



# From ideation to action with Agentforce

Learn from Salesforce customers who have deployed Agentforce, sharing 25 practical lessons from over 1000+ live implementations.

The image displays the Agentforce user interface and a 3D AI agent character. The interface includes a navigation menu with icons for a clock, a lightbulb, and a brain. The main content area features a 'Welcome to Agentforce' message and three primary actions: 'Build an Agent', 'Practice Creating an Agent', and 'Learn About Agentforce'. A detailed view of the 'Describe what you want Agentforce to do' form is shown, with fields for Name, API Name, Role, Company, and Agent User. The form also includes checkboxes for 'Assign user permissions based on agent force and actions' and 'Fetch event logs with conversation data'. The background features a stylized landscape with mountains, trees, and a blue sky with clouds.

**Agentforce**

### Welcome to Agentforce

Resolve customer cases, improve sales, and anticipate customer needs with autonomous agents. When you're ready to dive in, here are three ways to start using Agentforce.

- Build an Agent**  
Navigate to Agents in Setup and select **New Agent** to create one from scratch or start with a pre-built skill.
- Practice Creating an Agent**  
Follow detailed pointers and step-by-step instructions on Trailhead to practice setting up and deploying an agent.
- Learn About Agentforce**  
Check out an overview or dig into the details of setting up, launching, and fine-tuning Agentforce.

**Describe what you want Agentforce to do**

Name:  API Name:

Role:

You are an AI sales representative that contacts prospective leads and gives them product and pricing information.

Company:

Agent User:  **New Agent User**

Assign user permissions based on agent force and actions  
 Fetch event logs with conversation data

# A few words from a Customer Agentblazer



**Rudy Khoury**  
Chief Digital Officer,  
Fisher & Paykel

Agentic AI is a generational opportunity to redefine how we run our businesses, how we serve our customers and how we elevate our employees.

At Fisher & Paykel, we've been on the Agentforce journey for a year now - and we've learned that the real differentiators are mindset and pace. We set an ambitious vision and within weeks, we **doubled our digital self-serve resolution rate from 30% to 60%**, freeing time for our people to deliver more of the personalised world-class service that defines our brand.

That's the power of moving from idea to real business value. Every week you wait, your competitors are getting stronger, you're delaying taking your teams on the journey and you're missing opportunities to improve every area of your business.

So identify a valuable use case, follow the lifecycle, learn from the 1000s of customers already live, move fast, and learn quickly.

**Your next 90 days will set the pace – how will you use them?**

*Rudi*

**Introducing the five stages of Agent Deployment:**

**1 Ideate**

Every new agent starts with a valuable outcome. In this stage, you're identifying high-impact opportunities where an agent can step in – whether it's to automate a process, enhance a customer experience, or augment your team behind the scenes.

**2 Build**

Once you've got the use case, it's time to bring the agent to life. You'll configure its goals, the systems it connects to and how it handles tasks – all while aligning it to your business logic, data, and trust requirements.

**3 Test**

Before the agent sees the light of day, you put it through its paces. This is where you validate performance, simulate scenarios, get feedback, and make sure it's safe, accurate, and effective.

**4 Deploy**

It's go-time. The agent is live and working in the real world – helping customers, employees, or partners. This stage is all about ensuring adoption, monitoring usage and making sure it's delivering value, across your channels, from day one.

**5 Optimise**

Even after launch, the work continues. You'll monitor how the agent performs, learn from feedback, and keep tuning it to be faster, smarter, and more impactful.

**Agent Development Lifecycle**



For each stage in the Agent Development Lifecycle we'll be sharing **25 lessons learned from 15 ANZ Customer Examples**

# Introduction



Since we launched Agentforce in November 2024, we've seen industries test, learn, and transform how they serve customers, support employees, and engage partners.

With thousands of customers live globally – including more than 75 in ANZ – we're excited to share a series of lessons learned to help accelerate your journey towards becoming an Agentic Enterprise.

At Salesforce, we've mapped the various stages of the agent journey into a clear **Agent Development Lifecycle framework**.

What we've seen is that customers who skipped critical stages – like building before defining clear outcomes, or deploying without thorough testing – ran into the biggest roadblocks. Considering the full lifecycle helps avoid these pitfalls and accelerates time to value.

For each stage in the  
Agent Development Lifecycle we'll be sharing  
**25 lessons learned from**  
**15 ANZ Customer Examples**

## Authors and contributors



**Jonathan Hatcher**  
VP AI Solution Engineering, ANZ



**Jessica Raicebe**  
Senior Director, AI Product Growth



**Keran Wijetunga**  
Director AI Specialists



**Priya Kanjia**  
Platform, Product Marketing, APAC



**Cat Prestipino**  
Field Marketing Director, ANZ



**Jacob Hkeik**  
Agentforce & Data 360, Product  
Marketing, APAC


---

Please note that this document contains general information only and should not be taken as providing legal advice. All organisations should seek independent legal advice about their own privacy obligations before making any purchasing decisions. This document is based on Salesforce's interpretation of the regulations, is provided "as is" and is subject to change. We endeavor to assist our customers in their compliance journey, but we exclude liability for any representations or warranties made to our customers relating to regulation and its relationship to our solutions in our communications with customers.



# Contents

<b>Chapter 1</b>   Ideate.....	07
<b>Chapter 2</b>   Build .....	12
<b>Chapter 3</b>   Test .....	18
<b>Chapter 4</b>   Deploy.....	23
<b>Chapter 5</b>   Optimise .....	28
Conclusion and closing thoughts.....	33



# 1

## Ideate

Every new agent starts with a valuable outcome. In this stage, you're identifying high-impact opportunities where an agent can step in – whether it's to automate a process, enhance a customer experience, or augment your team behind the scenes.





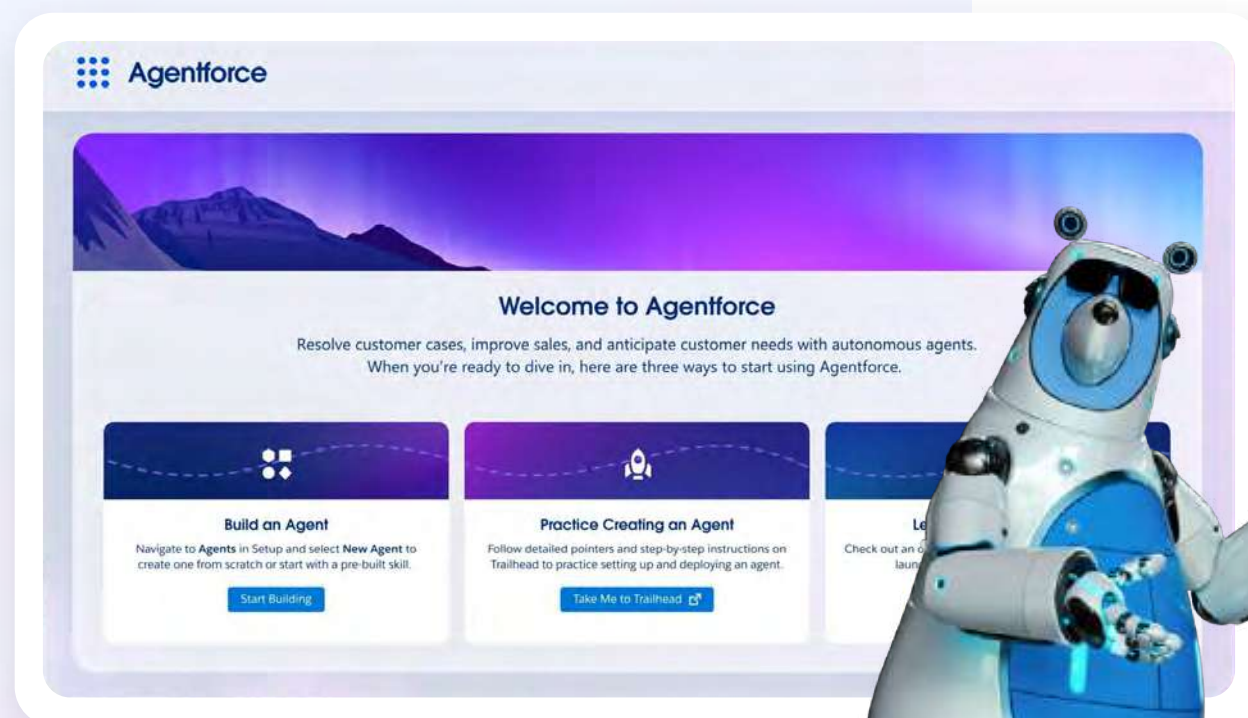
1

# Every successful Agentforce deployment begins with a clear vision in the Ideate stage.

This foundational phase is dedicated to identifying high-impact opportunities where an agent can deliver significant value - whether it be automating processes, enhancing customer experiences, or augmenting your existing teams.

Even with some of our largest customers, executives often play a crucial hands-on role in this stage. They help keep the process "raw" by engaging directly with teams and leaders to uncover real pain points, ensuring that enough time is spent to define truly valuable problems without letting the exercise become "corporate entertainment" that stalls execution.

By focusing on valuable outcomes from the outset, businesses can pinpoint the most promising areas for agent intervention and lay the groundwork for a truly transformative solution.





Let's explore the **5 lessons learned** from our customers during the **Ideate** phase:

### 1 Talk to the humans behind the change, they know where the opportunities are.

The most valuable, and most bought into, ideas emerge when you go beyond endless process maps and hear directly from the people closest to the work. Workshops that bring together customer, partner, and employee insight from across functions often uncover not just new agent opportunities, but the underlying barriers that must be solved first, such as resistance to change, system access, data availability, or integration gaps.

#### Customer Example

A major ANZ bank was looking for value from agentic technology in sales and service teams.

We ran a cross-functional workshop to surface pain points, value pools, and opportunities for agentification. What became clear was that some of the most pressing blockers weren't new use cases at all, but foundational challenges: limited access to third-party systems, data not readily available, and inconsistent quality.

