale. A renewal agent that applies an unapproved discount, for example, could affect hundreds of accounts before anyone spots the pattern.
- **Vendor lock-in gets deeper.** As teams build more agents within a SaaS platform, more of the enterprise’s workflows can become tied to that vendor’s technology. Over time, it becomes harder and more expensive to switch providers, or even to route work to a different AI model or agent platform. The risk is not just technical lock-in. As the cost of switching rises, the enterprise’s negotiating leverage falls, too.
- **Consumption costs run away.** Agentic SaaS is shifting the economics of software from predictable seat-based licenses toward models that charge for usage, actions, or outcomes. Most major SaaS vendors have already layered an AI consumption meter on top of existing seat pricing (while also bundling AI capabilities into new SKUs that carry a step-up at renewal). This makes costs harder to forecast: every new agent or workflow can generate

incremental expense as it runs. A single high-volume agent, running automated outreach or resolving customer cases, can incur costs at a pace that bears little resemblance to the headcount-based budgeting models that CIOs have historically used. Multiply that difference across teams building independently, and consumption can climb quickly, without the visibility or natural budget ceilings that are typical of traditional licenses. The result can be a very unwelcome surprise when the bill arrives.

For CIOs, the challenge now is to make distributed building both safe and scalable.

Own the Foundations, Not the Delivery

The answer is not to reclaim the builder role. If CIOs attempt to centralize agent development, business teams will either find themselves waiting in development queues again or figure out ways around the prescribed process. Nor should CIOs let functions build but then stipulate that IT must approve every agent one by one. That approach trades a development bottleneck for a governance bottleneck.

The stakes are becoming clear. Many agentic programs are being abandoned before they reach production, and in many cases for the same reasons: costs outrun the business case, teams lack a baseline to compare value against, and IT tacks risk controls onto the existing system instead of building them into the design. These are management failures, not model failures. Avoiding them means balancing control with velocity to make the safe path the fastest path. Done well, distributed building can dramatically expand the enterprise's capacity to innovate because every user becomes a potential builder.

Business functions should own the use cases they pursue, the agents they build, and the outcomes they produce. The CIO's job is to create an environment that lets users move quickly while ensuring enterprise-wide consistency and governance. This requires CIOs to own four foundations: the data architecture, the agent governance model, an enterprise-controlled orchestration layer, and the cost model underpinning the commercial strategy.

Data Architecture: Give Every Agent the Same Truth

The CIO needs to own a single, governed data architecture that every agent uses, regardless of which function built the agent. For CRM, the results will be one customer identity, common definitions for key business concepts, and a consistent view of each customer's engagements. Sales and service may build very different agents, but those agents should agree on who the customer is and what actions the agent can perform.

This means making customer data agent-ready rather than report-ready. Report-ready data answers "What happened?" Agent-ready data must answer "What am I allowed to do, and for whom, right now?" The CIO should own that common context and the rules governing how agents access and update it, rather than allowing each function to create its own for its agents. In addition, all functions should automate logging of agent interactions with the customer record, so they can trace and audit decisions and changes afterward.

A simple test: Could two functions independently build an agent and get the same answer about the same customer? If not, the foundation is not ready.

Agent Governance Model: Make Agents Earn Their Autonomy

Traditional access controls answer a familiar question: "What is this user allowed to see and do?" Agent governance adds another: "How much can this agent do on its own?"

Autonomy should be earned in graduated, revocable steps. An agent might begin by recommending an action for a human to approve, then progress to taking certain actions independently, and eventually operate autonomously within a clearly described scope. The CIO should define this progression across functions, including the boundaries and the evidence necessary to move from one level to the next. In addition, every agent should have a named human owner who is accountable for it. If an agent's performance deteriorates or if the agent breaches its boundaries, the organization should reduce its autonomy accordingly.

Agents require governance and oversight throughout their life cycle, not just at launch. Every agent should be registered, monitored, reviewed as it changes, and retired when it no longer delivers sufficient value. Audit trails, decision logs, cost and performance tracking, alerts for unexpected behavior, and a tested kill switch should be standard parts of the model.

