Gen25

06
NOV
26

AI service agent package: answering your FAQ

fast and decisive



© Gen25

AI service agent package: answering your FAQ fast and decisive

Service leaders are expected to put AI to work this year. The safest place to start is the questions your team already answers every day. Those answers are already written down in your knowledge articles, the volume is high, and the outcome is easy to measure.

Gen25's Salesforce FAQ Agent is an Agentforce service agent that answers customers only from your approved knowledge articles. It shows where each answer came from, and when the answer isn't there, it says so and hands over to your team.

Use case in practice

At a large Dutch retailer handling around 1,500 chat sessions a day, the self-service rate in chat rose from around 40% to around 70%. That is roughly 400 fewer contacts a day reaching a person.

This document explains why routine questions make the right first use case, why an agent that says less is worth more than one that guesses, and how scope, delivery and measurement work. It closes with a readiness checklist you can complete in a single conversation.

The pressure on service leaders

AI has moved from experiment to expectation in customer service.

91%	Service leaders feeling executive pressure to implement AI in 2026
AI rank among service leaders' priorities	#10 → #2 in one year
30% today → 50% expected by 2027	Share of cases handled by AI

Based on research by Salesforce and Gartner



The pressure is real, but so is the risk. A broad AI programme can take months before anyone sees a result, and one public mistake can set back trust inside the organisation for a long time. Service leaders need a starting point that is quick, safe and measurable.

Why routine questions are a valuable use case

The same questions reach your team every day, and the answers already exist. That combination makes routine questions the best first job for an AI agent.

- **The content is already there**. Returns, delivery times, opening hours, payment options: your team answers these from articles or from memory. Once that knowledge is in articles, an agent can use it.
- **The volume is high**. A small set of contact reasons typically drives a large share of contacts. Taking those off the queue changes the workload straight away.
- **The risk is contained**. Answering a policy question needs no access to orders, payments or customer records. That keeps scope, security review and integration effort small.
- **The result is measurable**. Contacts reaching your team, the self-service rate and answer accuracy can all be counted from the first conversation.

The business challenges behind it are familiar. Skilled people spend their time on simple questions, and the cost per contact keeps rising. Growth means hiring, while peaks are hard to staff. Knowledge leaves when people leave. Customers can't reach you outside office hours.



Say less rather than guess

The biggest risk of an AI agent is not that it can't answer. The bigger risk is that it answers something you never wrote. Take a question your articles don't cover: "Do you refund shipping costs on returns?"

An AI that guesses

- **Answer:** “Yes we refund shipping costs on every return within 14 days”
- **Where it came from**: General knowledge, not your articles
- **What happens next**: It sounds right, but it isn't your policy. Your team cleans it up later, often after a complaint.

The Gen25 FAQ Agent

- **Answer:** “I couldn't find that in our information. Shall I connect you with a colleague?"
- **Where it came from**: Only your articles
- **What happens next**: It is an honest answer. Your team picks it up from there.

That behaviour is a design choice, not luck. The agent's configuration contains one explicit rule: "Never provide generic information, advice or troubleshooting steps, unless retrieved from searching knowledge articles." That rule is part of the scope you sign.

The result is an agent that may answer slightly fewer questions, but whose answers you can stand behind. Automation must never cost customer satisfaction.

What we offer in the package

The agent sits in the chat window on your website and is built on Agentforce and Salesforce Messaging.

1. The customer asks a question in their own words.
2. The agent searches your knowledge articles, in Salesforce Knowledge or in sources landed in Data 360.
3. If it finds the answer, it replies and shows the source article.
4. If it doesn't, it says so honestly and offers a colleague.
5. The conversation is routed to your team's queue, with the chat history attached. Outside opening hours, it waits there until someone is online.

Where your content lives decides the size

Size S

- Duration: ~10 working days
- Your content: in Salesforce Knowledge

Size M

- Duration: ~15 working days
- Your content: plus one source outside Salesforce, landed in Data 360

Size L

- Duration: ~20 working days
- Your content: Plus two sources outside Salesforce, landed in Data 360



Each outside source adds five working days. More than two sources is a separate project.