Once we layered in the additional value that agents could unlock if these basics were addressed, those core data and integration fixes moved higher on the priority list. The exercise reframed ideation: fixing fundamentals first became the fastest path to agent-driven value.

### 2 Target real business outcomes, designing in lagging measures to give confidence.

Anchor every idea to a measurable result, whether it is faster resolution times, reduced manual effort, or improved customer satisfaction. Use existing CX research, service logs, or employee feedback to pinpoint where pain is sharpest. Starting from clear business outcomes keeps ideation focused and ensures that your agent roadmap is always tied to tangible value, not technology for technology's sake.

#### Customer Example

A large APAC-wide recruitment company were looking at Agentforce to increase recruiter productivity.

As we explored the Ideate phase with them, not only could we increase recruiter productivity (by saving them time doing monotonous tasks) but we could create a new 24/7 always-on digital agent experience that could attract new candidates and drive up revenue.

By targeting recruiter productivity as the outcome, they saved 7,500 hours per week that were redirected into deeper candidate relationships. Not only that, the always on agent can engage candidates in new markets, broadening reach, and adding a revenue uplift value driver to the ROI.





3

### Identify internal talent & agree on executive ownership.

Successful AI projects have two things in common: a senior executive sponsor with the authority to commit resources and unblock issues, and a clear owner responsible for day-to-day optimisation. Decide early how you will resource each initiative - in-house, with a partner, or as a hybrid - and ensure there is a named leader accountable for outcomes.

If the capability isn't quite mature enough internally then involving a partner can achieve two parallel benefits: accelerating live agents in your business and uplifting internal capability.

4

### Accelerate the path to production from existing investments.

Agents perform best when they are built on consistent, well-structured, and trusted data. Before exploring new ideas, examine the assets you already have, such as knowledge articles, product catalogues, or customer records. As you prioritise you'll have a great sense of what can be capitalised on right now to improve speed to market.

5

### Solve the right problems, the right way, with the right technology.

Agentic AI isn't the solution for every part of your business. Some opportunities are better served by fixed-logic automation, while others need an agent's ability to reason with context. During ideation, **be explicit: if the task requires a predictable calculation or a deterministic rule, use automation.**

This distinction prevents wasted effort, avoids performance issues, and ensures you are applying the right tool to the right problem. Agentforce can blend together reasoning with fixed logic and deterministic flows, helping to deploy the right technique for the right problem.

#### Customer Example

A large ANZ travel organisation delivered their first Agentic Proof of Concept by blending predictive, generative and agentic capabilities.

Existing predictive models were leveraged for sentiment analysis, case classification and priority scoring. Generative AI was used for summarisation and Agentic AI supported the natural language agent experience and automated booking assistance.

The blend of these 3 existing and new technologies delivered a seamless end to end customer experience.



# Customer Agentblazer Example

## FISHER & PAYKEL

### Company Profile

Fisher & Paykel is a New Zealand based retailer of luxury home appliances, selling all across the world. Personalised and consistent customer service is a hallmark of their brand.

### Process

They recognised early that Agentforce can play many roles in improving customer experience. Some of those ideas were more complex than others; for example, resolving product queries for their approximately 10,000 SKU product catalogue versus connecting Agentforce to an appointment booking system. Fisher & Paykel decided to move initially with a simpler use-case that would have immediate impact to the customer service teams. All they needed to do was teach Agentforce about its products and connect Agentforce to its existing knowledge base articles.

### Learning & Impact

This is a great example of moving fast, learning, and then applying those learnings to more complex use-cases. Fisher & Paykel has now **doubled its digital self-serve resolution rate from 30% to 60%**, unlocking human capacity in the contact centre to deliver that world-class CX everyone expects from Fisher & Paykel.



You can read more about this story here



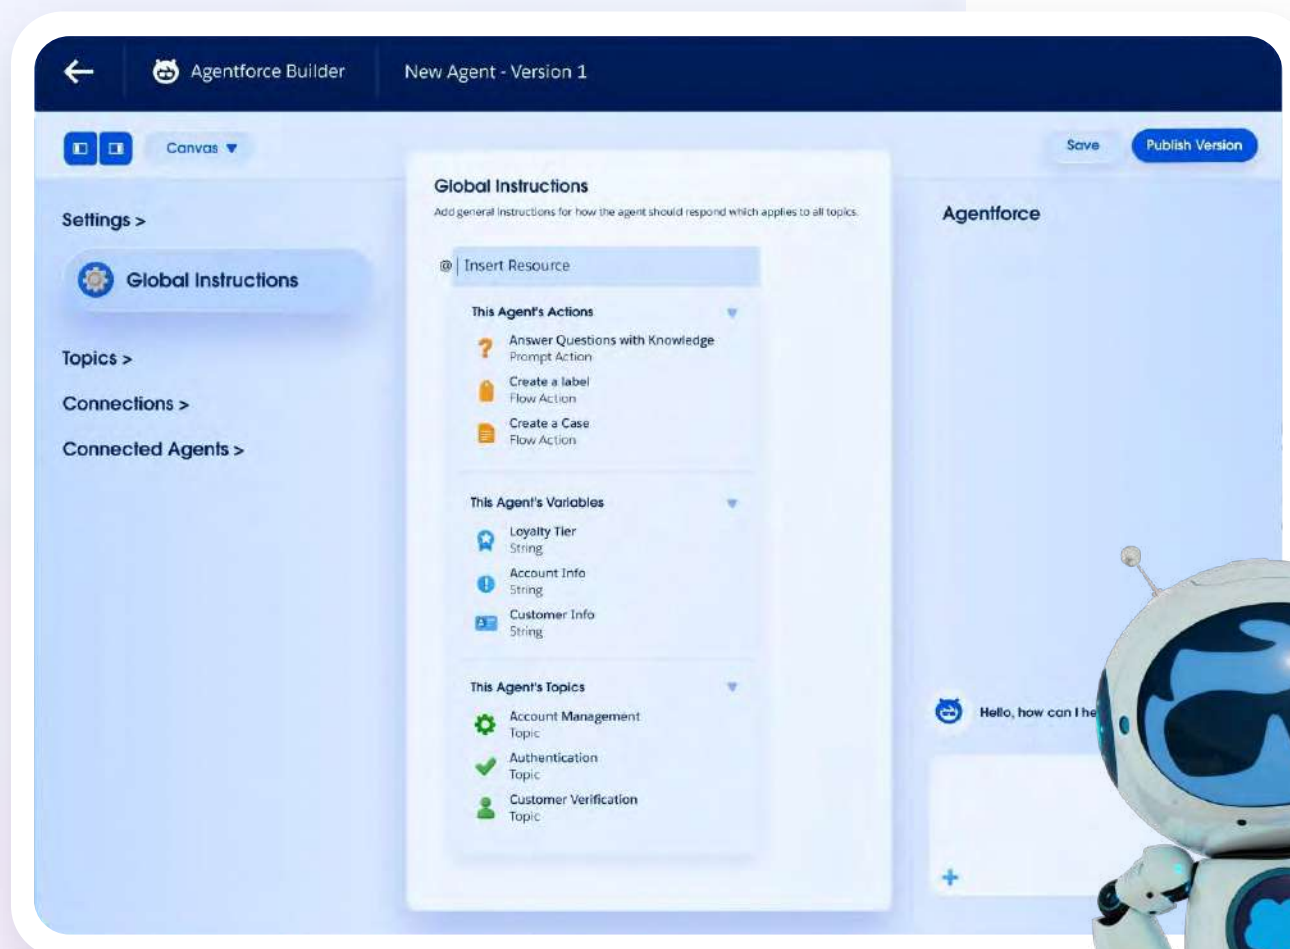
# 2 Build

Once you've got the use case, it's time to bring the agent to life. You'll configure its goals, the systems it connects to, and how it handles tasks – all while aligning it to your business logic, data, and trust requirements.




2

Now that you have your use cases defined, we move to the **Build** phase.



Best practices for agent design continues to evolve as the tooling get more intuitive and customers mature in their thinking around agent risk, performance, cost, and ownership.

Nevertheless, a key truth has held true over the last 12 months. The days of defining requirements and throwing it over the fence for I.T. to build are over. Building agents is a consultative process, where technology and the business teams work closely together to define the agent jobs to be done (JTBD), data sources, actions, guardrails and the user experience.





Let's explore the **5 lessons learned** from our customers during the **Build** phase:

**1** Plan for no, low, and pro code resources.

When pulling together your Agent delivery team it is useful to have a resource that is:

