Everforth Quinnox

The 5-Dimension Framework for Evaluating Salesforce Agentforce Partners



©2026 Everforth, Inc. All rights reserved.

Foreword

In twenty-two years across enterprise technology, software delivery, and digital transformation - fifteen of them inside the Salesforce ecosystem - I have watched several platform shifts arrive with the promise of changing everything. Most delivered incremental gains. Agentforce is different, and I say that as someone whose job depends on not overstating these claims to clients.

What sets Agentforce apart is simple to state and hard to engineer: for the first time, we are asking Salesforce to act on intent rather than instruction - to reason across an organization's data estate, decide a course of action, and execute it. That is a genuinely new category of system to design, govern, and operate, and I have seen first-hand how few organizations - and how few partners - are actually ready for it. The pattern repeats itself across every wave of automation that came before: the technology is sound, but data foundation, governance discipline, and change management get treated as afterthoughts. They are not afterthoughts. They are the programme itself.

This guide exists because the gap between a partner who can talk fluently about Agentforce and one who has actually operated it in production - through the data failures, the edge cases, the budget surprises - is very difficult to see from a pitch deck. The five dimensions that follow are drawn directly from where Agentforce programmes actually fail. Use them to separate genuine production depth from well-packaged certification - your organization's first Agentforce deployment will set the tone for every one that follows.



Murali G Sivanandam
Director - Salesforce,
Everforth Quinnox

Table of Contents

01 **Chapter 1:**
Introduction

02 **Chapter 2:**
The Six Challenges That Derail Agentforce Programmes

03 **Chapter 3:**
How the Right Partner Changes the Outcome

04 **Chapter 4:**
How to Select the Right Partner: The Five Evaluation Dimensions

05 **Free Download:**
The Agentforce Partner Evaluation Kit

The Bottom Line

Make your partner evaluation more rigorous

Download the Agentforce Partner Evaluation Checklist

Evaluate up to 3 Salesforce Agentforce partners across
5 critical dimensions and 35 evaluation criteria, with evidence prompts and a clear decision score.

35 evaluation criteria

Evidence prompts

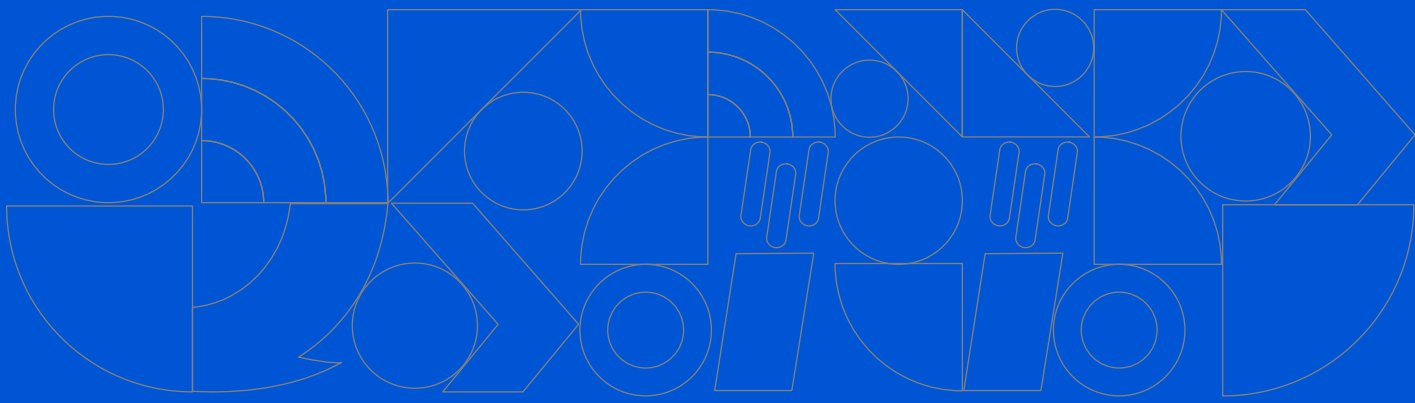
Compare up to 3 partners

Clear decision score

Keep the checklist open as you read this guide and use it to evaluate each partner against the criteria that matter most.

