



SALESFORCE AGENTFORCE

Quick Start

One agent, one job, done properly — and a clear view of whether it should go further.

SCOPE 1

Pilot

\$5,000

3 weeks

SCOPE 2

Most popular

Operational

\$13,500

5 weeks

SCOPE 3

Advanced

From \$22,000

Quoted after readiness

Start narrow, prove it, then widen

An agent is only as good as the data and content it reasons over, and the fastest way to lose confidence in AI is to deploy it broadly before it is reliable. Pilot proves one use case in a controlled setting. Operational takes it into daily use with grounding and governance. Advanced widens the scope, but only after a readiness assessment confirms the data can support it.

Brysa Limited

satish.thiagarajan@brysa.ai · +44 (0) 7867 865 102 · brysa.ai

What you will have at launch

A Brysa Agentforce Quick Start delivers a working agent scoped to one clearly defined job, tested against real scenarios, with explicit boundaries on what it will and will not attempt.

One defined job

An agent with a narrow, testable purpose.

Grounded answers

Responses drawn from sources you have approved.

Clear boundaries

Agreed topics it handles and where it hands off.

A safe escalation

A defined route to a human when confidence is low.

Tested behaviour

Evaluated against real scenarios before go-live.

Governed operation

Agreed ownership, monitoring and change process.

Included in every scope

- Kick-off, use-case definition and confirmation of the selected package scope
- Agent configuration with agreed topics, instructions and boundaries
- Escalation and handoff design so users always reach a human when needed
- Scenario-based testing against agreed acceptance criteria
- Brysa testing, coordinated client UAT and production deployment
- Administrator and user enablement appropriate to the selected scope
- Remote post-launch hypercare for the stated period and hour allowance
- A defined route into ongoing support through a Brysa Care Plan

STEP 0 · REQUIRED WHERE COMPLEXITY IS UNCERTAIN**Agentforce Data Readiness**

Where data quality or grounding content is uncertain, Brysa begins with a Data Readiness assessment. It reviews the objects, fields and grounding sources the agent would depend on and produces a readiness score, a remediation backlog and a firm fixed price for the build. The full fee is credited against any Agentforce package started within 60 days. It is required before Advanced, and recommended before Operational where content has not been reviewed.

\$7,500

2 weeks

Credited on proceed

Choose your scope

Start with the tier that matches how confident you are in your data and how far you want the agent to reach. Advanced is quoted only after a readiness assessment.

CAPABILITY	SCOPE 1 · PILOT	SCOPE 2 · OPERATIONAL	SCOPE 3 · ADVANCED
Fixed price	\$5,000	\$13,500	From \$22,000
Best fit	Proving one use case in a controlled setting	Taking a proven use case into daily operation	Customer-facing or multi-topic deployment
Data Readiness required	Recommended	Recommended	Required
Agents	1	1	Confirmed after readiness
Topics	Up to 2	Up to 3	Confirmed after readiness
Standard actions	Up to 2	Up to 4	Confirmed after readiness
Custom actions	1 simple action	Up to 5 actions	Confirmed after readiness
Knowledge grounding	Not included	Included, up to 3 source categories	Extended grounding
Deployment channel	Internal or controlled group only	Internal, or one employee-facing channel	Customer-facing channel available
External integrations	Not included	Not included	Assessed at readiness
Escalation and handoff	Basic handoff to a human	Defined escalation with agreed triggers	Extended escalation model
Testing	Up to 20 test scenarios	Up to 50 test scenarios	Confirmed after readiness
Guardrails and boundaries	Agreed topic boundaries	Topic boundaries and response guardrails	Extended guardrail model
Monitoring and reporting	Not included	4 reports and 1 dashboard	Confirmed after readiness
Governance handover	Not included	Operating and change guide	Full governance model
Training	Admin and user sessions	Admin, owner and user sessions	Expanded enablement
UAT	1 cycle	1 cycle plus retest	2 cycles plus regression retest
Hypercare	10 days, up to 3 hours	21 days, up to 6 hours	30 days, up to 8 hours
Typical delivery	3 weeks	5 weeks	Quoted after readiness

All limits are maximums within the fixed-price package and are confirmed in the configuration workbook after discovery.

What Brysa will not promise

No implementation partner can guarantee the accuracy of a generative agent, because accuracy depends on data, content and the questions users actually ask. What Brysa does commit to is a defined scope, tested behaviour against agreed scenarios, explicit boundaries, a working escalation path to a human, and an honest assessment of whether your data can support the use case at all.

SCOPE 1

Pilot

\$5,000

For proving one use case with a controlled group, before committing to a wider deployment.