**No code** for strategy,

**Low code** who understands all operational components of your project, as well as a

**Pro code** resource who can help dive into the underlying code if and where required.

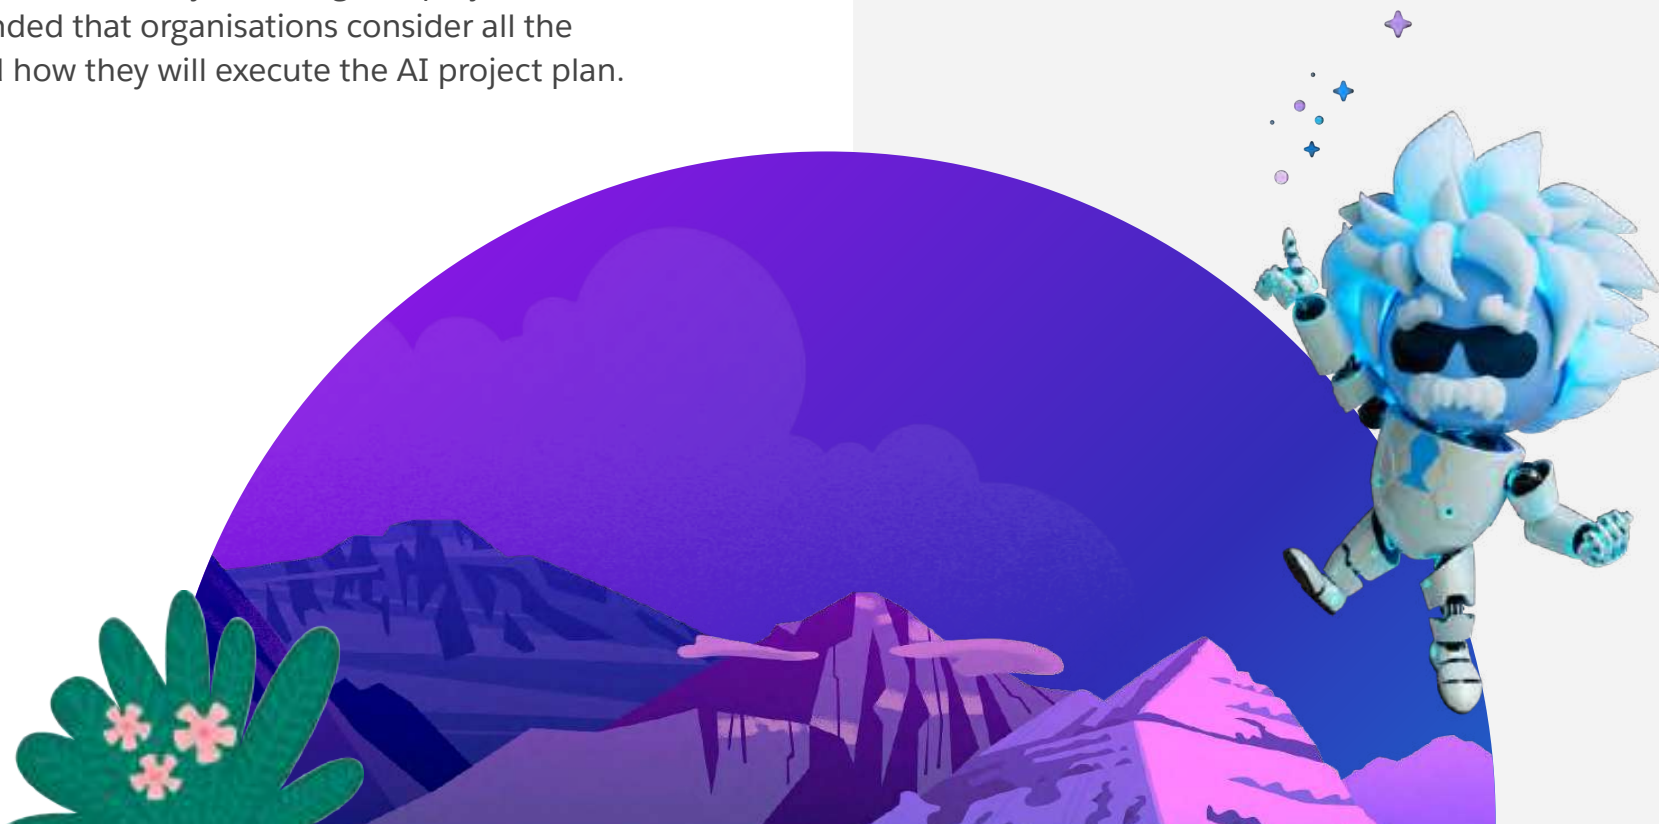
This fast tracks troubleshooting in your build and testing phase of deployment. While all three of the above components are not mandatory for a single deployment, it is heavily recommended that organisations consider all the user personas and how they will execute the AI project plan.

**2** Estimate consumption costs early in the process.

To set expectations with stakeholders, especially around budgeting - it's essential to map out how many expected queries and actions you wish your agent(s) to run.

Salesforce provides flexible methods to measure external/ internal/ operational use cases so they can be appropriately metered or bundled to avoid blowing out costs. Mapping this out before you build lessens the chances of being caught by surprise later on.

A common pitfall is scoping external conversations when the use case is high-volume internal conversations.





**3 Start small to move fast.**

Just like humans when starting a new job, agents should have a limited scope and require clear instructions in order to perform reliably. An all purpose general agent is fine for consumer AI, but in business, agents are operating in a deeper context including user permissions, business processes, and automated actions.

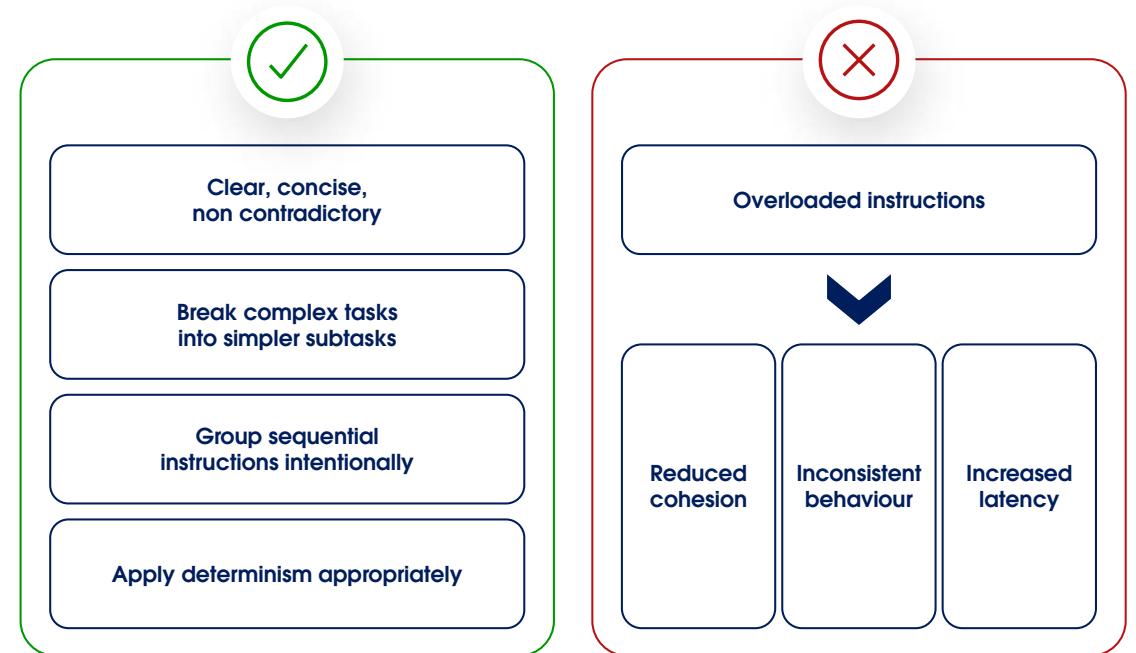
**Use precise instructions to ensure your agent performs reliably. This means keeping all instructions clear, concise, and simple.** To manage complexity, break down complex tasks into smaller, manageable subtasks, and group sequential instructions intentionally to ensure proper flow. Critically, you must avoid overloading the agent by providing too many instructions.

**Customer Example**

A **financial banking customer** significantly improved retrieval processing timeframes by truncating thousands of broad documents into concise, focused information required for their specific AI use case.



**Instruct your agent properly**





#### 4 Unstructured data needs your attention.

Customers have well established data governance practices in place for structured data (tables) however the quality of unstructured data has generally been neglected. In this world of agents, the quality of your source documents directly impacts agent accuracy. To achieve this, you need to ensure your content is:

**Thorough** - explains concepts thoroughly, provide examples, and structure content logically.

**Consistent** - follows company guidelines and a strict process to maintain consistency.

**Focused** - content that directly answers customer questions and is focused on the agent's core function.

##### Customer Example

Many of our **customers in ANZ** are turning to the power of Generative AI to help optimise their knowledge content and get it prepared for agentic use.

One **ANZ retailer** discovered that by simply standardising headers in their unstructured PDFs, Agentforce could retrieve information with far greater accuracy, turning a basic content clean-up into a powerful accelerator of business value.

#### 5 Business + IT building together is the winning formula.

Engaging with a non-deterministic & semi-autonomous agent is a big change, and users should understand how agentic AI works, how the agent has been built, and what it can do long before it moves into production. Customers that have brought the users along for the ride, from ideation through to deployment, have reported much higher rates of speed and adoption than those that have delayed user training.

##### Customer Example

One of our major **Telco customers in ANZ** built and deployed a customer facing agent that allows them to review and update their plans in less than 5 weeks time. While the speed itself is impressive, the reason they were able to move at such pace was due to the significant investment they had made prior in both their business and technical teams to understand agentic AI.

This went beyond executive sponsorship and training, they integrated their technical teams and change management specialists into a single delivery squad, who worked together in parallel rather than sequentially. This understanding that user adoption must be baked into your deployment process led to a highly successful agent go-live.



