

Softcat Plc

Salesforce Agentforce

- Cloud Software - Digital Marketplace - Lot 2b
- G-Cloud 15



CONTACT INFO

Dominic Clark

pstiq@softcat.com

Table of Contents

Salesforce Agentforce	1
1. Service Description	2
1.1. What is Agentforce?	2
1.2. Service Overview	2
1.3. Service Features	4
1.4. Service Benefits	16
1.5. Main components & Functions of this service	23
2. Information Assurance	26
3. Data Backup / Restore and Disaster Recovery	27
4. On-Boarding	29
4.1. Deployment	29
4.2. Configuration and Customisation	29
5. Off-Boarding	29
5.1. Termination	29
5.2. Data Extraction and removal	29
6. Pricing Summary	30
7. Service Constraints	30
8. Service Management	31
9. Customer Responsibilities	33
10. Client-side Technical Requirements	34
11. Planned Maintenance Windows	34
12. Training	35
13. Data Centre Locations	35
14. Performance	35
14.1. Service Levels	36
14.2. Incident Response Time	36
14.3. Incident Updates	36

Salesforce Agentforce

1. Service Description

1.1. What is Agentforce?

Agentforce is a comprehensive suite of tools designed to create and customise AI agents, as well as a collection of autonomous AI agents integrated across the Salesforce Customer 360 platform. An AI agent is a software entity that performs tasks on behalf of a user or another program using artificial intelligence techniques, such as processing natural language, learning from data, and making decisions based on predefined rules or learned patterns. These agents typically assist with specific tasks like customer service, data analysis, or recommendation systems. In contrast, an autonomous agent operates independently without human intervention, capable of making decisions and taking actions on its own based on its programming and the environment it operates in. These agents are frequently utilised in intricate systems where real-time decision-making is essential, such as in public safety operations, emergency response coordination, and certain types of automated public service systems. Agentforce integrates numerous advancements in workflow automation, data management, and generative artificial intelligence to enhance efficiency and effectiveness in these critical areas.

Autonomous AI agents within Agentforce are set to transform customer service for Government, Non-Profit, and Educational institutions. The Agentforce Service Agent, a next-generation AI agent for service channels, is powered entirely by large language models (LLM) rather than traditional declarative dialogs. This autonomous conversational AI assistant is always available, empathetic, and tailored to specific guidelines, facilitating effortless service and higher case resolution through dynamic, proactive conversational responses that both inform and initiate action.

These autonomous AI agents enhance employee productivity by handling tasks in service, sales, marketing, and commerce, thereby driving exceptional efficiency and customer satisfaction. They improve service experiences across all channels, 24/7, with natural, data-grounded responses, enabling swift and accurate case resolution. These agents operate independently, providing real-time assistance and automating routine tasks, which allows human agents to focus on more complex issues. The AI agents are highly customisable, enabling businesses to tailor their functionalities to meet specific needs and workflows, ensuring a seamless fit within existing processes.

Unlike traditional bots, Agentforce agents can be set up in minutes using pre-built templates and customised quickly with low-code solutions. Customers can define clear security parameters and ensure seamless escalation to human representatives when necessary. With Agentforce, users can build and customise autonomous agents that continuously improve their performance through self-learning, operate around the clock, and seamlessly transition to human support. These agents are dynamic and conversational, eliminating the need for manually built, predetermined conversational flows, making them significantly easier to set up.

1.2. Service Overview

5 Attributes of an Agent

- **Role:**
 - What is it trying to do?
 - What is its job to be done?
 - Determines which topics the agent can execute on and what are the instructions for those topics
- **Data:**
 - Refers to the information and insights the agent has access to.
 - Includes customer data, transaction history, and other relevant information.
 - Ensures the agent can make informed decisions and provide accurate assistance.
- **Actions:**
 - Encompasses the tasks and activities the agent can perform.
 - Includes both proactive and reactive measures to address customer needs.
 - Defines the standard operating procedures and workflows the agent follows.
 - Is it performing workflows? Connecting APIs into other systems? What are the specific capabilities tied to those topics and instructions? What are the specific capabilities the agent needs to have?
- **Guardrails:**
 - Sets the boundaries and limitations within which the agent operates.
 - Includes compliance with legal, ethical, and organisational policies.
 - Ensures the agent's actions align with the company's values and standards.
- **Channel:**
 - Specifies the communication mediums the agent uses to interact with customers.
 - Includes phone, email, chat, social media, and other platforms.
 - Ensures the agent can effectively reach and assist customers through their preferred channels.



1.3. Service Features

Agentforce, Salesforce's new AI-powered platform, brings significant benefits to customers by leveraging generative AI tools and large language models (LLMs) to create intelligent, autonomous agents. These AI agents communicate in everyday language, automating routine tasks, analysing data, answering questions, and providing support, thereby enhancing efficiency and effectiveness in the workplace:

- **Enhanced Productivity:** By automating routine tasks and handling everyday business interactions, Agentforce allows users to focus on more strategic activities, boosting overall productivity.
- **Natural Language Interaction:** Customers can interact with AI agents using natural language, making it easier to get relevant answers and support without needing specialised knowledge or training.
- **Secure and Relevant Responses:** The AI agents draw from secure, proprietary company data to provide accurate and relevant answers, ensuring that sensitive information is handled safely.
- **Real-Time Data Access:** Agentforce can securely access and act on real-time data from Salesforce, ensuring that decisions and actions are based on the most current information available.
- **Seamless Human-AI Collaboration:** The platform provides seamless handoffs to human agents when needed, ensuring a balanced approach between automation and human oversight, which enhances customer service and support.
- **Integration with Salesforce Ecosystem:** Sitting on Salesforce's Customer 360 and Data Cloud platform, Agentforce integrates with the core metadata platform and Data Cloud, leveraging real-time data to deliver trusted and accurate AI-driven solutions.

Agentforce's Features Narrative

In a world where generative AI tools are revolutionising the way we work, one groundbreaking development stands out: the creation of AI assistants, or AI agents. These intelligent agents, powered by large language models (LLMs), are transforming everyday business interactions by communicating in natural, everyday language.

Imagine a busy customer service department within a large public sector organisation, such as a big municipality. The team is inundated with routine tasks, data analysis, and a continuous flow of inquiries from the public. Now consider the ability to enhance customer service through innovative AI-powered solutions. Introducing Agentforce, Salesforce's innovative AI-powered platform, designed to assist public sector organisations with intelligent, autonomous agents that can support organisations to streamline operations. With Agentforce, the customer service team experiences a significant transformation.

These AI agents seamlessly handle routine tasks, analyse data, and provide instant support, making the team's work more efficient and effective. When a customer poses a question, the AI agent understands and responds in natural language, drawing from secure, proprietary company data to provide accurate and relevant answers.

The impact is immediate. Productivity increases as the AI agents take on everyday business interactions, freeing up the human team to focus on more strategic and complex issues. The

agents are versatile and capable of assisting across a wide range of workflows and tasks, ensuring that no query goes unanswered and no task is left incomplete.

Agentforce's ability to securely access real-time data from Salesforce is a game-changer. The AI agents act on the most current information available, making decisions and executing tasks with precision. And when a situation requires a human touch, Agentforce provides seamless handoffs to human agents, ensuring the perfect balance between automation and human oversight.

Residing on Salesforce's Customer 360 and Data Cloud platform, Agentforce integrates effortlessly with the core metadata platform and Data Cloud. This integration allows the AI agents to leverage real-time data, delivering trusted and accurate AI-driven solutions that customers can rely on.

For the customer service team, Agentforce is more than just a tool; it's a transformative force. It enhances productivity, ensures secure and relevant responses, and fosters seamless collaboration between AI and human agents. The team can now provide exceptional service, backed by the power of real-time data and intelligent automation.