Typical delivery: 3 weeks

WORKSTREAM	WHAT BRYSA WILL DELIVER	INCLUDED LIMIT
Use-case definition	— Kick-off and one workshop of up to 3 hours — Confirm one use case, the audience, the questions in scope and what good looks like — Agreed acceptance criteria and configuration workbook	1 use case 3 workshop hours
Agent configuration	— One agent configured with up to 2 topics — Agent instructions, tone and topic boundaries agreed and documented — Up to 2 standard actions and 1 simple custom action	1 agent 2 topics
Boundaries and escalation	— Explicit statement of what the agent will not attempt — Basic handoff to a human where a question falls outside scope — Fallback messaging for low-confidence responses	1 handoff route
Testing	— Up to 20 agreed test scenarios covering expected and edge-case questions — Documented results with observed behaviour for each scenario — One correction pass following the test results	20 scenarios
Deployment	— Deployment to an internal or controlled user group only — No customer-facing channel at this tier — Launch checklist and rollback position	Controlled group
Enablement	— One 2-hour administrator session covering configuration and change — One 60-minute user session covering effective use and limitations	2 sessions
Test and launch	— Brysa testing and one coordinated client UAT cycle — Production enablement for the controlled group	1 UAT cycle
Post-launch support	— Remote support for questions, defects and minor corrections within the agreed scope — Continues seamlessly into a Care Plan where selected	10 days 3 hours

Operational

\$13,500

For taking a proven use case into daily operation, with grounding, monitoring and a governance handover.

Typical delivery: 5 weeks

WORKSTREAM	WHAT BRYSA WILL DELIVER	INCLUDED LIMIT
Use-case definition	— Kick-off and up to 2 workshops totalling 6 hours — Confirm the use case, audience, escalation expectations and success measures — Agreed acceptance criteria, guardrails and configuration workbook	1 use case 6 workshop hours
Agent configuration	— One agent configured with up to 3 topics — Up to 4 standard actions and up to 5 actions in total including custom actions — Agent instructions, tone, guardrails and topic boundaries documented	1 agent 3 topics
Knowledge grounding	— Grounding configured across up to 3 approved source categories — Source precedence agreed where content overlaps or conflicts — Review of coverage gaps that would produce unreliable answers	3 source categories
Boundaries and escalation	— Defined escalation model with agreed triggers and routing — Response guardrails covering tone, scope and prohibited topics — Documented fallback behaviour	Defined escalation
Testing	— Up to 50 agreed test scenarios covering expected, edge-case and adversarial questions — Documented results and one correction and retest period	50 scenarios
Monitoring and governance	— Configure 4 reports and 1 dashboard covering usage, escalation and deflection — Operating and change guide covering ownership, review cadence and how to adjust the agent safely	4 reports 1 dashboard
Test, train and launch	— Brysa testing, one client UAT cycle and one correction and retest period — Deployment to internal use or one employee-facing channel — Administrator, owner and user enablement totalling up to 5 hours	1 UAT cycle 5 training hours
Post-launch support	— Remote support for questions, defects and minor corrections within the agreed scope — Continues seamlessly into a Care Plan where selected	21 days 6 hours

SCOPE 3

Advanced

For customer-facing deployment or wider topic coverage. Requires a completed Data Readiness assessment before pricing.

From \$22,000

Typical delivery: Quoted after readiness

WORKSTREAM	WHAT BRYSA WILL DELIVER	INCLUDED LIMIT
Readiness prerequisite	— Agentforce Data Readiness assessment completed at \$7,500, credited against the build — Readiness score, remediation backlog and firm fixed price for the implementation — Where readiness is insufficient, remediation is completed before the build begins	Required \$7,500 credited
Agent configuration	— Agent scope, topics and actions confirmed by the readiness findings — Extended instruction and guardrail model for a broader audience — Data access and sharing behaviour reviewed for external exposure	Confirmed at readiness
Knowledge grounding	— Extended grounding across the agreed source categories — Source precedence, currency and retirement rules documented	Confirmed at readiness
Boundaries and escalation	— Extended guardrail model appropriate to a customer-facing channel — Escalation routing to the right human team with agreed service expectations	Extended model
Testing	— Scenario coverage confirmed at readiness, including adversarial and safety testing — Two UAT cycles with correction and regression retest	2 UAT cycles
Governance	— Full governance model covering ownership, monitoring, review cadence and change control — Administration handover guide and up to 8 hours of enablement	8 training hours
Post-launch support	— Remote support for questions, defects and minor corrections within the agreed scope — Continues seamlessly into a Care Plan where selected	30 days 8 hours

Add-ons and ongoing care

Where an agent needs to reach a little further, these items can be added at a fixed price and confirmed in the Statement of Work.