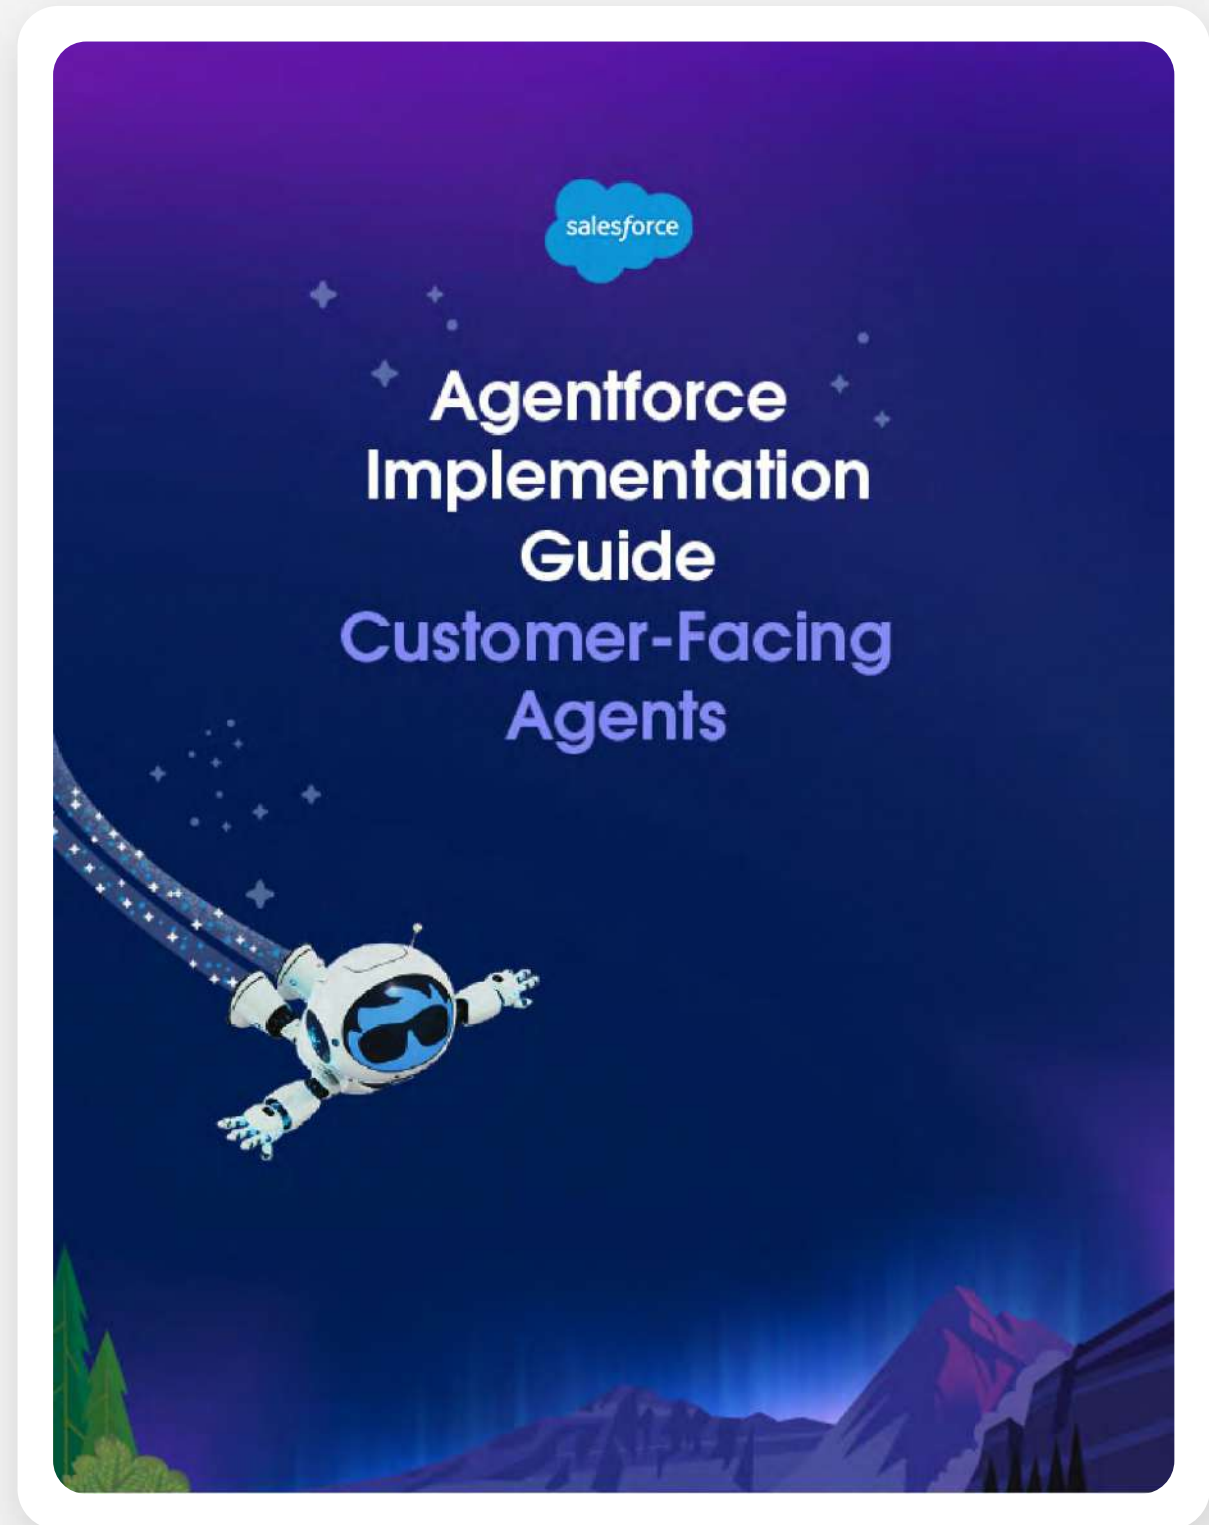



### A Key Resource for Builders:

#### Introducing the Agentforce Guide

We've aggregated the most actionable best practices and lessons in our:

 [Agentforce Implementation Guide](#)



# 3

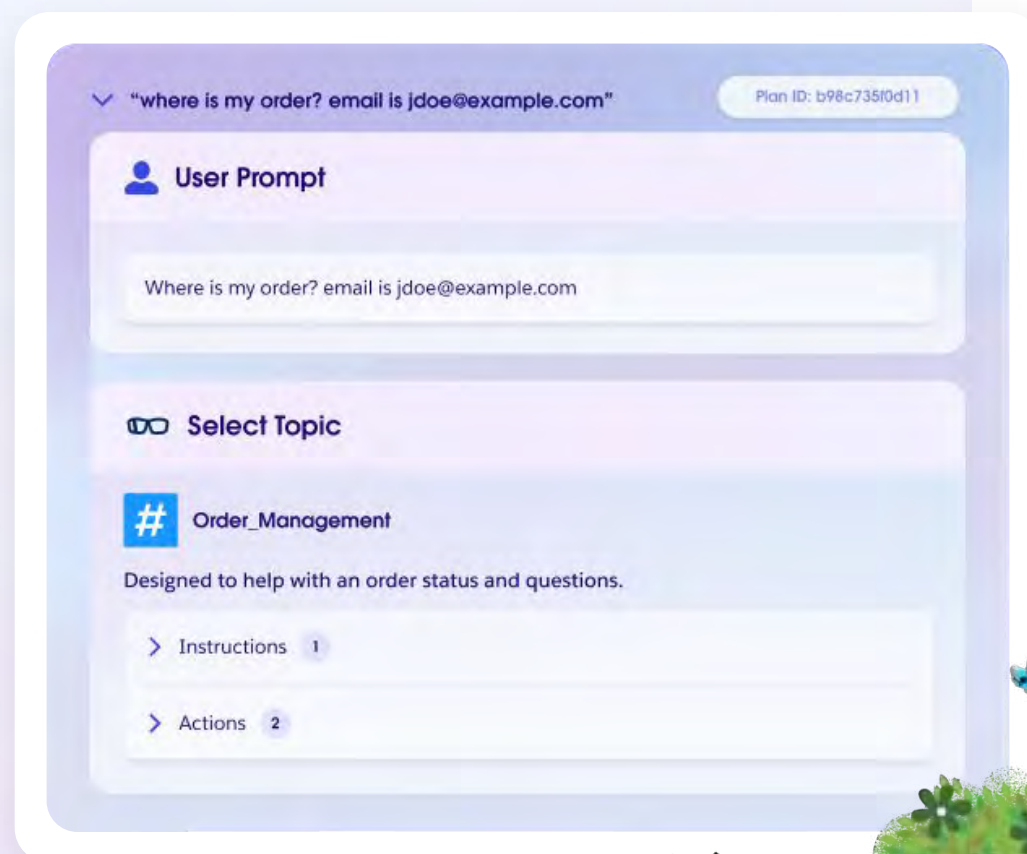
# Test

Before the agent sees the light of day, you put it through its paces. This is where you validate performance, simulate scenarios, get feedback, and make sure it's safe, accurate, and effective



3


With your configurations from the **Build** phase complete, we now shift focus to the **pre-production Testing** phase.



Testing agentic AI is fundamentally different from traditional automation. Rather than testing for explicit outcomes, we're looking to test for consistency in behaviour and establish acceptable thresholds of performance.

In order to properly test your agents, there are a few factors to consider. Firstly, test small, and scale up gradually as you see performance improve, and data points processed as expected. This can ensure you trial Agentforce at a pace that is relevant to your team's resources and expertise levels.

From there, you can continue to test levels of accuracy in relation to the agent responses within your pre-production environment.





Let's explore the **5 lessons learned** from our customers during the **Test** phase:

## 1 Create the right testing scenarios.

A good starting point is to create testing scenarios on what you want your agent to do, and what you don't want your agent to do - and optimise towards the former.

- ✔ Identify inaccurate intent classification, incomplete or incorrect action sequences, or unexpected agent responses.
- ✔ Salesforce's testing center uses Assertion based testing and LLM as a judge based AI assistive testing to help evaluate these metrics.

For our more agentially mature customers, the most common testing scenarios include:

### Topic classification issues

(e.g. triggers incorrect topics)

### Action execution problems

(e.g. calls wrong actions or unexpected actions)

### Knowledge retrieval issues

(e.g. missing knowledge or inclusion of unnecessary information)

### Response quality issues

(ungrounded information, not sourced)

### Guardrails and instruction issues

(e.g. disregards operational instructions), or

### Returns a response that lacks structured step-by step guidance

(generic advice, lost context, or repeated invalid responses).