In essence, Agentforce empowers organisations to navigate the complexities of modern business with ease, making work more efficient and interactions more meaningful. It's a testament to the incredible potential of AI, bringing a new era of productivity and innovation to customers everywhere.



Agentforce is built on the Salesforce Platform and is defined by a common metadata framework.


Agentforce agents are autonomous, proactive applications designed to execute specialised tasks

to help both your teams and your citizens. Agents use large language models (LLMs) to analyse and understand the full context of citizen interactions or an automated trigger, then autonomously reason through decisions on the next steps.

These agents generate responses that are consistent with your organisation's mission and guidelines using trusted data, including Salesforce CRM data, external data from Data Cloud, and more. Agents operate 24/7 across various platforms like self-service portals and messaging channels, handling tasks proactively within set guardrails. When faced with complex issues beyond their scope, they can escalate the matter to human agents, ensuring that queries are resolved efficiently and accurately.

The Agentforce out-of-the-box agents are easy to customise and deploy with clicks, not code. They can be set up in minutes, are easily scalable, and work around the clock across any channel. Examples include:

- **Service Agent** replaces traditional chatbots with AI that can handle a wide range of service issues without preprogrammed scenarios, improving citizen service efficiency.
- **Sales Development Representative (SDR)** engages with prospects 24/7, answering questions, managing objections, and scheduling meetings based on CRM and external data, allowing your organisation to focus on building deeper citizens relationships.
- **Campaign optimiser** automates the full campaign lifecycle, using AI to analyse, generate, personalise, and optimise marketing campaigns based on your organisation's mission.
- **Assistive Agent** supports users in the flow of their work with specific tasks by searching for data, creating action plans, and executing them to increase workplace efficiency.




Agentforce Architecture

Agentforce Building Blocks

Agentforce is a trusted conversational AI assistant seamlessly built into the Salesforce interface.

Your organisation can set up and launch agents for your users, and these are available in the flow of work in Salesforce. Your agent can perform business tasks on behalf of the users in your Salesforce org.

Salesforce Agentforce provides a set of tools and features that your organisation can use to create, customise, and deploy trusted agents and other innovative AI applications, complete with the right guardrails and supervision. Let's take a closer look and walk through the key components.



Salesforce Platform components for building agents and AI applications

Agentforce integrates various data sources, including CRM data and third-party data from external data sources, using Data Cloud and Mulesoft.

Salesforce metadata establishes universal rules that are enforced regardless of whether the data is accessed from traditional applications or agents. This includes permissions, sharing models, validation rules, and workflow automation to ensure data security and adherence to your organisation's security practices.

Metadata also allows LLMs to better understand the context and meaning of the data, which can lead to more accurate responses. For example, an LLM could use metadata to present CRM data to the user in a more useful and actionable way (ad hoc UI).

Reliable AI is dependent on high-quality and unified data. Salesforce Data Cloud brings together and unifies all your data — Salesforce and external, structured and unstructured — to ground AI with high-quality, relevant, and actionable information. Data Cloud provides your organisation with over 200 connectors and the ability to easily create custom ones to connect all of your data. Once your data is connected, unified, and harmonised, Data Cloud enables you to activate it at scale in AI agents, analytics, and other applications, delivering valuable insights and personalised experiences.

At the core of the architecture is the Agent, which processes and acts upon the integrated data. Agents operate based on specific Topics, which define the tasks to be performed. Instructions provide the necessary guidance and rules, while actions such as Flows, Apex, and Prompts are executed to support these tasks. Agents can handle multiple Topics, enabling them to perform a variety of functions.

Agentforce Agents are autonomous software systems that can analyse user or environmental inputs, identify tasks, reason through solutions, and orchestrate actions to complete them. Agents have different levels of autonomy. Assistive (partially autonomous) agents collaborate with humans to perform the task at hand. Autonomous agents operate independently without direct human supervision but with the robust guardrails described earlier in this post, and the ability to hand off tasks to humans as needed.

Actions are how an agent gets things done. An agent includes a library of actions, which are individual tasks the agent can perform. For example, if a user asks the agent for help with writing an email, the agent can launch an action that drafts and revises the email and grounds it in relevant Salesforce data. Salesforce provides some standard actions for common Salesforce tasks by default, and you can create custom actions specific to your organisation's mission.

A topic is a category of actions related to a particular job to be done. Topics contain actions, which are the tools available for the job, and instructions, which tell the agent how to make decisions.

When a user enters a question or request in the agent conversation, an agent selects a relevant topic. Based on the actions and instructions, the agent can either prompt the user for additional information (for example, a clarifying question or the required input for an action) or launch one or more actions. Salesforce provides a library of standard topics for common use cases, and you can create custom topics to meet your users' specific business needs.

The **Atlas Reasoning Engine** is the brain of Agentforce agents. It communicates with users and large language models (LLMs) to generate plans and execute relevant business actions. Essentially, it acts as a reasoning engine, determining the appropriate Topic based on user input.

Atlas Reasoning Engine orchestrates how an agent launches topics and actions during a conversation to accomplish a task for the user. When a user opens the Einstein panel and asks a question or enters an instruction, the planner service works with the LLM behind the scenes to carry out the request. Here's what it does.

- Interprets the user's request and classifies the utterance into a topic.
- Iteratively builds plans for accomplishing the user's goal.
- Finds and launches the right topics and actions to achieve the goal.

Thanks to Atlas Reasoning Engine, Agentforce's conversational automation solution can make proactive and intelligent decisions at scale with little or no human intervention. Several advancements make that possible, including:

Orchestration based on ReAct prompting - In the Reasoning and Acting (ReAct) mechanism, the system goes through a loop of reason, act, and observe until a user goal is fulfilled. This kind of looping approach lets the system consider any new information and ask clarifying questions or confirmations so that the user's goal is fulfilled as precisely as possible. This also leads to a much more fluid and natural language-like conversational experience.

Topic classification - When a user input comes in, it gets mapped to a topic, which contains the relevant set of instructions, business policies, and actions to fulfill that request. This mechanism helps define the scope of the task and the corresponding solution space for the LLM, allowing the system to scale effortlessly. Natural language instructions embedded in topics provide additional guidance and guardrails for the LLM, such as having actions to be executed in a certain sequence. These concepts allow agents to scale to thousands of actions safely and securely.

Use LLMs for responses - letting the LLM respond using the context in conversation history, allows for a far richer conversational experience. Users can ask for clarifications and ask follow-up questions to prior outputs, leading to an overall higher goal-fulfillment rate.

Thoughts/reasoning - Prompting LLMs to share thoughts or provide reasons for selecting certain actions prevents hallucinations. This provides visibility into the LLM's behavior, so that your organisation can fine-tune the agent to align with your mission. Reasoning is available in the Agent Builder canvas by default, and users can also prompt the agent with follow-up questions so it can explain its reasoning. This not only prevents hallucinations but helps build trust.

The Atlas Reasoning Engine is designed to enhance user interaction through faster, more capable, and more multi-turn conversations. Key features include:

- **Multi-turn chat:** The reasoning engine facilitates interactive communication with users by considering and adapting to added conversational context, enhancing the accuracy of the service provided.
- **Topic classification:** The reasoning engine classifies user utterances into topics based on predefined descriptions, ensuring relevant responses.
- **Instructions and actions:** Each topic includes specific instructions and actions, such as having actions to be executed in a certain sequence, to assist users accurately and efficiently.
- **Knowledge Retrieval:** The reasoning engine uses multiple techniques including advanced retrieval augmentation generation (RAG) which selectively uses multiple language models to iteratively refine the quality of queries, retrieving the most relevant knowledge, while also evaluating the quality of the response.
- **Searchable public data:** Agents can now securely access public data through the Einstein Trust Layer, expanding its knowledge base.

Agentforce is integrated with the Einstein Trust Layer, and the feature includes AI guardrails. Agentforce also respects standard Salesforce access controls.