Download the free checklist

Chapter 1:
Introduction



Agentforce: A New Class of Enterprise Software

Salesforce Agentforce, launched at Dreamforce 2024, is not an incremental update to Einstein or a rebranded chatbot. It is a fundamentally different class of enterprise software: an autonomous AI agent platform that can reason, decide, and act across business processes without requiring a human to initiate every step.

Where traditional Salesforce automation - Flows, Einstein bots, macros - executes fixed instructions, Agentforce agents operate on intent. You define the goal and the guardrails. The agent determines how to achieve it, which tools to invoke, which data to query, and when to escalate to a human. This shifts Salesforce from a system of record and engagement to a system of action.

The practical implications are significant. A service agent can autonomously handle a returns request from initiation to resolution - querying the order history, verifying eligibility, processing the refund, and updating the customer - without a human touchpoint unless a defined threshold triggers escalation. A sales agent can proactively identify at-risk accounts, synthesize signals from across the data estate, and generate a personalized outreach - before a sales rep has even opened their pipeline view.

> "We have over 1,000 paid Agentforce deals already, with more than 200 going live this week alone." — **said Marc Benioff, Salesforce CEO at the Agentforce 2.0 launch**

This is why Agentforce is moving so fast. The opportunity to remove cost from operations, accelerate revenue workflows, and create genuinely differentiated customer experiences is real and near-term. Enterprises that move decisively - with the right partner - will build advantages that are difficult to replicate. Those that move without adequate preparation will incur the costs of both implementation and recovery.





Why You Cannot Do This Alone and Why Most Partners Are Not Ready Either

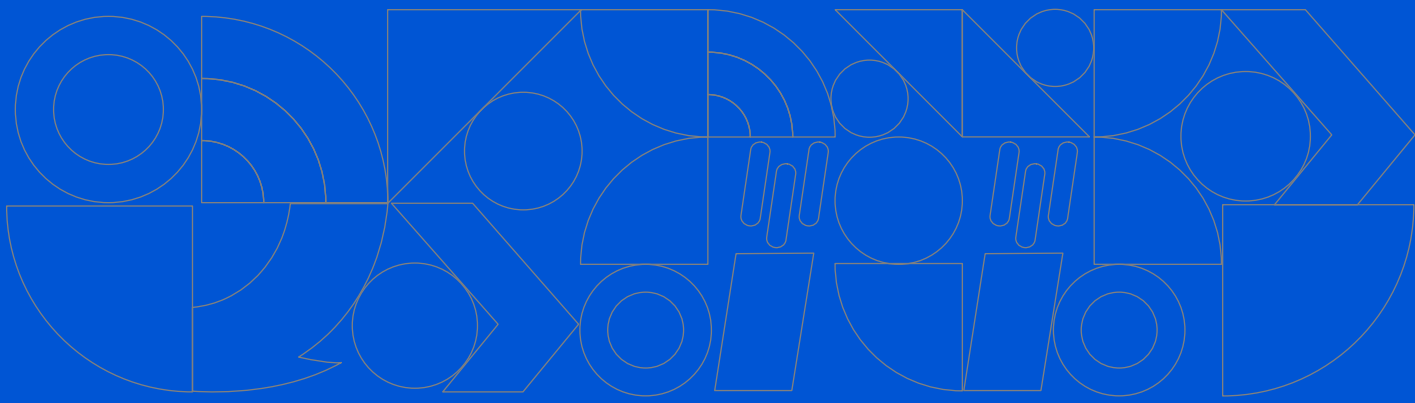
Agentforce is not a configuration project. It is not a matter of enabling a feature in your Salesforce org and connecting a few workflows. It requires a distinct set of capabilities that span data architecture, AI governance, change management, and commercial model design - disciplines that rarely sit together in a single internal team and that most Salesforce partners have not yet built to production depth.

The internal capability gap is real. Building and operating AI agents requires Data Cloud architecture expertise, understanding of probabilistic system behavior, prompt engineering discipline, compliance framework design, and ongoing agent performance management. Most enterprise IT functions are not staffed for this.

The partner market gap is equally real, and less well understood. Salesforce's rapid commercialization of Agentforce - **3,000+ registered partners within months of launch, backed by a $500M partner ecosystem fund** - has created a market where the gap between genuine expertise and well-packaged marketing is almost invisible in a proposal deck. A partner who certified their team six months ago and one who has co-developed agent architectures through multiple enterprise iteration cycles will submit proposals that look remarkably similar. Your procurement process needs to tell them apart.