A simple test: For every agent running across the organization, is it clear who owns it, what it is authorized to do, what it has actually done, and how to stop it within minutes?

Enterprise-Controlled Orchestration Layer: Preserve the Freedom to Choose

As agents proliferate, enterprises will rarely rely on a single provider. As a result, agents built across SaaS platforms and elsewhere will need to work together on processes that span multiple systems.

The orchestration layer serves as a strategic control point that determines which agent or platform handles a task and how work moves between agents or platforms. The CIO should decide whether the CRM vendor will control this layer or simply participate in it, rather than letting the answer emerge by default. In either case, the enterprise should retain control of the policies governing how agents operate across platforms.

The goal is to preserve choice. The enterprise should be able to route work to the model, agent, or provider that best suits the task, and it should be in a position to change that choice as technology, performance, or economics evolve. That flexibility preserves commercial leverage, too: the easier it is to move work elsewhere, the less the enterprise must depend on any single vendor.

A simple test: If the CRM provider changed its pricing model tomorrow, how much of the agent estate could move to an alternative platform within six months?

Cost Model: Make the Economics Visible

CIOs should take the lead in defining the cost model for agentic software: how agent costs are measured, how they connect to business outcomes, and how those economics change as usage grows. As SaaS pricing shifts toward consumption, cost visibility becomes increasingly important. A listed cost per seat tells you little when agent activity drives the bill. More useful measures tie costs to individual business results, such as a resolved case or a qualified lead.

The function that builds an agent should be accountable for what it costs and for the value it delivers. Charging consumption back to the function brings that accountability home. The CIO should provide a clear view of the relevant economics by tracking agent costs, forecasting how they are likely to grow, and measuring cost per business outcome.

The same information will strengthen the CIO's hand with vendors. Rather than negotiating primarily over seats and discounts, CIOs can challenge pricing based on the economics of the outcomes that the technology delivers.

CIOs should zero in on the fine print. Does the contract make clear what generates a charge? Does it keep rates predictable as consumption changes? Does it preserve the organization's ability to move workloads elsewhere if the economics no longer make sense?

A simple test: Is it possible to tell the CFO today what each AI agent costs per business outcome? If the CIO can identify the spending but not what it delivered, the economics are not yet managed.

Where to Start

CIOs do not need to perfect all four foundations before they start. Instead, they can begin with a few high-value use cases, identify which foundations the cases expose as weakest, and address those gaps first. Then they can further strengthen the foundations as adoption grows.

In practice, that means three moves.

- **Inventory what exists.** Most enterprises have more agents running than they are aware of. Inventory what has been built on the CRM stack, function by function, and give every agent an identity and a named owner. CIOs cannot govern what they cannot see.
- **Prioritize where the consequences are greatest.** Choose the two or three CRM workflows where a bad agent decision would have the greatest impact, and apply all four foundations there. Adjust the foundations as needed, then extend them across other workflows.
- **Make governed building easy to repeat.** Give functions approved tools and reusable approaches that enable them to build quickly within enterprise guardrails. Make costs visible to the functions that generate them, and hit the switch on chargeback. Manage agents throughout their life cycle, and retire those that no longer deliver value.

The clock is ticking. Within the next year, the agent estate could outgrow the ability of many enterprises to track it. This work is far easier to perform on a cohort of 50 agents than on a cohort of 500.

Agentic SaaS offers a compelling proposition to business users—and to CIOs, too. Technology leaders can expand their impact by enabling more of the enterprise to build.

Today's agent-building capabilities are only the beginning. With the right foundations, business teams can solve more of their own problems and move faster, safely and at scale. And innovation no longer has to wait in line.

Authors



Sakshit Teotia

Platinion Principal IT Architect
New York



Rajeev Singh

Managing Director, BCG
Platinion
Dallas



Matt Emara

Platinion Director
Toronto



Karan Gupta

Platinion Director
New York



Srikanth Srinivasan

Platinion Manager, Enterprise
Solutions
Atlanta



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.