



“How many AI initiatives in your organization made it to pilot... and quietly stalled before production?”

Deploy Your Agentforce by Summer '26

Turning AI Investment into Measurable Gains



Anand Khambete
Belgium Salesforce CTO and Certified Technical Architect





Top takeaways for you

- 01** To understand why it is the time now to move towards agentic
- 02** Golden principles for enterprise agentic architecture
- 03** Public sector use case - “Reimagining Case Participant Management with Agentforce”
- 04** Life science use case- “Evolving from Veeva Insights to Agentforce-Driven Next Best Actions”
- 05** A practical guide to help you to make your steps towards Agentic world.



Why 2026 Is The Year of Enterprise AI Agents



Not because AI suddenly got smarter, but because enterprises are finally ready

The pilot era is ending

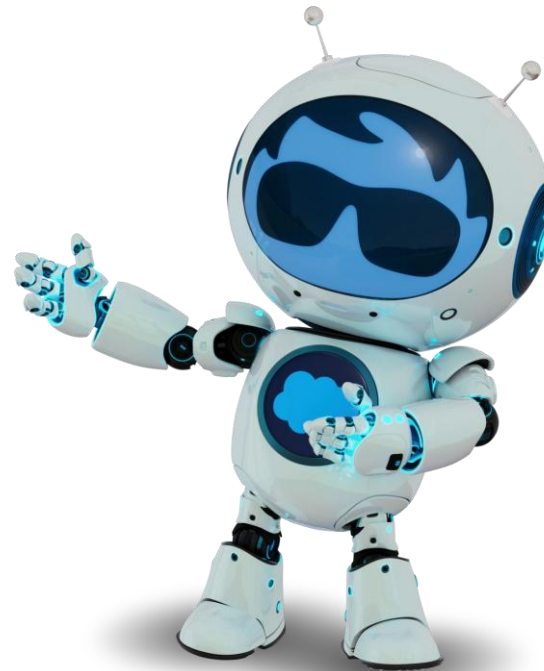
AI is moving from experimentation to **production accountability**.

ROI has replaced innovation hype

Boards and CFOs now demand **measurable outcomes, not demos**.

Competitive advantage is shifting fast

By end-2026, leaders will **run parts of the business on agents**, Others will still be piloting.





Agentic Enterprise Architecture

Intelligence alone is not the differentiator

System of engagement

Slack
 Slackbot | Canvas | Enterprise Search | Messaging & Huddles | Tableau Viz & Data Q&A

Any workspace

System of agency

Agentforce Customer & Employee Agents
 Voice, Web, Mobile & Messaging | Agentforce Script | Builder | Observability | Orchestration

Any agent
 MCP & A2A

System of work

Customer 360

Sales | Customer Service | Field Service | IT & HR Service | CCaaS | Mktg | Commerce | Supply Chain | Revenue Mgmt | Platform

Fin Serv | Health & Life Sciences | PubSec | Manufacturing | Consumer Goods | Comms | Nonprofit | Education | Energy & Utilities | Media

Any app

System of context

Data 360
 CDP | Federation | MuleSoft | Informatica | Tableau | Structured & Unstructured | Zero copy | Real-time

Any data lake or warehouse

Trust layer

OpenAI | ANTHROPIC | Gemini | LLaMA by Meta | Open source



5 golden principles for agentic enterprise



5 golden principles

Agentforce is not a one-time feature rollout — it is an operating-model shift

Success depends on how work gets redesigned, not which buttons get clicked.

Agentforce should not mimic today's processes — it should challenge them

Don't automate inefficiencies; rethink what work should exist at all.

Agentforce is not owned by IT alone — it is co-owned by business, IT, and governance

Clear ownership and accountability drive production success.

Agentforce does not replace decision-makers — it amplifies decision quality and speed

The goal is better decisions, faster execution, and fewer handoffs.

Agentforce is not about scaling AI — it is about scaling outcomes

Adoption metrics matter less than cycle-time reduction, cost savings, and revenue impact.



Use cases



Reimagining Case Participant Management with Agentforce

Public sector case management often spans multiple internal and external entities (citizens, NGOs, lawyers, agencies).