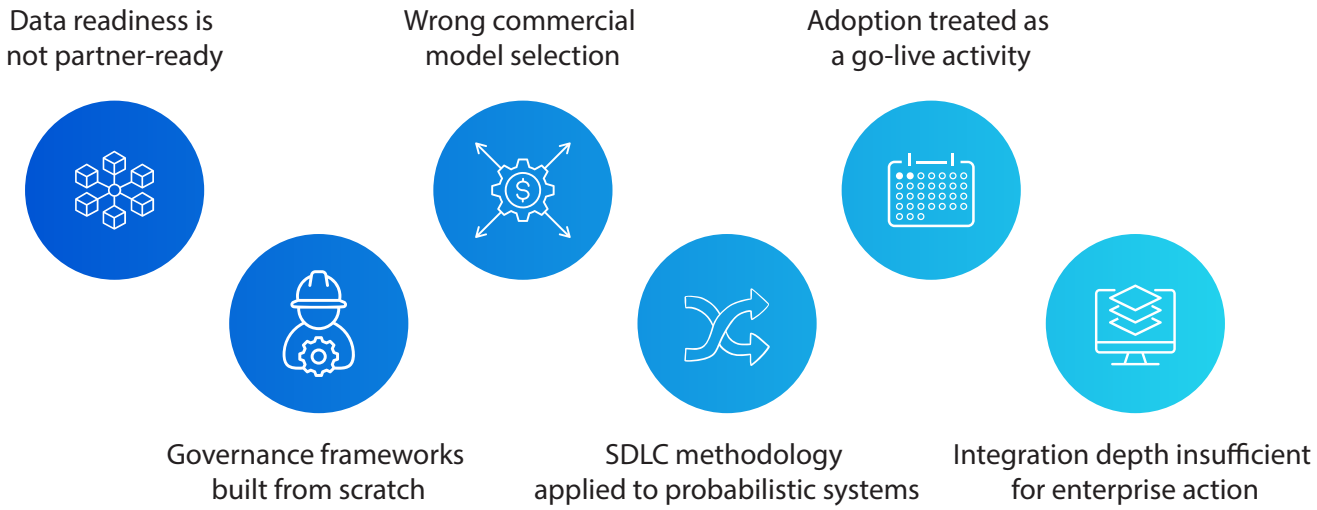
Chapter 2:

The Six Challenges That Derail Agentforce Programmes



Understanding what goes wrong - and why - is the most useful frame for evaluating what a partner must bring. These are not theoretical risks. They are the failure patterns that have already emerged across early enterprise Agentforce deployments.

6 Challenges That Derail Agentforce Programmes



01 Data readiness is not partner-ready

Most organisations have fragmented data across CRM, ERP, and service systems. Without a unified Data Cloud layer, Agentforce agents produce limited, inaccurate outputs. Partners frequently encounter this only after scoping - causing timeline and budget overrun.

Also Read: What Is AI-Ready Data & How to Assess Data Readiness for AI?

02 Governance frameworks built from scratch

Enterprises deploying AI agents into customer-facing workflows without pre-built governance expose themselves to regulatory risk, reputational damage, and incidents that erode executive confidence in the entire programme.

03 Wrong commercial model selection

Agentforce now has three pricing structures. The wrong model - or failure to model consumption at scale - can turn a well-scoped pilot into a budget crisis at production volumes.