In every size

- One chat channel on your website
- Answers from your knowledge articles
- Handover to your team, routed to your queue
- Tracking switched on from day one

In no size

- Rewriting or updating your knowledge
- Order status, returns and other lookups
- Voice, WhatsApp or a second channel
- Maintenance after handover

- **Licences you need**: Service Cloud, Agentforce, Data 360, Messaging and Salesforce Knowledge.
- **Add-on**: post-chat NPS. A short survey appears in the chat window the moment the conversation ends. Each score is saved on the chat in Salesforce, ready for reporting. It works on any Salesforce web chat, so you can compare AI-handled and human-handled conversations side by side.

From challenges to goals to KPIs

We agree what success means before go-live, then report against the same numbers before and after.

Challenge	Goal	KPI
The same questions reach the team every day	Self-service	Top contact reasons; self-service rate
Skilled people spend time on simple questions; cost per contact keeps rising	Lower cost to serve	Contacts reaching your team; contacts no longer reaching your team
Growth means hiring; peaks are hard to staff; knowledge leaves when people leave	Resilience	Peaks handled without extra staff
Customers can't reach us outside office hours	Accessibility and speed	Questions answered outside opening hours
Automation must not cost customer satisfaction	Customer satisfaction	Answer accuracy on your questions; NPS and CSAT held

The baseline is recorded in week 0. Testing on day 6 produces two figures: did the agent route the question correctly, and was the answer right. Tracking starts with the first live conversation.

Ten working days: the delivery method

The agent is live on day 7, and you have a delivery report on day 10. Delivery runs in seven steps across three phases.

Phase	Goal	Step	
Before day 1	Week 0	Baseline	Agree goals and KPIs; record today's numbers.
During project	1	Provisioning	Licences confirmed, org and users provisioned, AI accounts created with scoped permissions.
During project	2-3	Agent and subagent	Agent and prompts configured against your knowledge content.
During project	4-5	Messaging deployment	The chat channel configured and connected to your website. This is the largest single step.
During project	6	Test	Answers tested against the agreed topics. You approve.
Go-live	7	Deploy	A person deploys to production, and the agent goes live.
Adoption	8-9	Hypercare	Live traffic watched; answers tuned.
Adoption	10	Handover	Report on what customers asked, what the agent couldn't answer, KPIs against baseline, and next steps in four areas.
Adoption	30	Review	With a month of traffic, content gaps identified and next steps agreed.

What we need from you in week one: the team that owns your website, because the chat widget is deployed on your side.

Results and what it could mean for you

De Bijenkorf, a large luxury retailer in the Netherlands, runs the agent on web chat with around 1,500 chat sessions a day.

- **Self-service rate in chat:** from around 40% to around 70%, holding structurally
- **Contacts reaching a person:** around 400 fewer a day at comparable volume

To estimate the value for your organisation, multiply the contacts the agent takes over by your own cost per contact:

Research shows

Published Salesforce customer results for service agents range from 28% to 84% self-service, so around 70% is credible rather than unusual. What nobody publishes is how long it took to build.

Annual saving = chats per day x share handled by the agent x days per year x (cost per assisted chat - cost per AI chat)

We complete this calculation with your own numbers during the week 0 baseline.

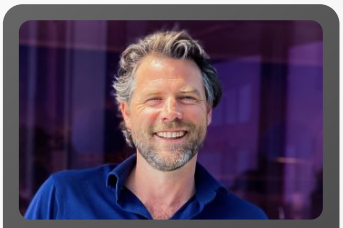


Your consulting partner, turning business challenges into impact

Implementing customer service AI should be fast, safe, and measurable. Built on Salesforce Agentforce, the Gen25 FAQ Agent resolves routine questions using only your approved knowledge articles. By choosing to hand over to your team rather than guessing, it protects customer trust while taking immediate pressure off your service queues.

Our fixed-scope delivery model moves you from baseline to a live deployment in just seven working days, backed by a clear performance report on day ten. Gen25 stays involved post-launch to monitor live traffic, identify content gaps, and steadily increase self-service rates—helping you lower costs and handle peak volumes with ease.

Ready to move from AI pressure to proof in ten working days? Let's discuss your challenges!



Bas Verkleij

sales@gen25.com

Some of the customers we serve