New Test													
Basic Information	<b>Evaluations</b>												
Test Conditions	Choose the evaluations you want to perform for this batch test.												
Test Data	<b>Default Evaluations</b>												
<b>Evaluations</b>	<table border="1"><thead><tr><th>Evaluation</th><th>Description</th></tr></thead><tbody><tr><td>Response Evaluation</td><td>Evaluates if the agent's response matches the desired response.</td></tr><tr><td>Topic Assertion</td><td>Assesses the agent's ability to select the correct topic.</td></tr><tr><td>Action Assertion</td><td>Assesses the agent's ability to select the correct actions.</td></tr></tbody></table>	Evaluation	Description	Response Evaluation	Evaluates if the agent's response matches the desired response.	Topic Assertion	Assesses the agent's ability to select the correct topic.	Action Assertion	Assesses the agent's ability to select the correct actions.				
Evaluation	Description												
Response Evaluation	Evaluates if the agent's response matches the desired response.												
Topic Assertion	Assesses the agent's ability to select the correct topic.												
Action Assertion	Assesses the agent's ability to select the correct actions.												
	<b>Response Quality Evaluations</b>												
	<table border="1"><thead><tr><th><input checked="" type="checkbox"/> Evaluation</th><th>Description</th></tr></thead><tbody><tr><td><input checked="" type="checkbox"/> Completeness</td><td>Evaluates if the response includes all necessary information.</td></tr><tr><td><input checked="" type="checkbox"/> Coherence</td><td>Evaluates if the response is easy to read and free of grammatical errors.</td></tr><tr><td><input checked="" type="checkbox"/> Conciseness</td><td>Evaluates if the response is short but accurate.</td></tr><tr><td><input checked="" type="checkbox"/> Latency</td><td>Tests the length of time, in milliseconds, it takes for the agent to generate a response.</td></tr><tr><td><input checked="" type="checkbox"/> Instruction Adherence</td><td>Evaluates if agents adhere to the provided topic instructions.</td></tr></tbody></table>	<input checked="" type="checkbox"/> Evaluation	Description	<input checked="" type="checkbox"/> Completeness	Evaluates if the response includes all necessary information.	<input checked="" type="checkbox"/> Coherence	Evaluates if the response is easy to read and free of grammatical errors.	<input checked="" type="checkbox"/> Conciseness	Evaluates if the response is short but accurate.	<input checked="" type="checkbox"/> Latency	Tests the length of time, in milliseconds, it takes for the agent to generate a response.	<input checked="" type="checkbox"/> Instruction Adherence	Evaluates if agents adhere to the provided topic instructions.
<input checked="" type="checkbox"/> Evaluation	Description												
<input checked="" type="checkbox"/> Completeness	Evaluates if the response includes all necessary information.												
<input checked="" type="checkbox"/> Coherence	Evaluates if the response is easy to read and free of grammatical errors.												
<input checked="" type="checkbox"/> Conciseness	Evaluates if the response is short but accurate.												
<input checked="" type="checkbox"/> Latency	Tests the length of time, in milliseconds, it takes for the agent to generate a response.												
<input checked="" type="checkbox"/> Instruction Adherence	Evaluates if agents adhere to the provided topic instructions.												
	<input type="button" value="Cancel"/> <input type="button" value="Generate Test Cases"/>												



**Salesforce Testing Center:** a feature available as part of Agentforce, which allows end users to perform several evaluations, including the ability to test their Agentforce results to monitor agent quality, troubleshoot utterances, and fine-tune agents.



## 2 Always test in pre-production environments.

Always make sure to test your agents in pre-production environments before publishing for wider internal or external use (Agentforce's testing center is deliberately available in pre-production stage).

### Customer Example

An ANZ higher education customer, retrospectively try to determine why they hit overages across a certain category, 90 days into an implementation. This could have been caught and investigated within the first week of user activity if they had been monitoring their consumption cards within the **digital wallet** as part of the testing phase.



Digital Wallet is a free account management tool that offers near real-time usage data for your consumption-based Salesforce products

## 3 Non determinism requires different testing criteria.

It's commonly understood that using LLM's as part of the Agentic experience means that answers will vary, even if the same question is asked. When running user acceptance testing for your agent, make sure that your success criteria is based on the ability of the agent to understand the intent, follow the defined instructions and guardrails and take the right action. It is important to set success criteria that is not too rigid and set the expectation to users that they can always ask a follow up question if they require additional information from the agent.

### Customer Example

One ANZ customer classified their agent testing as unsuccessful. Upon further review, it was discovered that they had explicit and narrow testing criteria, meaning that Agent responses were considered a failure if they did not provide an exact match to a predefined response. Establishing criteria that takes into consideration the non-deterministic nature of agentic AI will provide a fairer assessment of their capabilities.

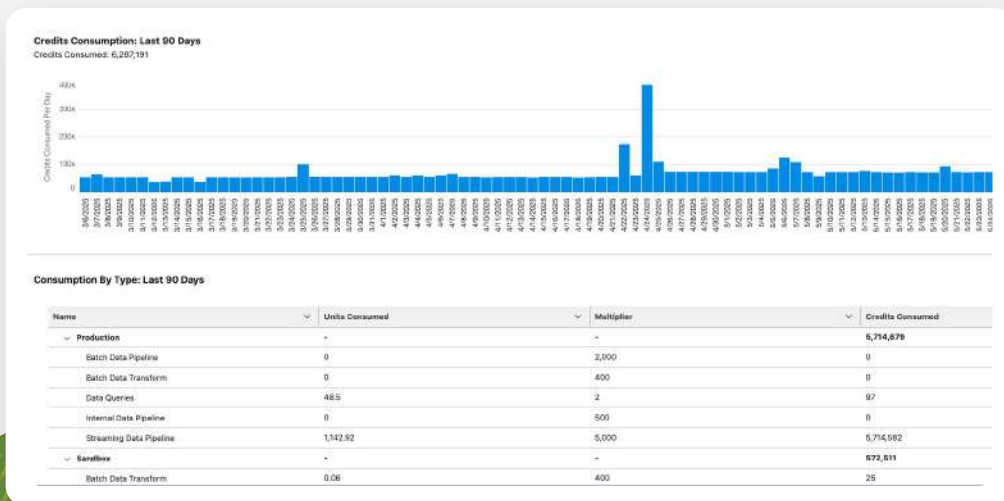




**4 Monitor Digital Wallet to help estimate production credits.**

Monitoring Digital Wallet is a useful way to forecast production costs and credit estimation. This will provide visibility into a baseline cost across various scenarios. Once you have this benchmark, you will be able to compare what the most complex, highest consumption use cases are.

Tracking consumption in Digital Wallet during testing will also help spot any potential issues during the testing phase. For example, if during routine testing, you see a spike in consumption correlating with the agent giving inaccurate responses or pulling incorrect information, there is likely a configuration issue within the agent instructions. Monitoring consumption trends can help pinpoint any issues to address during testing.



**5 Ensure instruction configuration and data sources have followed consistent governance practices.**

Inconsistent instructions and configuration is a common cause of issues with agent responses. You want to ensure that your instruction configurations and data sources are adhering to consistent governance practices. Instruction configuration should be clear, unambiguous, and - most importantly - consistently applied across different scenarios.

It is also important to ensure your data sources have an established governance practice that is being followed for consistency. Remember this applies to all the data sources the agent might be referencing including databases, APIs, and knowledge bases.



# 4

## Deploy

It's go time! The agent is live and working in the real world – helping customers, employees, or partners. This stage is all about ensuring adoption, monitoring usage, and making sure it's delivering value from day one.



4

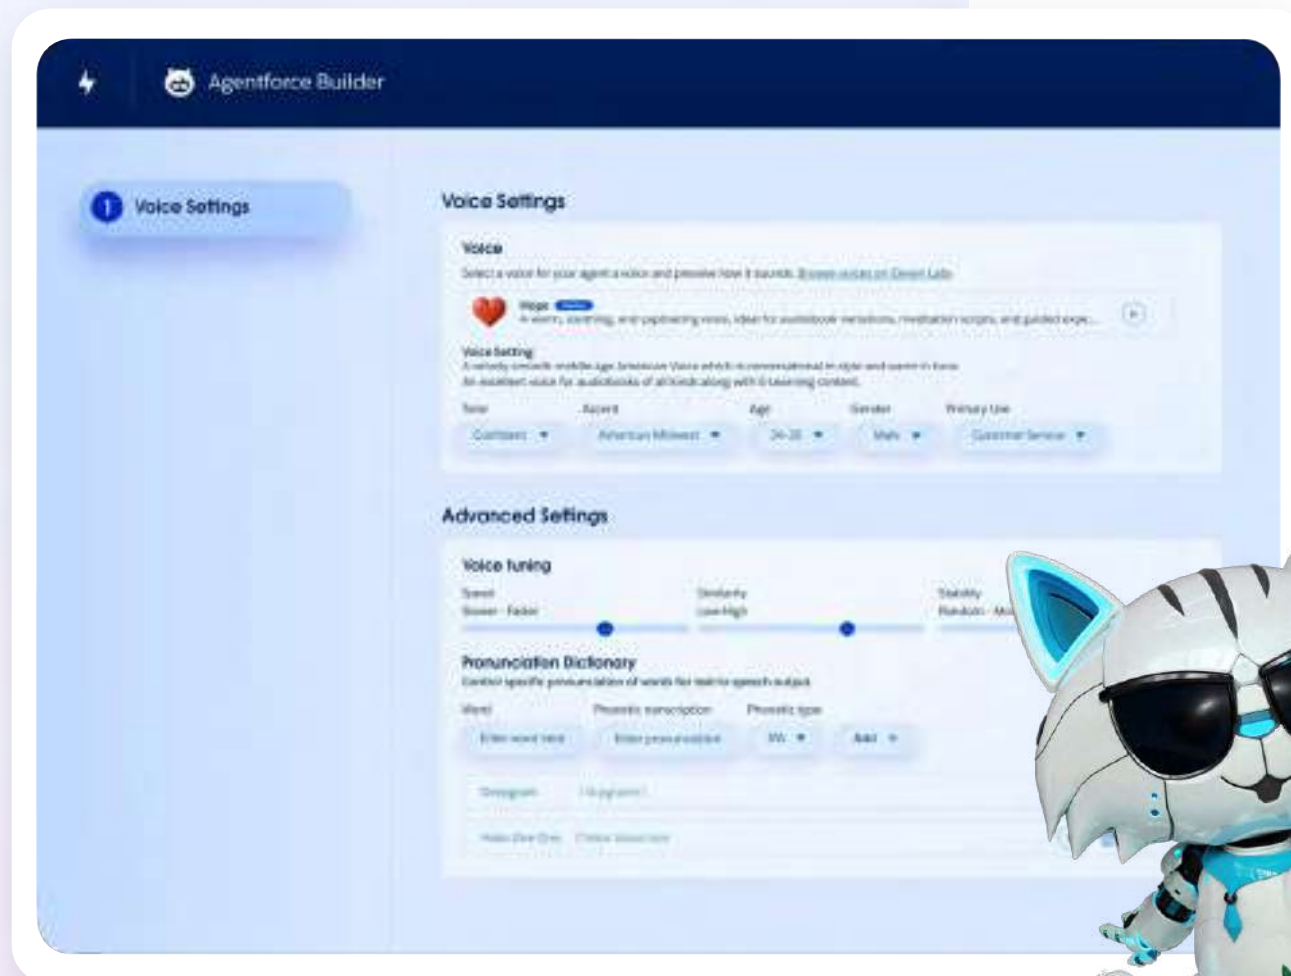
# Congratulations, you've joined our rapidly growing global cohort of Agentblazers.