04 SDLC methodology applied to probabilistic systems

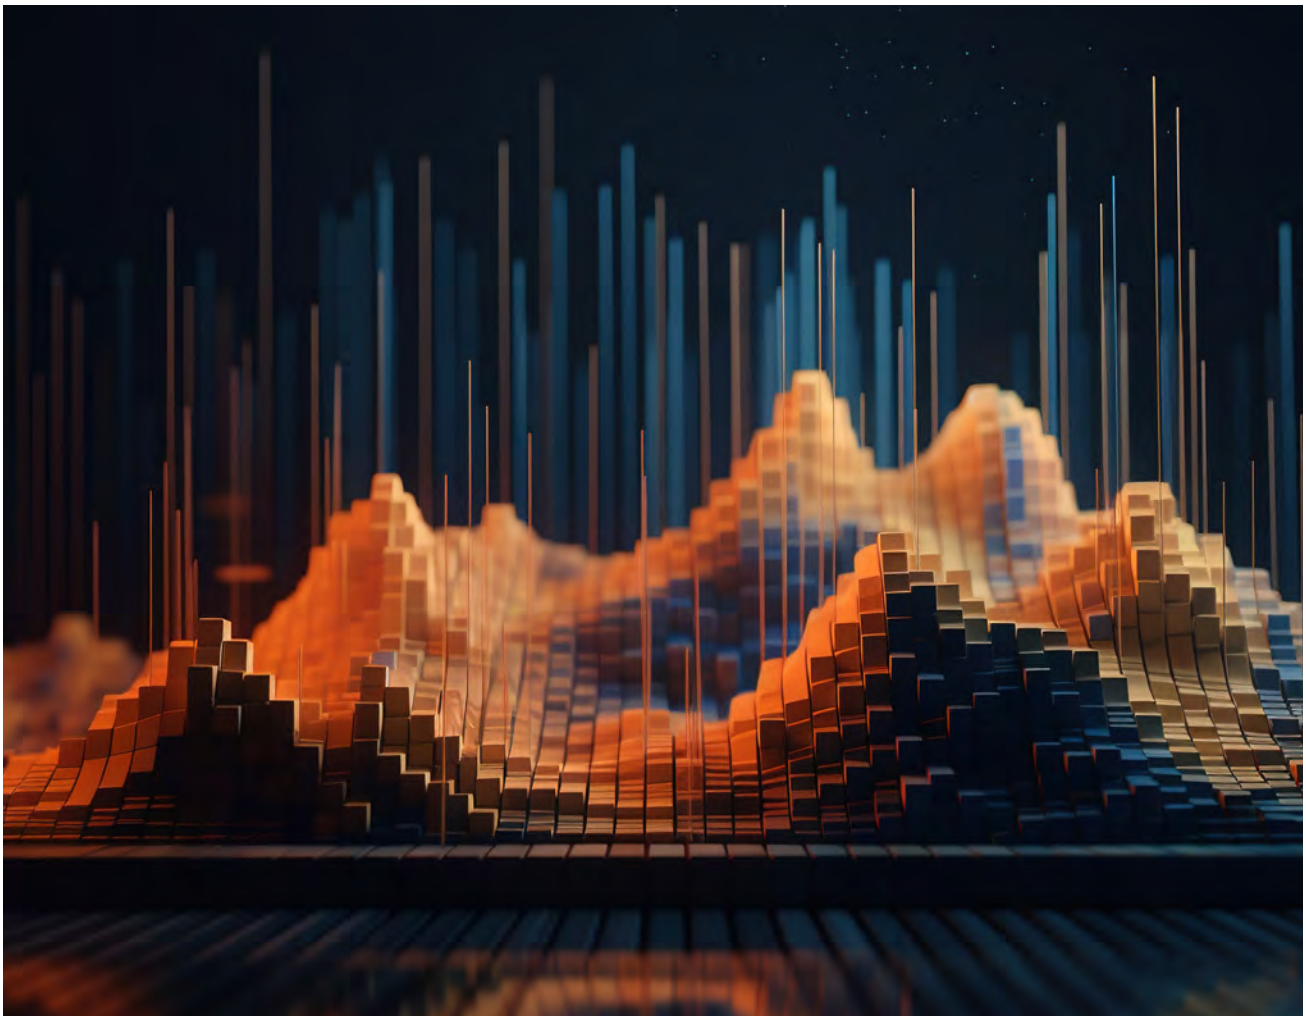
Traditional UAT and sprint-based delivery frameworks are structurally insufficient for AI agents that behave non-deterministically. Partners without ADLC (Agent Development Lifecycle) capability produce agents that pass testing and fail in production.

05 Adoption treated as a go-live activity

70% of digital transformations fail due to people and process issues. Organisations that treat change management as a post-launch task - rather than a parallel workstream from day one - consistently underperform on adoption KPIs.