Do you lack a **single, reliable overview of all case participants**?

Is **context lost** across systems, emails, and handovers?

Is adding or linking participants **manual, error-prone, and slow**?

Is your case progress delayed due to **administrative overhead**?



Reimagining Case Participant Management with Agentforce

Agentforce Solution

An Agentforce Case Participant Agent that:

- Provides a **real-time, consolidated view** of all participants linked to a case
- Allows users to **add or link individuals or organizations using natural language**
- Works across **internal and external participants** seamlessly

Benefits*

- 40–60% reduction in manual effort for adding/linking participants.
- 20–35% reduction in overall case cycle time.
- 30–50% improvement in data completeness.

* Bases on internal research and testing

The screenshot displays a chat window titled "Internal Agent". The chat history includes:

- A user question: "To which case is pierre lambert with national security number 85031245623 related to?"
- An AI response: "Pierre Lambert is related to the following cases:
1. **Case Number: 00001026**
• Role: Defendant
2. **Case Number: 00001027**
• Roles: Section Chairman, Counselor, Advocate General
Would you like to take any specific action regarding these cases?" (with thumbs up, thumbs down, and a folder icon)
- A user question: "Who is participant in case 00001026"
- An AI response: "Here are the participants in case **00001026**:
Legend:
• Internal Participant (green dot)
• External Participant (red dot)
1. **Pierre Lambert** (red dot)
• Role: Defendant
2. **City Social Services** (green dot)
• Role: Section Chairman
3. **Sophie Martin** (red dot)
• Role: Lawyer
• Employer: Law & Co.
4. **Alex Green** (red dot)
• Role: Plaintiff
Let me know if you need more details!"

At the bottom, there is a text input field with the placeholder "Describe your task or ask a question..." and a send button.



Evolving from Legacy Veeva Insights to Agentforce-Driven Next Best Actions

In a Veeva-driven model, identifying the next best action for an HCP is highly manual, fragmented, and inconsistent.

Field reps must **manually review HCP interaction history** (calls, emails, CLM activity)

Prone to **missed opportunities and inconsistent HCP experience**

Next best actions are **subjective and rep-dependent**



Deploy an Agentforce intelligent agent to deliver real-time, compliant and personalized next best actions for each HCP.

Agentforce Solution

- Analyses **complete HCP engagement history** across channels.
- Consider regional compliance and regulatory constraints.
- Recommends **next best actions in real time**.
- Ensures every interaction is **scientifically relevant and compliant**.

Benefits

10–15% increase in HCP engagement rates.

30–40% reduction in time spent preparing for HCP interactions.

10–20% increase in rep productivity.