Amazing, your agent is live in Production.

But the work doesn't stop here because an agent is only as effective as the trust it holds within your organisation. This is the moment where agents and humans need to work together, and for many of our customers, the success of this first deployment will have ongoing ramifications for further investment in agentic AI.

In the deploy phase it is absolutely essential that users feel informed, heard and supported in what can be a tectonic shift in how they engage with their work. Issues need to be resolved quickly and robust business continuity processes (BCP) should be adhered to.

Channel strategy is also a critical component of your agent deployment, ensuring it is surfaced in the right channels (web, email, voice, messaging) with a clear path of escalation.





Let's explore the **5 lessons learned** from our customers during the **Deploy** phase:

### 1 Keep your agent in a state of hyper-care for the first couple of weeks.

Even though you've rigorously tested your agent pre-production, the reality with agentic AI is that something will come up that you didn't plan. At this stage, in order to build user confidence and adoption, it's critical to keep your agent in hypercare. This means daily use of your observability toolkit, having a clear and communicated method for users to report feedback or raise errors and making rapid (micro-) adjustments to quickly resolve issues.

#### Customer Example

A recent **internal employee agent release within the Salesforce Slack instance** had a structured hyper-care period of two weeks. During this period our teams performed daily performance checks for latency and response quality, issue tracking with clear prioritisation categories (critical, moderate and minor).

Support was facilitated by a dedicated Slack channel for rapid issue reporting and a clear escalation process.

This proved highly effective, resulting in seven issues being resolved and closed (one critical, one moderate and five minor) before transitioning to steady state support.

### 2 Test user permissions with a small group in Production before rolling out further.

These settings dictate a user's ability to access the agent and define how agent usage is metered. While these requirements are fully documented on [help.salesforce.com](https://help.salesforce.com), customers frequently overlook this step in their pre-deployment checklist, resulting in frustrating deployment errors. To mitigate this risk, always validate access and consumption by testing with a small, contained subset of users before announcing the effective 'go-live.'

This validation is also crucial for preventing unexpected metering surges.

#### Customer Example

A **customer in NZ** recently deployed an internal sales agent and was surprised by a significant spike in their Einstein Requests consumption, despite having purchased user licenses for unmetered internal usage.

Upon review, they discovered the users were not assigned the Unmetered User-Based AI permission set, which authorises the use of certain AI features without consuming credits.

This pre-check is now a mandatory step in their release management process to avoid future challenges.





3

Ensure you have escalation processes in place so your customers aren't stranded if something goes wrong.

Many of our customers have deployed service agents on their websites. Agents that are capable of answering general FAQ's, retrieving customer information and triggering actions like submitting a refund or generating a lead.

In circumstances where the agent can't understand the customer request or it encounters an unexpected error, best practice recommends that it is escalated to a human agent for resolution. Customers can configure the standard 'escalation' topic in Agent Builder as well as design their flows to allow the agent to raise a case in the situation where there are no human agents available.

### Customer Example

One of our **global customers** deployed a service agent on their website and found that 5% of conversations were being dropped without escalating to a human for further review.

Upon further investigation it was clear that the escalation logic was incorrect and that customer queries were routed to a queue without any human agents assigned to monitor it.

This can be avoided with thorough escalation testing pre-deployment as well as using the escalation fallback feature which ensures there are no abandoned chats.

The image displays two overlapping screenshots from the Agentforce interface. The top screenshot is the 'Voice Settings' page, which allows users to customize the voice for their agent. It features a 'Basic' section with a 'Voice' dropdown set to 'Ethan', a description, and options for Gender (Male), Tone (Confident), Accent (American), and Age (26). Below this is a 'Preview Message' field and a 'Listen' button. The 'Advanced' section includes 'Voice Tuning' sliders for Speed, Similarity, and Stability, and a 'Pronunciation Dictionary' for defining specific words and phrases.

The bottom screenshot is a 'Voice Preview' chat window. It shows a conversation where the AI agent, 'Ethan', introduces itself: 'Hello, I'm Agenforce. I'm an AI Agent for Electra. I specialize in customer support. How can I help you today?'. A customer asks, 'Which model in the Sybil series would you recommend?'. The agent responds: 'Great question! The Sybil series includes a few models, each suited to different needs. Are you looking for something that's best for everyday use, high performance, or advanced features?'



4

### Gain stakeholder trust with internal deployments first.

If internal stakeholders or sign off remains hesitant, consider deploying with employee agents and internal use cases first. Focusing on smaller, familiar use cases internally will build trust and confidence in how an agent can help the business and will feel like a less risky use case to start with for stakeholders who may be concerned with launching directly with a customer facing use case.

It is also a useful way to get real world, company specific feedback and data to help mold your external use cases. Starting with internal use cases that employees will be the primary users of could help build more internal champions and support as well.

#### Customer Example

An **ANZ customer** recently deployed their first Agentforce use case internally to review call transcripts for automated quality assurance.

The journey from ideation to deployment took just 2 months, and they're already experiencing significant value with the ability to run quality checks on a much higher percentage of their call volumes than ever before.

This is a great example of delivering an internal high value use case to build confidence in agentic AI before deploying further.

5

### Focus on intuitive user experiences.

For customised user intuitive experiences, consider deploying with rich formatting and multi modal options. Presenting the agent's responses in a visual, easy to understand format will support usability and adoption.

For example, instead of the agent presenting back lines of text that the user needs to read and analyse, presenting action buttons or cards can prompt a user to immediately act on the agent's suggestion or response. This will help make the agent experience more interactive for the end user.

#### Customer Example

A **leading Australian retailer** has transformed their customer experience by deploying Agentforce across their website with 'adaptive response formats' enabled, allowing them to deliver the rich visual interactions that today's customers expect.

When browsing their extensive catalog of electrical appliances, customers can now engage with an agent that presents product information through interactive carousels rather than static text responses.

This approach ensures there's no degradation in experience between their website and the agent interface.



# 5

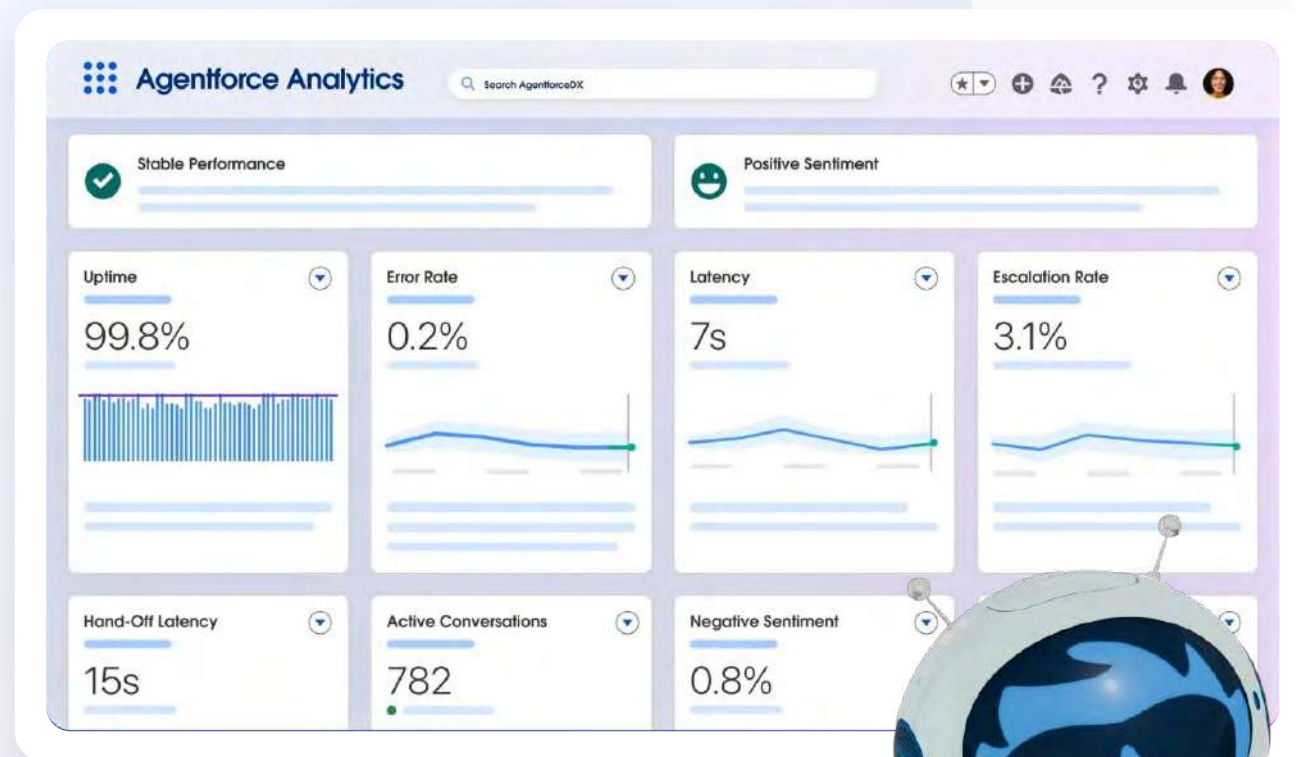
# Optimise

Even after launch, the work continues. You'll monitor how the agent performs, learn from feedback, and keep tuning it to be faster, smarter, and more impactful.



5

# Your agent is live and delivering great value to users.



A few tweaks are required here and there, but by and large your agent is performing as expected.