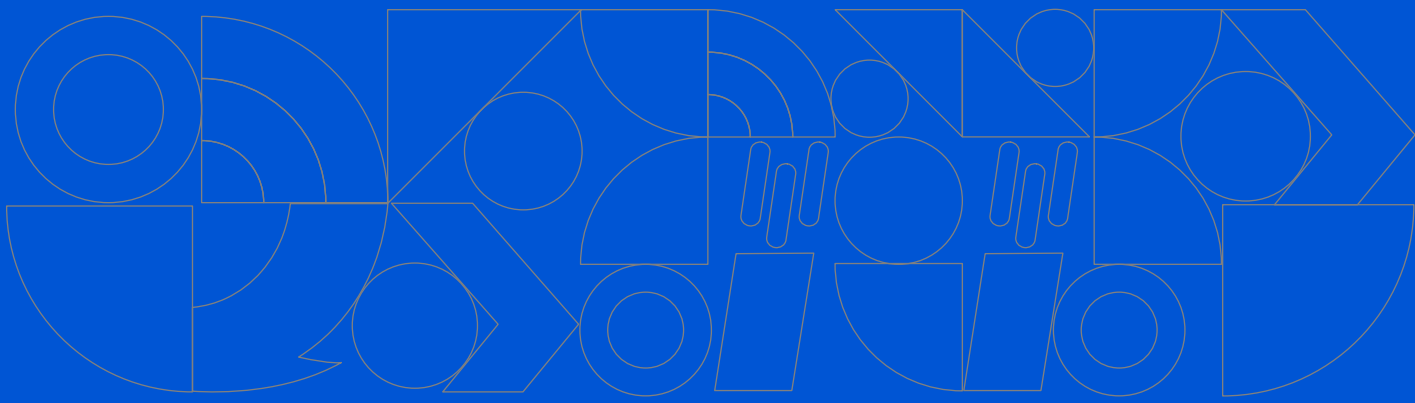
06 Integration depth insufficient for enterprise action

Agents that can reason but cannot act - because the partner's integration capability stops at Salesforce-native connections - deliver a fraction of the intended business value.



Chapter 3:

How the Right Partner Changes the Outcome



The right Agentforce partner does not simply execute a deployment. They de-risk the conditions that cause deployments to fail. The difference between a partner who delivers a working agent and a partner who delivers a working programme is the difference between a feature and a capability your organization can build on.

Here is what that looks like in practice across the dimensions that matter most.

01 They start with your data, not your use case

Before scoping a single agent, a capable partner runs a Data Cloud readiness assessment. They map your data estate, identify gaps in unification across CRM, ERP, and service systems, and define what is achievable given your current data maturity. This is not a delay to deployment - it is the difference between an agent that performs and one that hallucinates. Data quality problems discovered after go-live cost significantly more to remediate than those discovered before scoping.

02 They arrive with governance, not a promise to build it

Enterprise AI agents operating in customer-facing or operational context requires guardrails, audit trails, escalation protocols, and documented acceptable-use policies - before they go live. A partner with genuine deployment depth arrives with a pre-built governance framework. They adapt it to your regulatory context and use case. They do not propose building it as part of your engagement.
Ask any shortlisted partner directly: 'Can you share your AI governance framework document?' The answer, and how long it takes to produce, tells you everything.

03 They build to the ADLC, not the SDLC

AI agents are probabilistic systems. They behave differently across identical inputs. The ADLC (Agent Development Lifecycle) - formally defined by Salesforce across five phases — is a delivery methodology built for this reality. It replaces sprint-based UAT with behavioral testing, edge-case simulation, and continuous monitoring loops. Partners who apply traditional SDLC methodology to Agentforce produce agents that pass testing and fail in production.