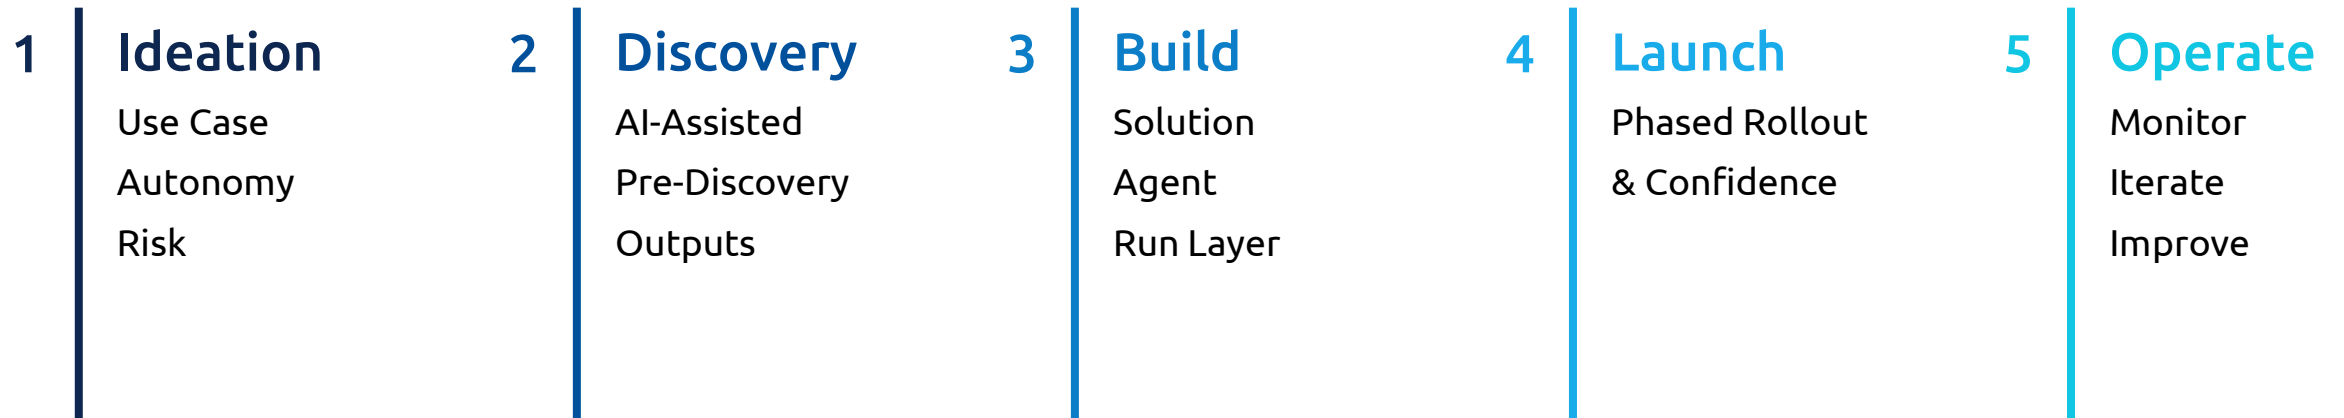


A practical guide towards an agentic enterprise : from today to Summer 2026



Five phases. One connected journey.

From first intent to in-production — and iterating





1 | Ideation

Use case

What problem are we solving?

What type of capability is needed?

Each type drives different architecture, grounding, and testing choices.

Autonomy level

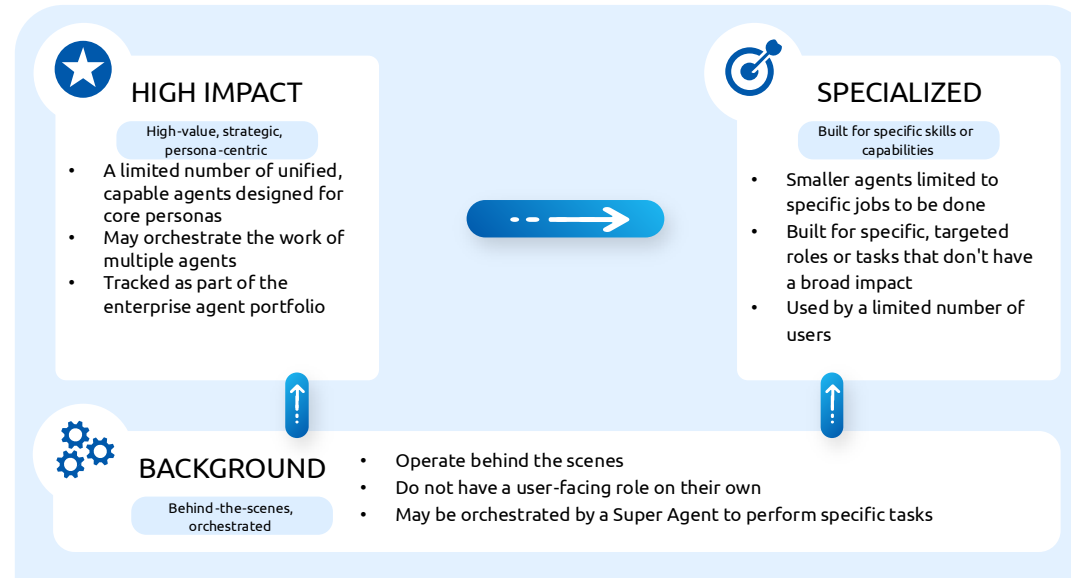
Where does this sit on the spectrum?
Human-in-the-loop, assisted, semi-autonomous, or fully autonomous?