The Einstein Trust Layer enables Agentforce to use any LLM safely by ensuring that no Salesforce data is viewed or retained by 3rd-party model providers. The Trust Layer lets your organisation benefit from generative AI without compromising citizens data. It also lets you use trusted data to improve generative AI responses.

- **Secure gateway:** Agentforce accesses models through a secure gateway that enforces security and privacy policies consistently across different model providers.
- **Data grounding:** The Trust Layer ensures that generative prompts are grounded and enriched in trusted data.

- **Data masking and compliance:** Before the request is sent to the model provider, it goes through a number of steps including data masking, which replaces personally identifiable information (PII) data with anonymised data to ensure data privacy and compliance.
- **Zero-data retention:** Your organisation data is never retained by a third-party LLM provider. Salesforce has zero-retention agreements with model providers, which means providers will not persist or further train their models with data sent from Salesforce.
- **Demasking, toxicity detection, and audit trail:** When the output is received from the model, it goes through another series of steps, including demasking, toxicity detection, and audit trail logging. Demasking restores the real data that was replaced by fake data for privacy. Toxicity detection checks for any harmful or offensive content in the output. Audit trail logging records the entire process for auditing purposes.

Agentforce includes ethical guardrails to minimise AI hallucinations and security guardrails to prevent threats and malicious attacks, such as prompt injections.

Built on the Salesforce metadata architecture, Agentforce is accessible through various digital channels, including Experience Cloud sites as a Chat component, and external channels like SMS, WhatsApp, and Facebook Messenger. This ensures that Agentforce can engage with users across multiple platforms.

[Trusted AI and Agents Impact Report](#)


Salesforce released its first Trusted AI and Agents Impact Report, outlining its approach to AI governance, product design, and testing to ensure Agentforce is deployed responsibly. This report is intended to guide other organisations on building trustworthy AI systems. Key points include:

- **Governance is paramount:** Salesforce emphasises establishing strong governance and ethical principles from the outset for their AI agent development.
- **Trusted AI Principles:** They have developed and evolved trusted AI principles since 2018, including guidelines for generative AI and AI agents.
- **Decision-Making Structures:** Robust structures, such as board committees and advisory councils, ensure responsible AI deployment.
- **Ethical Use Policies:** Acceptable Use Policies (AUP) and AI AUP are in place to guide ethical customer use of Salesforce's AI products.
- **Product Design and Development:** Trust is integrated into product design, particularly with Agentforce, which is built on the Einstein Trust Layer for data privacy, security, and ethical standards. "Trust patterns" act as guardrails, and accessibility is considered from the start.

Agent Studio

To create and configure agents, your organisation can use Agent Studio and Agent Builder. Agent Studio includes Agent Builder for creating custom AI actions to accomplish specific business tasks, Prompt Builder for building and activating custom prompts in the flow of work, and Model Builder, where users can build or import a variety of AI models.


Agent Builder helps every organisation configure and customise Agentforce copilots to their mission. Your organisation can use existing Salesforce configurations (such as Apex, Flow, and MuleSoft APIs), and new generative AI components like prompts, to enable the agent to complete tasks across any Salesforce application or external system.



Agent Studio's Agent Builder

Prompt Builder allows the creation of custom, reusable AI prompts without coding, simplifying complex processes and driving business innovation. This not only broadens the use of generative AI beyond conversational interfaces, but also allows organisations to design and repurpose prompts for use across other experiences. For example, a custom prompt can be seamlessly embedded in a contact record as a button, enabling an agent in the contact center to get a snapshot of all escalated cases for a citizen in one click.


Unlike other solutions that limit businesses to a single Large Language Model (LLM), Einstein 1 Studio provides the flexibility to connect to a variety of AI models. Additionally, Model Builder is a no-code, low-code, and pro-code way for organisations to build their own predictive AI models, trained on their data. For generative AI, Model Builder allows customers to select from LLMs managed by Salesforce, or bring their own models. Your organisation can use predictive and generative AI models and services from Salesforce partners, including Amazon Web Services (AWS) via Amazon Bedrock and Amazon SageMaker, Anthropic, Azure OpenAI, Cohere, Databricks, Google Cloud's Vertex AI, and OpenAI, and train or fine-tune selected models on Data Cloud data without moving or copying data.



Agent Studio's Model Builder

AI-Assisted Agent Generation

The AI-Assisted Agent Generation tool enables you to swiftly create an agent and bring your ideas to life by automatically generating topics, actions, and agent metadata based on a simple natural language description of the agent and your company. This streamlined process allows you to focus on refining and customizing the agent to meet your specific needs, ensuring a more efficient and effective implementation. By leveraging this tool, you can quickly shape your agent to align with your business objectives, enhancing both the agent's functionality and the overall customer experience.



AI-Assisted Agent Generation


Testing Center

With the Testing center and workbench you can generate synthetic test data based on your use case, run a variety of tests, and have several OOTB KPIs to review and compare the test results.

Deploying an agent in the Testing Center involves setting up a testing environment that includes the necessary agents to execute your tests. This process typically includes:

- **Agent Configuration:** Configure the agents with the appropriate settings and parameters to match your testing requirements.
- **Environment Setup:** Ensure that the testing environment is properly set up, including network configurations, access permissions, and any necessary integrations.
- **Test Execution:** Deploy the agents and execute the tests according to your test plan. Monitor the agents' performance and collect data for analysis.

By following these steps, you can effectively use the Testing Center and Workbench to generate synthetic test data, run comprehensive tests, and analyse the results using OOTB KPIs. This approach helps ensure that your implementation is robust, scalable, and ready for production deployment.



The screenshot displays the 'Testing Center' interface in Einstein Studio. It shows a table titled 'Simple Topic Test' with columns for '#', 'Utterances', 'Expected Topic', 'Agent Response', and 'Topic Test'. The table contains 16 rows of test data, with a summary row at the bottom indicating an overall 'Pass' status and a score of 85.6%.

#	Utterances	Expected Topic	Agent Response	Topic Test
1	Where is my order?	Order Management	We are currently processing your...	Pass
2	Can I change the shipping address...	Shipping & Handling	To change the shipping address f...	Pass
3	How can I track my package?	Shipping & Handling	You can track your package by e...	Fail
4	I need to update my payment m...	Billing	To update your payment method...	Pass
5	What is the status of my return?	Returns	The status of your return can be...	Fail
6	Do you offer expedited shipping...	Shipping & Handling	We offer expedited shipping opti...	Fail
7	Can I cancel my order and get a r...	Refunds	To cancel your order and receive...	Pass
8	How do I create an account on y...	Account Management	To create an account on our web...	Fail
9	I forgot my password, how can I...	Account Management	If you forget your password, you...	Pass
10	What is your return policy?	Returns	Our return policy can be found o...	Pass
11	Do you offer international shipp...	Shipping & Handling	We offer international shipping t...	Fail
12	Can I add a gift message to my o...	Order Management	You can add a gift message to y...	Pass
13	How do I update my billing addre...	Billing	To update your billing address, L...	Fail
14	Is there a way to expedite the pr...	Shipping & Handling	To expedite the processing of yo...	Fail
15	What payment methods do you...	Billing	We accept payment methods suc...	Fail
16	Can I change the size/color of an...	Product Information	To change the size/color of an it...	Pass
Summary				Pass 85.6%

Testing Center

How an Agent Takes Action

Agents take action and adhere to guardrails using natural language descriptions that outline the tasks and operational boundaries. Here's a summary of how they operate.

1. The agent first receives a trigger, which can be a conversation with an employee or citizen, a change in data, or an automation.
2. The agent uses the LLM and natural language descriptions to identify the context and select a topic that best fits the job to be done, including the scope, data required, and necessary conditions.
3. Depending on the task, an agent selects and chains actions. Those actions are executed via flows, apex classes, APIs, or direct prompts.