04 They design for adoption from day one

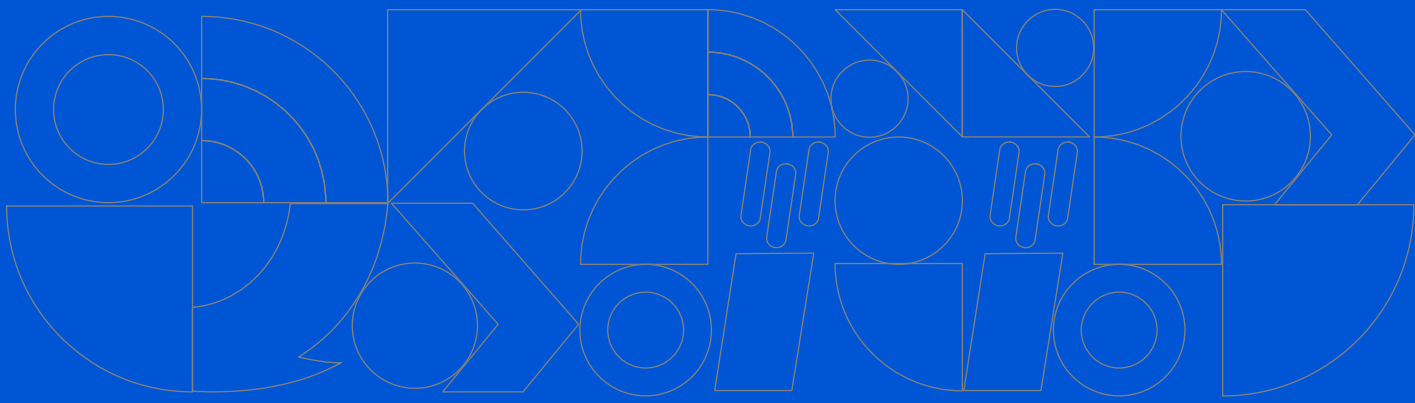
The right Agentforce partner understands that successful AI projects depend on user adoption, not just technical delivery. From the first sprint, they work on change management alongside implementation - aligning business teams, redesigning processes, training users, and measuring adoption to ensure the AI delivers real business value.
Partners who treat change management as something to address only at go-live often fall short. The most successful partners start adoption efforts alongside technical implementation from day one.

05 They model consumption costs before scoping, not after

Agentforce pricing has evolved through three models since launch. The current default - Flex Credits at $0.10 per action - is more predictable than the original $2 per conversation model, but total cost of ownership is still a function of workflow complexity, action counts, and usage volume. A commercially capable partner builds a consumption model before they write a statement of work. Those who skip this step are either inexperienced or transferring financial risk to your budget.

Chapter 4:

How to Select the Right Partner: The Five Evaluation Dimensions



The five dimensions below map directly to the challenges outlined above. They are the criteria that separate capable partners from well-marketed ones. Each maps to a specific section of the Partner Evaluation Checklist included at the end of this guide.

The Five Evaluation Dimensions



Evaluating Data and Integration Depth

Verified Competency in the New Partner Era

The Agent Development Lifecycle (ADLC)

Speed to Value and Safety Guardrails

Navigating the Consumption Economy

01 Architectural Proficiency: Evaluating Data and Integration Depth

Agentforce agents are only as intelligent as the data they can reason over. This is not a platitude - it is a hard architectural constraint. Salesforce's own documentation is explicit: Data Cloud is the foundational data layer that unifies customer signals across every touchpoint, and without it, agent outputs are confined to what exists in standard Salesforce objects. A partner who cannot demonstrate deep Data Cloud maturity is, in practice, building a capped agent.

The architectural questions to press during evaluation are precise. How does the partner approach data harmonization across ERP, ITSM, and CRM systems before agent deployment? What is their methodology for building a unified data model that agents can traverse in real time? Can they demonstrate Retrieval-Augmented Generation (RAG) grounding against external enterprise data - a capability Salesforce enabled through External Objects in Prompt Builder in 2025?

Integration depth is equally non-negotiable. Agentforce operates within a wider enterprise technology stack - connecting to identity providers, external APIs, legacy systems, and third-party platforms via MCP (Model Context Protocol). Partners who are Salesforce-native but weak on cross-platform integration architecture will hit the same wall in every deployment: the agent can reason, but it cannot act on the systems that matter most to your business.

What to evaluate

- Can they show a Data Cloud readiness assessment process conducted before agent scoping - not after?
- Have they documented MuleSoft and API orchestration experience at enterprise scale?
- Can they demonstrate RAG grounding using External Objects - a production example, not a demo?
- How do they handle multi-cloud agent architectures where data sits across Sales, Service, and Data Cloud simultaneously?

02 Technical Mastery: Verified Competency in the New Partner Era

Agentforce certification validates baseline knowledge. It does not validate production experience, architectural judgement, or the ability to navigate the failure modes that only emerge at enterprise scale. The distinction that matters is between certified teams and deployed teams.

A partner with genuine technical depth will provide architecture diagrams from production deployments, demonstrate pro-code extensibility beyond low-code configuration, speak fluently to multi-agent orchestration and A2A communication patterns, and show you their agent testing infrastructure. Those without it will describe what they do in SDLC language applied to a fundamentally different type of system.