At this point, many customers start to think about their next agent, but just like humans, agents have the ability to get better at their current job and develop new skills. This is an opportunity to consider how your current agent can deliver even more impact and value to your users.

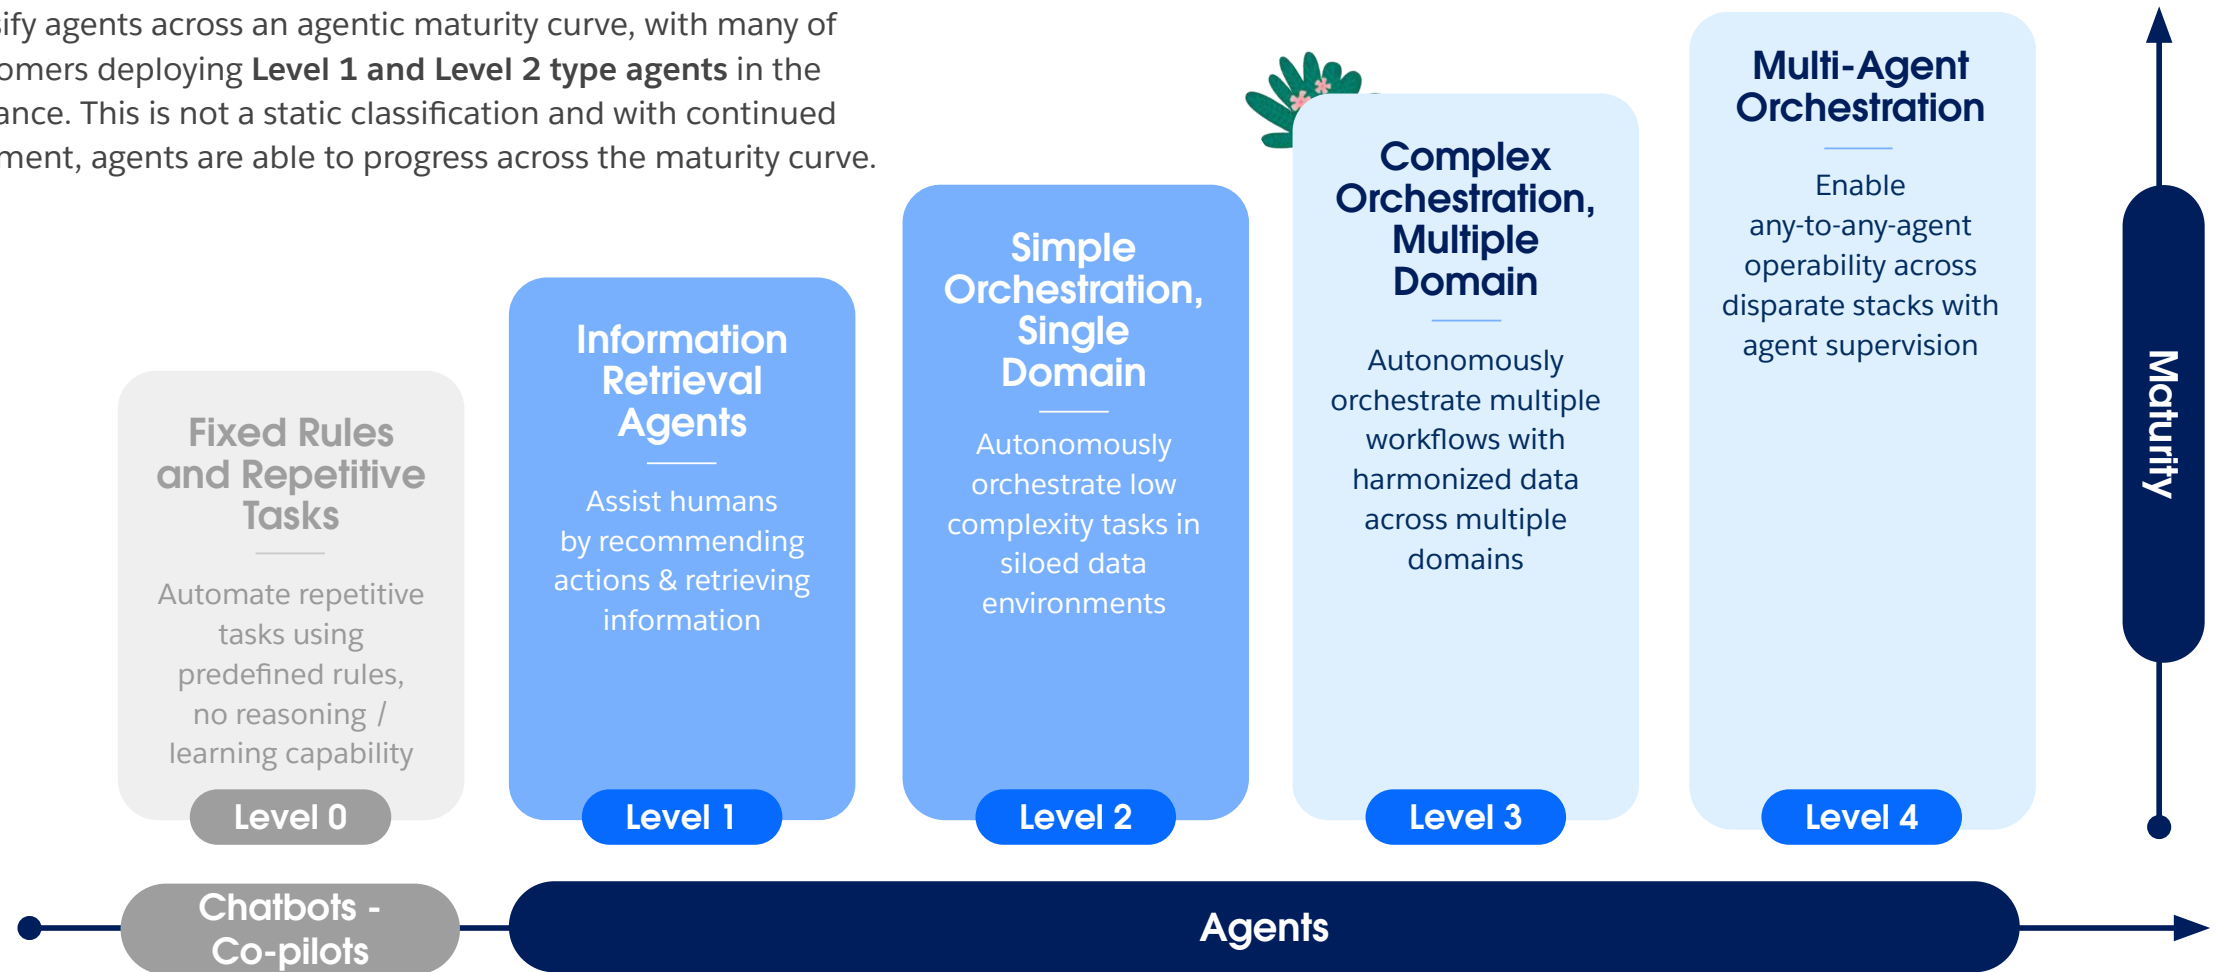




Let's explore the **3 lessons learned** from our customers during the **Optimise** phase:

**1** It's not always about the next agent. Sometimes it's about the next skill for your current agent.

We classify agents across an agentic maturity curve, with many of our customers deploying **Level 1 and Level 2 type agents** in the first instance. This is not a static classification and with continued development, agents are able to progress across the maturity curve.



The agent experience provided on <https://help.salesforce.com/> has been a continuous process of fine tuning which has led to significantly improved agent outcomes.

Transitioning the agent customer experience to a fullscreen experience saw a 22% increase in adoption rates. The agent progressed from a Level 1 agent to Level 2 with the integration of case creation, resulting in a 17% decrease in abandonment rates. Over time, the agent has been optimised to include in-line citations, consolidated data entry for case creation, and multi-language support.