This sets risk posture, approval process, and go-live criteria. Not a build-phase decision.

Risk tolerance

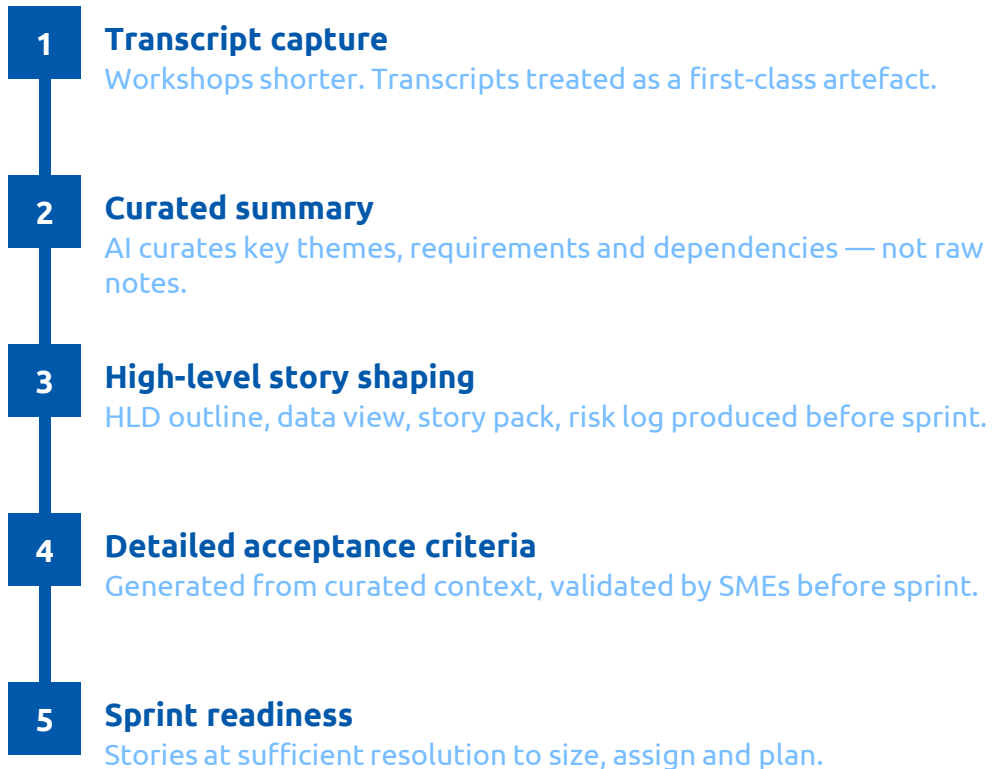
What is the tolerance for hallucination, incorrect actions, or missed edge cases?
What are the failure modes? Who owns the risk?

If not agreed before build, governance will be retrofitted and fragile.





The multi-stage delivery pattern



Four outputs required before build

- **HLD Outline**
End-to-end system boundaries, interfaces, data flow, security posture
- **Data & Grounding View**
What data exists, what grounding strategy, what retrieval approach
- **Initial Story Pack**
Use case stories at sufficient resolution to size and sprint-plan
- **Risk & Assumption Log**
Named owner, residual risk, explicit constraints



1 | Ideation

2 | Discovery

3 | Build

1

Data Layer

Data structures, objects, and models in place. Clean, accessible, permissioned.

3

Prompt Templates

Start simple. Persona separate from task. Explicit NO DATA FOUND exit. Version controlled.

5

Agent Configuration

Persona, scope, safety instructions applied. Edge cases documented and regression-tested.

2

Einstein Retrievers

Vector search, semantic retrieval, knowledge indexes configured and tested.

4

Topics & Actions

Classification descriptions accurate. Flow/Apex actions handle logic. Templates handle output.

6

Test & Tune

Quality evaluation with business SMEs. Prompt versioning. Conversation log review.

Thank you for attending.

Make it real.