Key verification questions

- Walk us through the Agentforce deployment at a comparable scale. What was the architecture, what broke in UAT that didn't break in the sandbox, and how did you resolve it?
- Can you share architecture diagrams from at least two live deployments? Redacted by client is acceptable - sandbox demos are not.
- Who is your named Agentforce practice lead, and how many full-time Agentforce specialists are allocated to our engagement?



03 Methodological Rigour: The Agent Development Lifecycle (ADLC)

Salesforce defines five key phases in the Agent Development Lifecycle (ADLC) for building, deploying, and managing AI agents. Knowing these phases is one thing but successfully applying them requires hands-on implementation experience.

> "The ADLC is a wholly new approach to developing software that accounts for the massive differences between developing probabilistic reasoning systems versus traditional deterministic software. When used properly, this methodology guarantees successful, robust, reliable results."
>
> **- Murali G Sivanandam,** *Director - Salesforce, Everforth Quinnox*



When evaluating a partner, ask them to explain how they use the ADLC in real projects and how they manage each phase from planning to ongoing operations.

#	Phase	What a capable partner does
01	Ideation & Design	Define agent purpose, persona, tools, and guardrails before a single line of configuration. Intent specification errors propagate through every subsequent phase.
02	Development (inner loop)	Iterative build combining low-code configuration and pro-code extensibility. Agent Script DSL enforces deterministic guardrails where probabilistic behaviour is unacceptable.
03	Testing & Validation	Behavioural testing, edge case simulation, hallucination benchmarking. Traditional UAT is structurally insufficient for non-deterministic systems.
04	Deployment	Phased rollout: internal test → controlled pilot → production. Each phase generates behavioural data that feeds the outer loop.
05	Monitoring & Tuning (outer loop)	STDM session trace analysis, routing failure detection, continuous prompt and tool optimisation. This is ongoing improvement — not maintenance.

A partner with strong ADLC experience will explain how they monitor AI agents from day one, test them against real-world and edge-case scenarios before launch, and use clear criteria to decide when the agent is ready to move from internal testing to production. A less experienced partner is more likely to rely on traditional software development practices like sprint cycles, user acceptance testing (UAT), and regression testing. While these are important, they were designed for predictable software and are not enough to manage AI systems, which produce probabilistic rather than fixed outcomes.

04 Operational Strategy: Speed to Value and Safety Guardrails

With Agentforce, enterprises can launch a pilot for a focused use case in as short as 4–6 weeks. But scaling AI across the business takes more time because it requires strong governance, clean and reliable data, and effective change management.

A partner that promises only speed risks deploying an AI solution that isn't ready for real-world use. On the other hand, a partner that focuses only on caution without providing a clear roadmap to business value may slow progress unnecessarily. The right partner delivers quick wins while building the foundation for long-term, enterprise-wide success.

Speed to value



- Pre-built use-case prioritization methodology - mapping your processes to agent archetypes and ranking by value and readiness before the SOW is written.
- Pilot-to-scale pathway with explicit go/no-go criteria at each phase transition - not a project plan that assumes success.
- Honest time-to-first-value benchmarks from comparable engagements, with ranges - not best-case projections.

Safety guardrails



- Pre-built AI governance framework - not scoped as a new deliverable. Ask to see it before the engagement is signed.
- Documented hallucination mitigation methodology: RAG grounding, prompt constraints, output confidence scoring.
- Designed escalation protocols with defined failure rate thresholds and pre-production testing.
- Compliance artefact library for your regulatory context - financial services, healthcare, or public sector.
- Pre-written incident response playbook for AI agent failures in production.

4–6 wks

Typical deployment timeline for a scoped Agentforce pilot vs 6–12 months for equivalent in-house builds

Source: Salesforce

98%

Of enterprises investigating AI automation, but only 20% feel prepared to deploy at scale

Source: TruSummit

90%

Of CIOs say managing AI costs is limiting their ability to drive value

Source: Salesforce

05 Commercial Insight: Navigating the Consumption Economy

This is where most enterprise buyers are least prepared - and where the costliest surprises emerge. Agentforce runs on consumption-based pricing, not traditional per-seat licensing. Salesforce has iterated through three pricing models in under 18 months. The right model for your deployment depends on use case type, volume predictability, and team structure and a commercially capable partner will build a cost model for your specific context before scoping begins.