4. Agents dynamically plan and execute tasks while strictly following predefined guardrails. They also have built-in mechanisms for harm and toxicity detection, using the Einstein Trust Layer, ensuring they avoid engaging in inappropriate or harmful activities.

An AI agent is an intelligent system that can understand and respond to citizens' inquiries without human intervention. They rely on machine learning and natural language processing (NLP) to handle a wide range of tasks, from answering simple questions to resolving complex issues. Most importantly, AI agents can continuously improve their own performance through self-learning.

In particular:

- **Perception and data collection.** AI agents start by gathering data from a variety of sources including citizens interactions, transaction histories, and social media. This data is crucial for understanding the context and nuances of citizens' queries. Advanced AI agents can integrate and process data in real time, providing them with the most up-to-date information to handle inquiries effectively.
- **Decision making.** Using sophisticated machine learning models, AI agents analyse the collected data to identify patterns and make decisions. For example, they can determine the most appropriate response to a citizen query based on past interactions and the current context. This decision-making process is enhanced by the agent's ability to learn from previous experiences and refine its responses over time.
- **Action execution:** Once a decision is made, AI agents can execute the required action. This could involve answering a citizen query, processing a request, or escalating a complex issue to a human agent. The execution is designed to be seamless and efficient, ensuring that citizens receive timely and accurate responses.
- **Learning and adaptation:** AI agents continuously learn from each interaction, refining their algorithms to improve accuracy and effectiveness. They update their knowledge base and use feedback to enhance future interactions. This continuous learning capability ensures that AI agents remain effective and relevant, even as citizens' expectations change and your organisation mission evolves.

By combining these capabilities, AI agents can handle a wide range of service tasks autonomously, such as making recommendations, suggesting solutions or next steps, and engaging in follow-up interactions. This frees up human agents to focus on more complex and value-added activities.

1.4. Service Benefits

Agentforce is a comprehensive solution designed to enhance the efficiency and effectiveness of citizen interactions by leveraging the powerful capabilities of the Salesforce platform. It integrates seamlessly with existing Salesforce features such as flows, Apex, and APIs, enabling streamlined processes and enhanced automation without the need for extensive reconfiguration or additional investments. Key Benefits of Agentforce include:

Seamless Integration

- Agentforce utilises existing Salesforce infrastructure, ensuring smooth integration with platform features like flows, Apex, and APIs.

- This integration allows for streamlined processes and enhanced automation, reducing the need for extensive reconfiguration or additional investments.

Human-AI Collaboration

- Agentforce combines human expertise with artificial intelligence to provide the best possible customer experience.
- AI handles routine inquiries and tasks autonomously, freeing up human agents to focus on more complex interactions.
- This collaboration improves efficiency and enhances the quality of customer service by ensuring human agents are available for personalised support.

Customer 360 Data

- Agentforce is grounded in Customer 360 Data, providing AI and human agents with access to high-quality, trusted, and harmonised data through Data Cloud.
- This comprehensive view of the customer's history and preferences allows for more accurate and relevant responses, resulting in personalised and effective citizen interactions.

Flexible Agent Modes

- Agentforce supports both autonomous and assistive modes for agents.
- Autonomous agents handle tasks independently, reducing the workload on human agents and increasing overall efficiency.
- Assistive agents provide support and guidance to human agents, enhancing their capabilities and ensuring they have the necessary information and tools to deliver exceptional service.

Trust and Security

- Agentforce ensures ethical and secure interactions through features like ethical guardrails, data masking, zero retention policies, and Hyperforce.
- This focus on trust and security is crucial for maintaining citizen confidence and ensuring sensitive information is handled appropriately.

Openness and Extensibility

- Agentforce is designed to be open and extensible, allowing for integration with ISV actions, prompts, topics, and the ability to bring your own language model (BYO LLM).
- This extensibility ensures that Agentforce can adapt to the unique needs of different organisations and can be customised to fit specific workflows and requirements.

Low-Code Tools

- Agentforce provides a set of low-code tools including Agent Builder, Prompt Builder, Model Builder, and Flow Builder.
- These tools allow users to build, customise, test, and manage Agentforce Agents with ease.

Differentiating Characteristics

- **Trusted:** Agentforce protects your data using the Einstein Trust Layer and the same metadata, permissions, and sharing models as traditional Salesforce applications.
- **Powerful:** Agentforce Agents utilise industry-leading Salesforce applications to deliver transformative experiences across sales, service, commerce, marketing, and industries.
- **Grounded in Unified Data:** Agentforce Agents deliver more accurate and relevant outcomes by grounding AI in all the relevant data made available and unified by Data Cloud.
- **Low-Code Tools:** Agentforce Agents can be built, customised, tested, and managed using a set of low-code tools including Agent Builder, Prompt Builder, Model Builder, Flow Builder, and more.




Agentforce brings together Humans + AI + Data + Actions

Efficiency and Productivity Gains with Agentforce

- **Automation of Routine Tasks:** Agentforce automates repetitive and time-consuming tasks such as data entry, scheduling, and basic customer inquiries. This reduces the manual workload on employees, allowing them to focus on more strategic and value-added activities.
- **Faster Response Times:** With AI agents handling routine queries and tasks in real-time, response times to customer inquiries and internal requests are significantly reduced. This leads to quicker resolutions and improved customer satisfaction.
- **24/7 Availability:** AI agents can operate around the clock, ensuring that tasks and customer support are handled even outside of regular business hours. This continuous operation enhances overall productivity and service levels.
- **Data-Driven Decision Making:** By analysing data and providing insights, Agentforce helps employees make informed decisions quickly. This reduces the time spent on data gathering and analysis, speeding up decision-making processes.

- **Streamlined Workflows:** Agentforce can integrate with various systems and tools, streamlining workflows and reducing the need for manual data transfers between different platforms. This integration minimises errors and enhances efficiency.
- **Embed Agentforce in any workflow:** Embedding Agentforce into your workflows allows you to leverage the power of conversational AI across various business processes. By connecting agents to Flows, Apex, and APIs using Invocable Actions, you can create a seamless and efficient automation experience. This integration ensures that your workflows are dynamic and responsive to changing conditions.
 - **Grounded in Your Data:** Agentforce gathers the necessary context to handle complex requests, ensuring that your automations are both accurate and relevant. This means that the agents can pull in data from various sources within your Salesforce environment to provide comprehensive and informed responses, enhancing the overall effectiveness of your workflows.
 - **Simple Configuration:** Utilise familiar low-code tools and APIs to configure headless agents, making the setup process straightforward and accessible. This simplicity allows even those with minimal coding experience to set up and manage agents effectively, reducing the time and effort required to deploy powerful automation solutions.
 - **Event-Driven and Proactive:** Invoke Agentforce from Flows, APIs, or events to trigger automation, enabling proactive responses and actions based on real-time data and events. This capability allows your workflows to react instantly to changes and events, ensuring timely and appropriate actions are taken without manual intervention.



Embed Agentforce into any Workflow

How does Agentforce enhance employee performance?

- **Focus on High-Value Tasks:** By offloading routine and repetitive tasks to AI agents, employees can concentrate on high-value tasks that require human creativity, critical thinking, and strategic planning. This shift not only improves job satisfaction but also enhances overall performance.
- **Reduced Cognitive Load:** Automating mundane tasks reduces the cognitive load on employees, allowing them to work more efficiently and with greater focus. This leads to higher quality work and increased productivity.
- **Enhanced Collaboration:** Agentforce facilitates better collaboration by providing real-time data and insights that can be easily shared across teams. This ensures that everyone has access to the same information, improving coordination and teamwork.
- **Skill Development:** With routine tasks automated, employees have more time to engage in professional development and skill-building activities. This continuous learning enhances their capabilities and prepares them for more complex roles.
- **Improved Accuracy and Consistency:** AI agents reduce the risk of human error in routine tasks, ensuring higher accuracy and consistency in operations. This reliability enhances overall performance and builds trust in the system.
- **Seamless Human-AI Collaboration:** Agentforce provides seamless handoffs between AI agents and human employees when needed. This ensures that complex or sensitive tasks are handled appropriately, combining the efficiency of AI with the nuanced understanding of human agents.