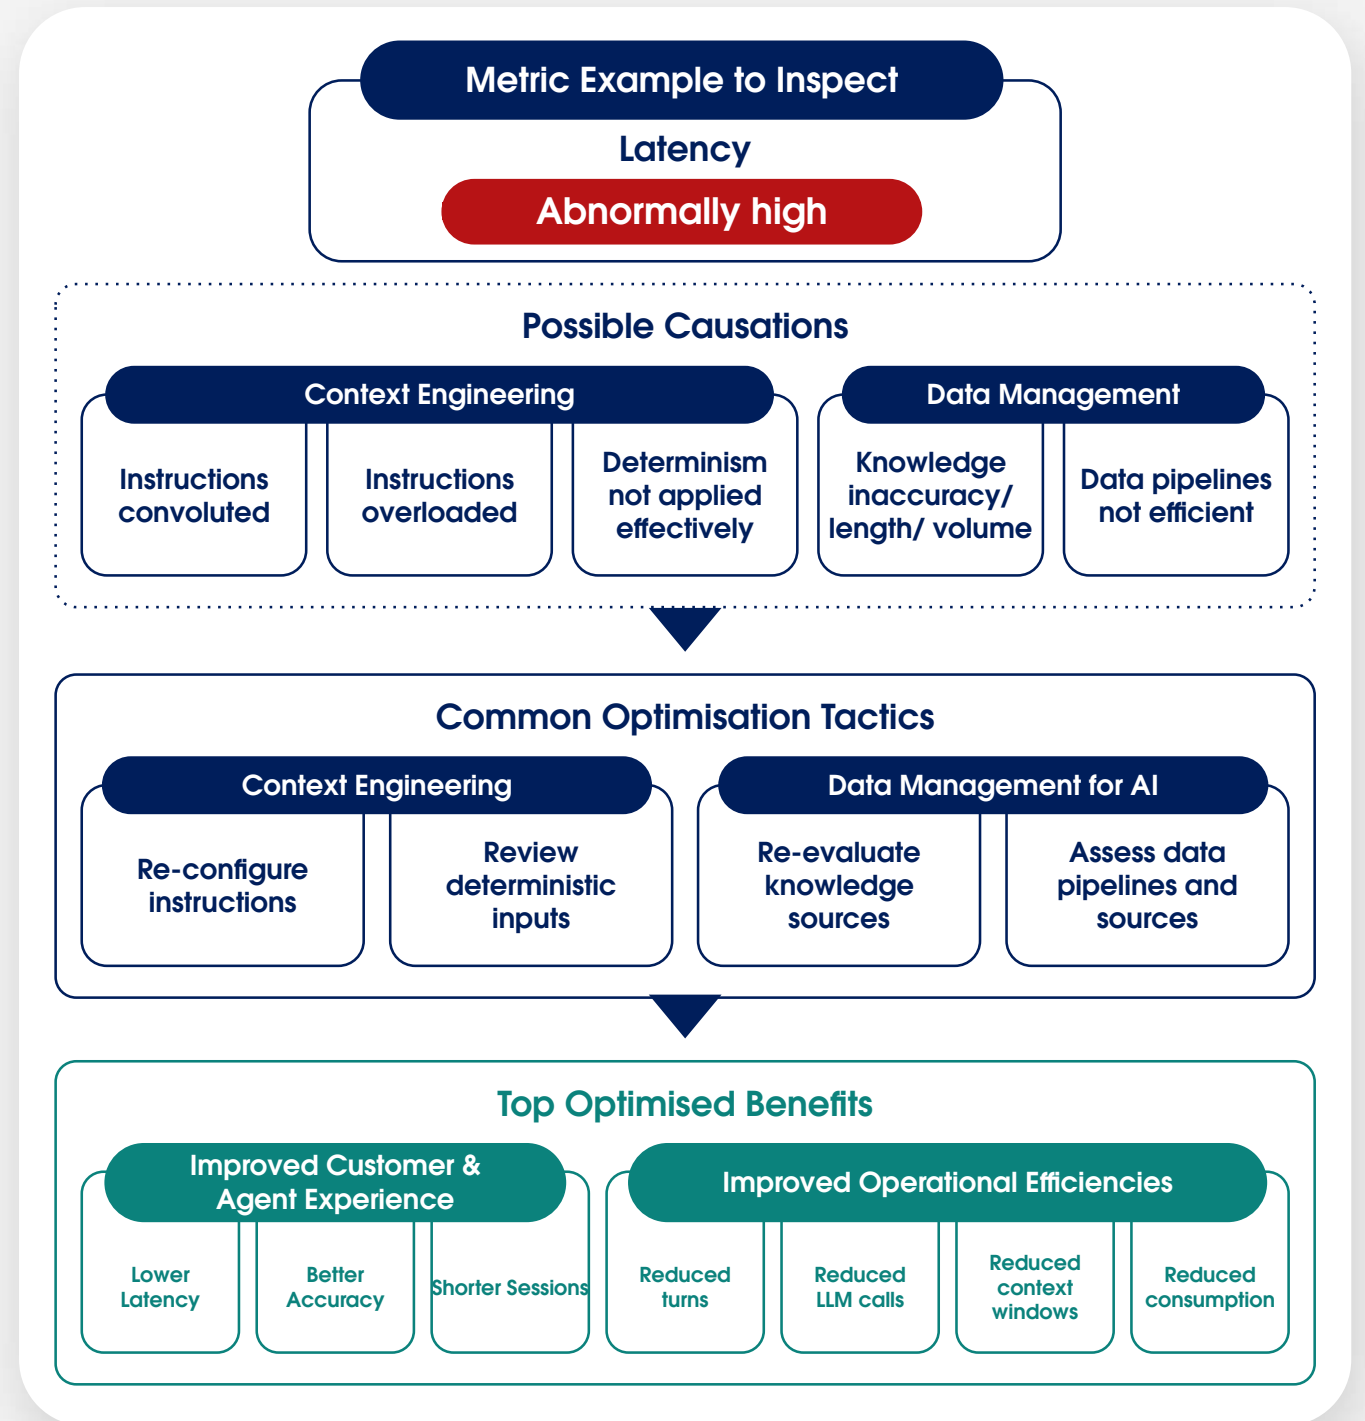


## 2 Agent latency is a joint responsibility.

At Salesforce, as an Agentic platform we have a responsibility to ensure that our infrastructure is running at low latency to allow real time engagement. But latency is also the responsibility of our customers in how they design their agents.

Agent deployment latency can be traced to inefficient configuration, where incorrectly created deterministic tasks written in natural language create overloaded LLM requests.

By optimising the agent setup to include deterministic variables and logic, our teams immediately saw a reduction in LLM calls, improved response accuracy, and shorter, more efficient sessions (fewer turns). This refinement streamlined instructions and reduced context window size, ensuring faster, more predictable outcomes.





3

### Use trace session data to refine your agents and develop your agent roadmap.

Using Agentforce's Observability dashboards, you can view how users are interacting with the agent, and the types of questions they are asking. Customers are leveraging this data to optimise agents by adding these 'utterances' into Agent Builder which is helping deliver significantly better agent responses. Based on the types of questions asked of the agent by customers and employees, there is an opportunity to consider your next agent 'skill' or next agentic use case.

The screenshot displays the Agentforce Studio interface. At the top, there's a search bar and navigation tabs for 'Agentforce Studio' and 'Insights'. Below this, a 'Moment' is selected with ID '7GBHt00000KF3bPBG'. The main area is divided into two sections: 'Session Log' on the left and 'Question regarding cancellation policy' on the right. The 'Session Log' shows a chat conversation with a user and an agent. The 'Question regarding cancellation policy' section includes a table with the following data:

Agent Name	Total Interactions	Average Agent Latency
Car Rental Service Agent	2 interactions	134 ms

Additional details include Moment Start Time (01/31/2025, 11:34 AM), Moment End Time (01/31/2025, 11:36 AM), Moment Elapsed Time (00:24), Moment Tags (Cancellation Request, Unresolved, Order Management), and Actions Triggerged (Retrieve Knowledge). A Quality Score of 'Low' is also displayed.

### Customer Example

One customer saw that 80% of their service agent intake was for billing purposes. They allowed them to create a pipeline of agents to build, including a billing enablement agent for internal staff, a billing specific agent for external use, and an FAQ agent that specialised in billing for new starters.

4

### Flexible change management processes means faster optimisation.

Customers that have flexibility built into their release management processes are able to make these rapid adjustments much faster than those that continue to follow rigid guidelines (i.e. pushing minor changes through DEV, SIT and UAT environments first).

### Customer Example

One such example is an ANZ customer that wanted to modify their agent to respond in bullet points for a particular topic. All this change required was an additional instruction in Agent Builder, and was thus considered low risk and easy to reverse. Admins should be empowered to make minor agent changes in Production during the hypercare phase, with higher risk changes to follow the standard change management practices.

5

### Keep across the latest Salesforce release notes.

Innovation is coming at a rapid pace on the Salesforce platform, and capabilities that might not have been available to your agent during the build and test phase might now be ready to turn on. For example, many of our customers that built information retrieval agents at the beginning of the year largely relied on PDF documents as their unstructured data sources. Support for additional file formats including csv, excel, powerpoint, and email have been added in recent releases (as well as a number of connectors) making it easy to ground your agent with additional unstructured data and increase the depth of its knowledge.



You can always see the latest release notes on the Salesforce Help site

# Conclusion and closing thoughts



# From Idea to Impact



The nature of work is changing. The **Agentic Enterprise**, where human employees and AI agents work together, is fast becoming a reality. The result unlocks new opportunities for productivity, customer connection, and growth.

Across thousands of Agentforce deployments and Salesforce itself as customer zero, the **structured Agent Development Lifecycle** has proven essential:



Ideate



Build



Test



Deploy



Optimise


This framework moves organisations efficiently from an initial concept to a high-impact live agent. The most impactful agents will be the ones that address real, clearly defined business problems and are built on the fundamentals, specifically the foundational data underlying the agent.

The work does not stop once the agent is live. There must be a considered process to continuously improve agents including ongoing monitoring to adapt the agent based on qualitative feedback from users and quantitative data about the agent's overall performance.

The long term success of any AI project hinges on the involvement of executive sponsors and the business process owners, to validate the agent's effectiveness and ensure maximum business value.

**By embracing this framework, you can move beyond simply adopting new technology to truly transforming into an Agentic Enterprise.**

# Summary of Lessons Learned



## Ideate

1

Talk to the humans behind the change, they know where the opportunities are.

2

Target real business outcomes, design in leading measures to give early confidence.

3

Identify internal talent & agree on executive ownership.

4

Accelerate the path to value from existing investments.

5

Solve the right problems, the right way, with the right technology.



## Build

1

Plan for no, low and pro code resources.

2

Estimate consumption costs early in the process.

3

Start small to move fast.

4

Unstructured data needs your attention.

5

Business + IT building together is the winning formula.



## Test

1

Create the right testing scenarios.

2

Always test in pre-production environments.

3

Non determinism requires different testing criteria.

4

Monitor Digital Wallet to help estimate production credits.

5

Ensure instruction configuration and data sources have followed consistent governance practices.



## Deploy

1

Keep your agent in a state of hyper-care for the first couple of weeks.

2

Test user permissions with a small group in production before rolling out further.

3

Ensure you have escalation processes in place so your customers aren't stranded if something goes wrong.

4

Gain stakeholder trust with internal deployments first.

5

Focus on intuitive user experiences.



## Optimise

1

It's not always about the next agent. Sometimes it's about the next skill for your current agent.

2

Agent latency is a joint responsibility.

3

Use trace session data to refine your agents and develop your agent roadmap.

4

Flexible change management processes means faster optimisation.

5

Keep across the latest Salesforce release notes.





The information provided in this report is strictly for the convenience of our customers and is for general informational purposes only. Publication by Salesforce, Inc. does not constitute an endorsement. Salesforce.com does not warrant the accuracy or completeness of any information, text, graphics, links, or other items contained within this guide. Salesforce.com does not guarantee you will achieve any specific results if you follow any advice in the report. It may be advisable for you to consult with a professional such as a lawyer, accountant, architect, business advisor, or professional engineer to get specific advice that applies to your specific situation.