Original model

$2 per conversation

Launched Oct 2024. Unpredictable at enterprise scale — a single query could trigger 8+ backend processes. Significant buyer backlash.



Flex Credits — current default

RECOMMENDED
$0.10 per action

$500 per 100,000 Flex Credits. Introduced May 2025. Charges per discrete action — not conversation. Ties cost to real business outcomes.



Per-user licensing

$125 / user / month

Unlimited internal agent usage. Often combined with Flex Credits for external-facing use cases.

Flex Credits and Conversations cannot be used simultaneously in the same Salesforce org. Total deployment cost includes Data Cloud licensing, base edition costs (Salesforce raised Enterprise/Unlimited pricing by 6% in 2025), training, and ongoing agent management - all of which frequently sit outside the initial partner scope. A partner who does not proactively surface these costs is either commercially inexperienced or commercially convenient.

Three pricing models. All running at the same time. On the same product. When the market hasn't converged on how to buy something, you don't force one model. You let customers self-select into the model that matches how they want to buy.

The strongest commercial signal: ask any shortlisted partner whether they will accept partial outcome-based fees tied to defined KPIs - deflection rate, CSAT improvement, resolution time. Willingness to engage with this, even partially, signals genuine confidence in their delivery quality. Reflexive refusal, without a substantive reason, signals the opposite.

FREE DOWNLOAD - Agentforce Partner Evaluation Kit

TAB 1:
Organizational Readiness Assessment

Score your organization across 5 readiness dimensions before shortlisting any partner. 22 scored factors with guidance, priority flags, and an overall readiness band (High / Moderate / Low). Identifies the gaps partners will charge you to remediate - so you can address them first or scope them explicitly.

TAB 2:
Partner Evaluation Checklist

Score up to 3 partners across the same 5 dimensions with 35 evaluation criteria; evidence prompts for each criterion, and an automated overall score with a decision band (Strong Shortlist / Conditional / Weak / Do Not Shortlist). Built to be shared with executive stakeholders as a procurement artefact.

Download Now

The Bottom Line

The partner you choose for Agentforce is not simply a vendor selection- it is a strategic decision that will influence how your organization adopts, governs, and scales AI over the coming years. The five dimensions outlined in this guide - architectural proficiency, technical mastery, methodological rigor, operational strategy, and commercial insight - provide a practical framework for distinguishing genuine Agentforce expertise from marketing-driven positioning.

As enterprises move from AI pilots to enterprise-scale autonomous operations, success will depend on selecting partners that can bridge Salesforce innovation with real-world business transformation. The most effective partners bring not only technical implementation skills, but also deep experience in data modernization, governance, integration, and change management.

At Everforth Quinnox, we believe Agentforce represents the next evolution of enterprise automation. By combining Salesforce expertise with AI-led engineering, enterprise integration, and business transformation capabilities, we help organizations accelerate their journey from AI experimentation to tangible business outcomes - while minimizing risk and maximizing long-term value.

Connect with our Agentforce experts for a personalized demo and see how the right implementation partner can help your organization get it right from day one.
Schedule a Demo Now

About Everforth Quinnox

Everforth Quinnox is an AI-first Digital Always organization orchestrating value for forward-looking businesses. Through AI-powered platforms and intelligent engineering teams, we accelerate digital transformation, redefine total experience, and improve productivity and operational efficiency.

Our AI-first workforce turns challenges and opportunities into a competitive edge, fueling growth, sparking innovation, and building resilience for accelerated time to value.

At the core of Everforth Quinnox's innovation is the Services as Software (SaS) model, transforming traditional IT services into an intelligent, self-evolving ecosystem. Powered by Agentic AI, this ecosystem autonomously designs, tests, orchestrates, and manages operations, continuously learning and improving to deliver scalable, resilient, and outcome-driven business value.

Everforth Quinnox is headquartered in the United States, with a strong global delivery presence across the US, UK, and India. Our work is consistently recognized by leading industry analysts for its real, tangible impact and is strengthened by a robust partner ecosystem.

For more information please contact:
marketing@quinnox.com | www.eveforthquinnox.com

©2026 Everforth, Inc. All rights reserved.