Compatibility with Third-Party Applications and Other Tools and Systems

Agentforce is designed to be highly compatible with both third-party applications and other tools and systems your organisation uses. Its robust integration capabilities ensure that you can incorporate it into your existing technology ecosystem, enhancing productivity and efficiency without significant disruptions.

Compatibility with Third-Party Applications

Agentforce's compatibility with a wide range of third-party applications ensures that organisations can seamlessly integrate Agentforce into their existing technology stack without significant disruptions.

- **API Integrations:** Agentforce supports robust API integrations, allowing it to connect with various third-party applications. This means you can integrate it with CRM systems, marketing automation tools, customer support platforms, and more.
- **MuleSoft for Agentforce:** MuleSoft Topic Center allows developers to use natural language to create Agentforce topics and actions from MuleSoft APIs. The MuleSoft API Catalog simplifies AI agent development by centralising access to diverse APIs from MuleSoft, Salesforce, and Heroku within Agentforce. The MuleSoft Agentforce connector allows developers to integrate Agentforce into existing workflows, enabling APIs and integrations to trigger agents for data enrichment and action execution across systems.
- **AppExchange:** Salesforce's AppExchange offers a vast marketplace of third-party applications that are pre-integrated with Salesforce. Agentforce can leverage these integrations to extend its functionality and provide a more comprehensive solution.
- **Custom Integrations:** For applications that are not available on AppExchange, custom integrations can be developed using Salesforce's extensive API capabilities. This flexibility ensures that Agentforce can work with virtually any third-party tool your organisation uses.
- **AgentExchange:** The Salesforce [AgentExchange](#) is a marketplace for Agentforce, their digital labor platform. AgentExchange allows partners and developers to build and sell AI

components like actions, templates, and topics. It features over 200 partners and hundreds of pre-built solutions to help businesses create and deploy AI agents. The platform aims to expand the capabilities of Agentforce and participate in the digital labor market.



Compatibility with Other Tools and Systems

Agentforce is built on the Salesforce Platform, which is known for its strong integration capabilities. Here's how it ensures compatibility with other tools and systems:

- **Salesforce Ecosystem:** Agentforce integrates seamlessly with other Salesforce products, including Customer 360, Data Cloud, Sales Cloud, Service Cloud, and Marketing Cloud. This ensures a unified experience across all Salesforce tools.
- **Data Integration:** Agentforce can access and act on real-time data from Salesforce and other connected systems. This real-time data integration ensures that the AI agents have the most up-to-date information to make informed decisions and execute tasks.
- **Middleware Solutions:** For more complex integration scenarios, middleware solutions like MuleSoft (a Salesforce company) can be used to connect Agentforce with various enterprise systems, databases, and applications. This ensures smooth data flow and interoperability across different platforms.
- **Standard Protocols:** Agentforce supports standard integration protocols such as REST, SOAP, and OData, making it easier to connect with a wide range of systems and applications.
- **Custom Development:** Salesforce's platform allows for custom development using Apex, Visualforce, and Lightning components. This means you can build custom connectors and integrations tailored to your specific needs.

Agentforce API


The **Agentforce API** empowers developers to seamlessly embed Agentforce into their applications and workflows, enabling limitless automation and creating dynamic experiences across their entire ecosystem.

Key Capabilities. With the **Agentforce API**, developers can:

- **Spin Up Headless Agents Easily:** Enable automation and interactions without the constraints of a user interface.
- **Seamlessly Connect Agents:** Integrate agents with existing platforms, workflows, and Lightning Web Components (LWCs) using standardised API endpoints.
- **Create a Distributed Ecosystem:** Facilitate smooth agent-to-agent transfers, whether the agents are built with Agentforce or other platforms.

Key API Features

- **Synchronous and Streaming Endpoints:** Provide full or tokenized responses to suit different interaction needs.
- **Easy Session Lifecycle Management:** Start sessions, send and receive messages, and end sessions with ease.
- **Secure Authentication:** Utilise JWT tokens tied to your organisation for secure and authenticated interactions.
- **Contextual Interactions:** Pass context variables to enable tailored and relevant agent experiences.



Agentforce API

How Public Sector Organisations can benefit from Agentforce

Government organisations can significantly enhance their operations and service delivery by leveraging Agentforce. This platform is highly customisable, allowing it to meet the specific needs and data quality requirements of each organisation. Its versatility makes it a valuable tool for various governmental functions.

Personalised responses to prospective students, making the process more efficient. By analysing data on student preferences and behaviors, institutions can tailor their outreach efforts to attract the best candidates. This can lead to a more effective recruitment process and a higher yield of accepted offers.

In terms of student services, AI agents can respond to and resolve cases faster, allowing staff to focus on more complex issues. This can enhance the overall student experience by providing timely and accurate support for academic, financial, and personal concerns. Additionally, AI can assist in prospect research for recruiting or fundraising, providing valuable insights to staff about prospective students or donors. This can enhance the institution's ability to attract and retain top talent and secure necessary funding.

Furthermore, Agentforce can support personalised learning initiatives. By analysing student performance data, AI can identify areas where students may need additional support and recommend tailored learning resources. This can lead to improved academic outcomes and a more engaging learning experience.

Additionally, AI-driven insights can help educators identify trends and patterns in student performance, enabling data-driven decision-making and continuous improvement in teaching and learning practices.

EMEA Government

For governments in the EMEA (Europe, Middle East, and Africa) region, Agentforce offers several benefits. The platform's AI capabilities can help improve interactions with citizens by providing quick and accurate responses to inquiries. This can enhance citizen satisfaction and trust in government services. Additionally, Agentforce's data analytics capabilities can help EMEA governments identify trends and insights that can inform policy decisions and strategic planning.

In the healthcare sector, Agentforce's predictive analytics capabilities can help improve provider productivity and personalise care, leading to better patient outcomes. By analysing patient data, healthcare providers can identify trends and potential health risks, enabling proactive care and reducing the burden on healthcare systems. This can result in more efficient use of resources and improved patient satisfaction. For example, predictive models can forecast disease outbreaks, allowing for timely interventions and resource allocation.


Moreover, Agentforce can support EMEA governments in managing large-scale public projects. By leveraging AI and data analytics, governments can monitor project progress, identify potential bottlenecks, and ensure that projects are completed on time and within budget. This can lead to more efficient infrastructure development and better utilisation of public funds.

1.5. Main components & Functions of this service

Agents are autonomous, proactive applications designed to execute specialised tasks to help employees and customers. Agents use large language models (LLMs) to analyse and understand the full context of customer interactions or an automated trigger, then reason through decisions on the next steps autonomously.

These agents generate responses that are consistent with your company's brand voice and guidelines using trusted business data, including Salesforce CRM data, external data from Data 360, and more. These agents are capable of operating 24/7 across various platforms like self-service portals and messaging channels, handling tasks proactively within set guardrails. When faced with complex issues beyond their scope, they can escalate the matter to human agents, ensuring that queries are resolved efficiently and accurately.

The Agentforce out-of-the-box agents are easy to customise and deploy with clicks, not code. They can be set up in minutes, are easily scalable, and work around the clock across any channel. Here are just some of the agents available.



- Service Agent replaces traditional chatbots with AI that can handle a wide range of service issues without preprogrammed scenarios, improving customer service efficiency.
- Sales Development Representative (SDR) engages with prospects 24/7, answering questions, managing objections, and scheduling meetings based on CRM and external data, allowing your sellers to focus on building deeper customer relationships.
- Sales Coach provides personalised role-play sessions for your sales team, using Salesforce data and generative AI to help sellers practice pitches and objections tailored to specific deals.
- Merchandiser assists your ecommerce merchandisers with site setup, goal setting, personalised promotions, product descriptions, and data-driven insights, simplifying daily tasks.
- Buyer Agent enhances the B2B buying experience, helping your buyers find products, make purchases, and track orders via chat or within sales portals.
- Personal Shopper acts as a digital concierge on your ecommerce sites or messaging apps, offering personalised product recommendations and assisting with search queries.
- Campaign optimiser automates the full campaign lifecycle, using AI to analyse, generate, personalise, and optimise marketing campaigns based on business goals.
- Agentforce supports employees in the flow of their work with specific tasks by searching for data, creating action plans, and executing them to increase workplace efficiency.

Each agent has parameters that define the job it can do and the action it can take.

- **Role:** An agent's purpose. This defines the job to be done and the broader goals the agent should achieve on your team.
- **Knowledge:** The data an agent needs to be successful. This could include company knowledge articles, CRM data, external data via Data 360, public websites, and so on.
- **Actions:** The goals an agent can fulfill. This is the predefined task an agent can execute to do its job based on a trigger or instruction. For example, it could run a flow, prompt template, or Apex.
- **Guardrails:** The guidelines an agent can operate under. These can be natural-language instructions telling the agent what it can and can't do, when to escalate to a human, or could come from built-in security features in the Einstein Trust Layer.
- **Channels:** The applications where an agent gets work done. This can be your website, CRM, mobile app, Slack, and more.

Agentforce Reasoning Engine, the Brain of an Agent

The Agentforce Reasoning Engine is designed to enhance user interaction through faster, more capable, and more multi-turn conversations. Here's a brief overview and key features.

- **Multi-turn chat:** The reasoning engine facilitates interactive communication with users by considering and adapting to added conversational context, enhancing the accuracy of the service provided.
- **Topic classification:** The reasoning engine classifies user utterances into topics based on predefined descriptions, ensuring relevant responses.
- **Instructions and actions:** Each topic includes specific instructions and actions, such as verifying order details or obtaining further information, to assist users accurately and efficiently.
- **Knowledge Retrieval:** The reasoning engine uses multiple techniques including advanced retrieval augmentation generation (RAG) which selectively uses multiple language models to iteratively refine the quality of queries, retrieving the most relevant knowledge chunks, while also evaluating the quality of the response.
- **Searchable public data:** Agents can now securely access public data through the Einstein Trust Layer, expanding its knowledge base.

How an Agent Takes Action

Agents take action and adhere to guardrails using natural language descriptions that outline the tasks and operational boundaries. Here's a summary of how they operate.

- The agent first receives a trigger, which can be a conversation with an employee or customer, a change in data, or an automation.
- The agent uses the LLM and natural language descriptions to identify the context and select a topic that best fits the job to be done, including the scope, data required, and necessary conditions.

- Depending on the task, an agent selects and chains actions. Those actions are executed via flows, apex classes, APIs, or direct prompts.
- Agents dynamically plan and execute tasks while strictly following predefined guardrails. They also have built-in mechanisms for harm and toxicity detection, using the Einstein Trust Layer, ensuring they avoid engaging in inappropriate or harmful activities.

2. Information Assurance

Salesforce holds a suitably scoped ISO:27001 certificate that covers the majority of its core services, please refer to the Salesforce Security, Privacy and Architecture (SPARC) documentation [here](#) for further product specific information. In addition, for country and product specific compliance please refer to this [link](#) for current status.

Salesforce.com EMEA Limited and its affiliates are committed to achieving and maintaining customer trust. Integral to this mission is providing a robust security and privacy program that carefully considers data protection matters.

In the UK Salesforce has completed an external CHECK Security assessment performed against its core services by a CREST member organisation, Salesforce also holds a current Cyber Essentials certificate, again completed by a CREST member organisation and can provide a response to the NCSC 14 security principles upon request.

In accordance with the EU Data Protection Directive and implementing national legislation, the Salesforce Processor BCR (Binding Corporate Rules) is intended to provide an adequate level of protection for Personal Data during international transfers within the Salesforce Group made on behalf of Customers and under their instructions, further information is available here <https://secure.sfdcstatic.com/assets/pdf/misc/Salesforce-Processor-BCR.pdf>

Salesforce adheres to the principles of the EU-U.S. Privacy Shield framework with respect to personal data submitted by Salesforce's customers in reliance on the Privacy Shield to the following online services: Sales Cloud, Service Cloud, Force.com, Communities, Chatter, Site.com, Database.com, Analytics Cloud, Financial Services Cloud, Health Cloud, Heroku, Pardot and Configure, Price, Quote (CPQ) <https://www.salesforce.com/assets/pdf/misc/privacy-shield-notice.pdf>

In addition, we currently hold the following certifications:

Geographical Recognition

- TRUSTe Certified Privacy Seal
- Japan Privacy Seal from the Japan Information Processing Development Corporation (JIPDC)
- TUV Certificate
- FedRAMP
- Cyber Essentials
- Privacy Shield

Global Audit Compliance

- ISO 27001, 27018
- SSAE 18/ISAE 3402
- SOC-1
- SOC-2
- SOC-3 (SysTrust)
- PCI-DSS
- Cyber Essentials

Learn more here: [Salesforce Services Trust and Compliance Documentation](#)

Relating to the Government Security Classification (GSC) scheme and the National Cyber Security Centre (NCSC), Salesforce has various Public Sector customers, holding a variety of data at OFFICIAL and OFFICIAL Sensitive levels. Further Salesforce has various documents to assist organisations assessing the Security posture of Salesforce such as a response to the NCSC 14 Security principles. Please contact us for up-to-date information and access to additional information.

3. Data Backup / Restore and Disaster Recovery

Redundancy and Scalability

The Salesforce service is built for high availability across 3 availability zones within the region. Each availability zone includes the following:

- Multiple network carriers for customer connectivity
- Multiple ISPs for customer transit and internal replication
- Multiple dedicated connections for DR/BCP
- Redundant cloud-based networking infrastructure
- Web, Application, API, Cache, Search, Index, Query and Batch servers are load balanced and live within multiple availability zones within the region, with failover capabilities
- Database servers are replicated both within the same availability zone for performance reasons, and across 3 availability zones to ensure availability.

Extensive use of high-availability servers and network technologies, and a carrier-neutral network strategy, help to minimise the risk of single points of failure, and provide a highly resilient environment with maximum uptime and performance. The Salesforce Services are configured to be N+1 redundant at a minimum, where N is the number of components of a given type needed for the service to operate, and +1 is the redundancy. In many cases, Salesforce has more than one piece of redundant infrastructure for a given function.

Salesforce provides Enterprise class recoverability through replication of data across three distinct availability zones.

Disaster Recovery

Production data centres are designed to mitigate the risk of single points of failure and provide a resilient environment to support service continuity and performance. The Covered Services utilise

secondary facilities that are geographically diverse from their primary data centres, along with required hardware, software, and Internet connectivity, in the event Salesforce production facilities at the primary data centres were to be rendered unavailable.

Salesforce maintains a disaster recovery plan that supports a robust business continuity strategy for the production services and platforms. Our platform is configured in a high availability mode and redundant fashion, and our Global Product Operations team monitors the software around-the-clock to ensure top performance. Product-specific information is available in the Security, Privacy and Architecture (SPARC) documentation for the respective services available on our website (<https://www.salesforce.com/company/legal/trust-and-compliance-documentation/>), and customers can view detailed policy documentation at <https://compliance.salesforce.com/en/disaster-recovery-bcp>.

Backup

All networking components, network accelerators, load balancers, Web servers, and application servers are configured in a redundant configuration. All Customer Data submitted to the Covered Services is stored on a primary database server with multiple active clusters for higher availability. All Customer Data submitted to the Covered Services is stored on highly redundant carrier-class disk storage and multiple data paths to ensure reliability and performance. All Customer Data submitted to the Covered Services, up to the last committed transaction, is automatically replicated on a near real-time basis to the secondary site and backed up to localised data stores. Backups are verified for integrity and stored in the same data centres as their instance. The foregoing replication and backups may not be available to the extent that Customer has managed packages that are uninstalled by Customer's administrator during the subscription term because doing so may delete Customer Data submitted to such services without any possibility of recovery.

To help customers routinely back up their data, Salesforce offers several native options that are available for no additional cost to customers. Salesforce provides tools like Data Loader and the API as a method for customers to manually restore their data. It is important to note the order in which data is restored, so that relationships and the connection to related records can be preserved.

The following documentation and options are available to customers as a method of backing up their data:

- Data Export Service: Manual or scheduled exports of your data via the UI. Export Backup Data from Salesforce:
https://help.salesforce.com/s/articleView?id=xcloud.admin_exportdata.htm&type=5..
- Data Loader: Manual on-demand exports of your data via the API. Export Data:
https://help.salesforce.com/articleView?id=exporting_data.htm&type=5.
- Report Export: Manual on-demand exports of your data via reports. Export a Report:
https://help.salesforce.com/articleView?id=reports_export.htm&type=5.
- Salesforce Backup and Restore product:
<https://www.salesforce.com/platform/data-backup-recovery/>

Restore

Salesforce offers several native options at no additional cost to help customers routinely back up their data. Depending on your edition, your org can generate and export backup files on a weekly or monthly basis. Individual users can view and restore their deleted records from the Recycle Bin, where all deleted records reside for 15 days before being deleted from the system. Admins also have access to an org-wide Recycle Bin to restore or permanently delete records during the 15-day window. Salesforce also supports manual restore via Data Loader and API.

As an add-on solution, Salesforce Backup provides additional backup and restore capabilities. You can configure and maintain automatic backups, make a backup on demand, and quickly find and restore data. You can control the order in which data is restored to preserve relationships and connections to related records.

You can find additional information and best practices for backup and recovery at <http://sfdc.co/BackupAndRecovery>.

4. On-Boarding

Describe how easy and simple it is for a customer to On-Board to our service

4.1. Deployment

Salesforce is a SaaS (Software as a Service) solution and users access the tools via a web browser. Therefore, no software or hardware deployment is necessary. Implementation takes the form of assigning usernames and passwords to staff and (if needed) completing webinar-based training on the application.

4.2. Configuration and Customisation

The service provides many configuration options, including integration with custom applications. Details can be found [here](#)

5. Off-Boarding

Off-boarding from the service is as simple as exporting and downloading your data if you choose to do so. We provide your data in an industry standard readable format to make it as easy as possible for you to migrate to another service if you wish to do so.

For further details of data extraction and removal, please refer to section 5.2 below.

5.1. Termination

Details of termination options and implications are contained in our terms.

The Supplier would never terminate except for buyers' material breach of the Call-Off Contract and the terms.

5.2. Data Extraction and removal

With a commercial model that requires customer trust to be at the core of everything we do, Salesforce also helps ensure that our customers can exit our service with open and transparent

processes and clearly defined commercial considerations. As such, should for any reason a customer wish to cease using the Salesforce service then outlined below are some of the key technical considerations for ending the subscription service and extracting both the customer's data and the technical investment made in the platform. Further detail relating to an exit strategy can also be discussed on understanding the nature of any potential solution

- Within 30 days post contract termination, customers may request return of their respective Customer Data submitted to the Covered Services (to the extent such data has not been deleted by Customer). Salesforce shall provide such Customer Data via a downloadable file in comma separated value (.csv) format and attachments in their native format. Note that Customer Data your organisation submits to Einstein Analytics instance groups for analysis is derived from other data to which your organisation has access, for example, data stored by your organisation using Service Cloud, Sales Cloud, third party applications, etc.
- Data export is provided at no additional charge. Full details are contained in our terms.

Salesforce understands the intentions of G-Cloud, and is keen to support customers adoption of cloud technology, however that should also include knowing how to leave a given cloud technology. Salesforce has an exit strategy paper that can be used to help plan a path away from Salesforce. Notwithstanding that Salesforce has one of the lowest attrition rates in the industry, we remain committed to your success. Continued investment is how we sustain 3 upgrades every year and a constant delivery on new innovation in every update.

6. Pricing Summary

Comprehensive pricing details can be found in our separate pricing document.

The service has various pricing levels and bandings based on type of service required and the volume needed. Please refer to pricing documents for detailed pricing.

7. Service Constraints

Here are some key considerations we wish to highlight; a full list is contained in our Supplier Terms.

Usage Limits

- Services and content are subject to usage limits, including, for example, the quantities specified in order forms.
- A user's password may not be shared with any other individual.

If you exceed a contractual usage limit, we may work with you to seek a reduction in your usage so that it conforms to that limit. If, notwithstanding our efforts, you are unable or unwilling to abide by a contractual usage limit, you will execute an order form for additional quantities of the applicable services or content promptly upon our request.

Customisation and configuration of the service is independent of the underlying infrastructure; upgrades do not impact any changes you may have made. As a result of this independence it is

typical for customers to customise their service to their own requirements. As the service is SaaS based, Salesforce is aware of which users are using which features etc., therefore we are aware of any potential impact a feature deprecation would have to customers and can work with them should there be a need for a feature deprecation.

8. Service Management

Your success is our success

Customer success is a top priority for Salesforce. Every customer gets a Standard Success Plan for online support and training. Our most successful customers take advantage of our Premier Success Plans to achieve an 80% higher return on their Salesforce investment. Large enterprise customers can benefit from Signature Success, our highest level of service to support their most critical business demands.

Premier and Signature Success packages can be purchased through G-Cloud specialist cloud services.

Standard Success

Every Salesforce customer gets a Standard Success Plan for online support and training. Our Standard Success Plan, included with each licence, provides:

- Success Communities to share with other customers
- Guided Journeys on how to use Salesforce
- Circles of Success Interactive Events
- 12/5/365 online case submission
- Response in two business days
- Trailhead online training
- “Getting Started” online training catalogue

Standard Success is for companies that need standard guidance in getting started with Salesforce. If you need a faster response, 24x7 support coverage, and/or a comprehensive training solution, we recommend our Premier Success Plans.

Premier Success

Our most successful customers take advantage of Premier Success to achieve an 80% higher return on their Salesforce investment.

The Premier Success Plan provides specialised guidance whether you have how-to questions, experience technical issues, need troubleshooting, or want to increase the value you get from Salesforce.

Benefits include:

- 24/7 online and phone support by senior support analysts
- Expert coaching sessions
- Reviews of your platform health and business value
- Developer support
- Expert Assistance

Specialised Guidance

- A discount on all Trailhead Academy courses and certifications
- Quicker response times than Standard Success. When critical issues arise, our skilled support engineers respond within one hour.

Expert Assistance

Expert coaching sessions are specialised engagements designed to help you get more value from Salesforce products. With Premier Success, you can attend webinars on a specific topic and watch coaching videos and then have an individual follow-up session to dive deeper.

We also offer personalised sessions with Salesforce experts to help you overcome obstacles and drive long-term success. More than 200 options cover a range of needs and interests across Salesforce products.

To get real-time answers to your questions, Premier Success includes live Q&A sessions with Salesforce experts, addressing topics from adoption and how-tos to best practices. Specialised technical support is also available for admins and developers to troubleshoot custom code issues.

Specialised Guidance

To help ensure that you see continuing success with Salesforce, Premier Success includes personalised guidance and insights. Through periodic reviews and check-ins, we evaluate your platform health and value maturity. These technical and business reviews assess key areas of platform performance, prioritise areas for growth, and set and track progress against your targets with quantifiable success metrics.

Signature Success

Extend your team with Signature Success. Rest assured knowing our certified experts are here to help you maintain your Salesforce solution.

Signature Success is the highest level of support from Salesforce and provides a high-touch experience led by a named expert who acts as an extension of your team. Customers with Signature Success benefit from increased performance and productivity through all the features of Premier Success plus:

- 15 minutes initial response for critical issues.
- The fastest response times from our most skilled support engineers.
- A designated Technical Account Manager.
- Customer Success Score for actionable insights and recommendations.
- Technical Health Reviews.
- Proactive Monitoring and Key Event Management.

Technical Account Manager (TAM)

Account management by a named champion. An assigned TAM provides consistent advocacy and guidance with a deep understanding of your business. These highly experienced technical experts provide support case oversight and escalation, weekly meetings, and tailored solution guidance.

Proactive services

Signature Success includes 24/7 monitoring that is tailored to your configuration. Your TAM coordinates with Proactive Services Engineers experts and helps ensure you prevent or mitigate potential issues identified, before they can create business disruption. If anything does go wrong, you can expect early alerts, remediation, and reviews to eliminate recurring root cause patterns.

Available at all times and offers the fastest case response times, including 15 minutes for severity 1 issues via phone or chat through our Help site.

Accelerators

Accelerators are quick, personalised work sessions that solve specific Salesforce challenges. The list of accelerators available varies depending on the type of plan you chose; there are no accelerators available with the standard plan. A customer can have unlimited accelerators as long as only one is being delivered at a time.

Learn more: <https://www.salesforce.com/services/success-plans/overview/>

There is no financial recompense model for not meeting service levels.

9. Customer Responsibilities

Here are some key responsibilities we wish to highlight; a full list is contained in our Salesforce Supplier Terms.

Management

Using our online tools, you will need to manage your usage of and access to the service, for example user accounts, installed applications, sites, etc.

Compliance

In using the service, you need to ensure your compliance with our terms, which include:

- Being responsible for the accuracy, quality and legality of your data and the means by which you acquired your data. Using commercially reasonable efforts to prevent unauthorised access to or use of the service, and notify us promptly of any such unauthorised access or use.

Usage Restrictions & Information Security

You must:

- Ensure that only information of an appropriate security classification is placed into the service.

You must not:

- Use the service to store or transmit infringing, libellous, or otherwise unlawful or tortious material, or to store or transmit material in violation of third-party privacy rights.
- Use the service to store or transmit malicious code.
- Interfere with or disrupt the integrity or performance of any service or third-party data contained therein.

- Attempt to gain unauthorised access to any service or content or its related systems or networks.
- Permit direct or indirect access to or use of any service or content in a way that circumvents a contractual usage limit.

10. Client-side Technical Requirements

Salesforce is a SaaS (Software as a Service) solution and users access the tools via a web browser. Therefore, no hardware or software installation is necessary.

Browser requirements

Salesforce supports a range of popular browsers. Learn more: [Supported browsers](#)

Internet access

Salesforce is designed to use as little bandwidth as possible, so that the service performs adequately over high-speed, wireless and mobile Internet connections.

While average page size is on the order of 90KB, Salesforce supports compression as defined in the HTTP 1.1 standard to compress the HTML content before it is transmitted as data across the Internet to a user's computer. The compression often reduces the amount of transmitted data to as little as 10KB per page viewed, due to the lack of image content. The site was designed with minimum bandwidth requirements in mind, hence the extensive use of colour coding instead of images. Our average user also is known to view roughly 120 pages from our site per day. However, it is best to measure any page that has been customised, especially if Visual Force components have been added to the page, to get an accurate measurement of the page size.

Our application is stateless; therefore, there are no communication requirements in the background once the page loads like traditional client server applications e.g. Outlook, therefore, once the page loads there are no additional bandwidth requirements until a user queries or writes information to Salesforce. Further information here

<https://help.salesforce.com/articleView?id=000004958&type=1>

11. Planned Maintenance Windows

When maintenance is scheduled, Salesforce publishes the dates and times of the maintenance windows on trust.salesforce.com. Premier Alerts are sent via email when the maintenance windows are posted to trust.salesforce.com. In the event of planned maintenance that requires customer action in advance, such as updating network settings in preparation for additional login pools, Salesforce endeavours to communicate via email to system administrators of your organisation months prior to the maintenance. If emergency system maintenance is required, customers may be notified less than one week in advance.

There are two types of maintenance at Salesforce:

- System maintenance is for sustaining the security, availability, and performance of the infrastructure supporting Salesforce services.

- Release maintenance is for upgrading Salesforce services to the latest product version to deliver enhanced features and functionality. There are three different kinds of release maintenance: major releases, patch releases, and emergency releases.

Major release maintenance dates and times are posted on trust.salesforce.com approximately one year before the release date. Major release maintenance occurs three times per year.

Patch releases and emergency releases are used to deliver scheduled and ad hoc application fixes and are typically seamless to customers. Whenever possible, patches and emergency releases are deployed during off-peak hours and without downtime. You can see our preferred maintenance schedule at

https://help.salesforce.com/apex/HTViewSolution?id=000176208&language=en_US or learn more at <http://trust.salesforce.com>

12. Training

We provide webinar-based and instructor led training for the majority of our services. Details can be found in our separate listing for Salesforce Training also listed on the Digital Marketplace. Further information can also be found here:

https://www.salesforce.com/services-training/training_certification/training-by-cloud.jsp

13. Data Centre Locations

For the Salesforce Services (services branded as Force.com, Site.com, Sales Cloud, Service Cloud, and Chatter), if a new 'Org' of the Salesforce Service is provisioned for a Customer with a billing address in the UK, Customer Data in that 'Org' is stored in data centres in England.

In addition, Salesforce may store information in data centres located outside of EEA, such as identifying information, relating to the Customer's instances(s) of Salesforce Services and users for the purposes of operating the Salesforce Services, such as facilitating the login process and the provision of customer support.

Such identifying information as provided by the Customer in its provision of user accounts shall only include the following personal data about users: first and last name, email address, username, phone number, and physical business address.

For other services which are not Salesforce Services, Salesforce uses various data centres located throughout the world. The location of Salesforce data centres for all services is specified in the Salesforce Infrastructure and Subprocessors Documentation, available under the following link - <https://www.salesforce.com/company/legal/trust-and-compliance-documentation/>

14. Performance

Full details are contained in our terms.

14.1. Service Levels

We will use commercially reasonable efforts to make the services available 24 hours a day, 7 days a week, except for:

- Planned downtime (see section 13).
- Any unavailability caused by circumstances beyond our reasonable control, including, for example, an act of God, act of government, flood, fire, earthquake, civil unrest, act of terror, strike or other labour problem (other than one involving our employees), Internet service provider failure or delay, non-Salesforce application, or denial of service attack.

14.2. Incident Response Time

The incident response time varies from 15 minutes to two days depending on the level of support selected. Details can be found in section 8

14.3. Incident Updates

Incident update intervals vary based on the level of support selected, ranging from every 30 minutes to longer periods. Details can be found in section 8.

THE CONTENT OF THIS DOCUMENT IS FOR INFORMATIONAL PURPOSES ONLY AND A CUSTOMER'S RIGHTS AND REMEDIES ARE GOVERNED EXCLUSIVELY BY THE G-CLOUD 15 CALL-OFF CONTRACT AND SUPPLIER TERMS. ANY RESPONSE TIMES OR CUSTOMER OUTCOMES STATED HEREIN ARE INDICATIVE ONLY.