Fixed-price add-ons

ADD-ON	FIXED PRICE	TYPICAL EFFECT ON TIMELINE
Additional topic	\$1,900	+2 to 3 days
Additional grounding source category	\$1,300	+1 to 2 days
Additional custom action	\$1,000	+1 day
Monitoring report pack (4 reports, 1 dashboard)	\$1,000	+1 to 2 days
Additional 25 test scenarios	\$1,000	+1 day
Additional standard action	\$700	+0.5 day
Additional training session (90 minutes)	\$700	No change
External system integration for an action	Quoted separately	Assessed first

Add-on prices are the same whether agreed at Statement of Work stage or raised during delivery, and are invoiced with the second instalment.

Brysa Care Plans

Agents need tending. Questions change, content goes stale, and behaviour drifts as topics widen. A Care Plan begins the day hypercare ends and gives you a route to adjust instructions, refresh grounding and review escalation. For any agent in daily use we recommend Partner.

	ESSENTIALS	ACTIVE	PARTNER
Monthly fee	\$750	\$1,600	\$3,000
Included hours	Up to 3 hours	Up to 7 hours	Up to 15 hours
Reactive support	Yes	Yes	Yes, priority
Small changes and fixes	Yes	Yes	Yes
Scheduled change window	Not included	Monthly	Fortnightly
Business review	Not included	Quarterly	Monthly
Roadmap and optimisation input	Not included	Not included	Included
Named Brysa administrator	Not included	Not included	Included

Commit at Statement of Work stage

Where a Care Plan is agreed alongside the Quick Start, Brysa applies a 10% discount to the first six months. Care Plans run on a rolling basis with 30 days' notice, and unused hours do not carry forward.

How the Quick Start is delivered

Brysa uses a structured delivery sequence so the agent is scoped narrowly, tested honestly and deployed with a clear route back to a human.

1	2	3	4	5
Define	Ground	Configure	Test	Deploy
Agree one use case and what a good answer looks like.	Confirm the sources the agent may reason over.	Build topics, actions, guardrails and escalation.	Run agreed scenarios and correct observed behaviour.	Release to the agreed audience and begin hypercare.

Typical timetable

PERIOD	PILOT	OPERATIONAL	ADVANCED
Week 1	Define and configure	Define and ground	Readiness assessment
Week 2	Test, correct and UAT	Configure topics and actions	Remediation where required
Week 3	Train, deploy and hypercare	Grounding, guardrails and escalation	Build phase, per confirmed plan
Week 4	Not applicable	Scenario testing and correction	Build phase, per confirmed plan
Week 5 onwards	Not applicable	UAT, governance, train and deploy	Testing, UAT, governance and deployment

When the clock starts

The delivery period begins once the Statement of Work is signed, the required Agentforce and Salesforce licences and access are available, the grounding sources have been identified and approved for use, the nominated decision-maker is confirmed and, where required, the Data Readiness assessment has been completed. Client delays move the delivery dates accordingly.

Decision and feedback cadence

- Brysa records all agreed decisions in the configuration workbook.
- The client provides one consolidated response at each review point.
- Feedback and approvals are expected within two working days unless otherwise agreed.
- New requirements raised during UAT are assessed as additional scope and priced from the add-on menu where applicable.

Definitions and responsibilities

These definitions matter because agent scope is easy to describe loosely and expensive to deliver loosely.

TERM	WHAT IT MEANS IN THIS PACKAGE
Agent	One configured Agentforce agent with a defined purpose, audience and set of topics.
Topic	One area of responsibility with its own instructions, scope boundaries and associated actions.
Standard action	One action available from the Salesforce-provided library, configured against agreed objects and fields.
Custom action	One action built declaratively using Flow or standard configuration. It does not include Apex, external callouts or third-party integration.
Grounding source category	One class of approved content the agent may reason over, such as Knowledge articles or a defined set of records.
Test scenario	One documented question or interaction with an expected outcome, used to evaluate agent behaviour before release.
Guardrail	One agreed constraint on what the agent may say or do, expressed in its instructions and verified by testing.
Escalation	A defined route by which a user reaches a human, including the trigger conditions and the receiving team.

What Brysa needs from you

- Appropriate Agentforce and Salesforce licences and administrator access
- A nominated business owner for the agent, not only a technical contact
- Approved grounding content that is current and free of known contradictions
- A written description of the use case and the questions in scope
- Agreement on what the agent must never attempt, confirmed in writing
- A named human team to receive escalations, with agreed response expectations
- Named users to carry out UAT and provide one consolidated issue list
- Acceptance that agent accuracy cannot be guaranteed and requires ongoing tending

What is not included

The following work sits outside all three standard scopes unless it is expressly added to the Statement of Work. Where an item is needed, Brysa will scope and price it separately.

- Agentforce, Salesforce, Data Cloud or third-party licences and subscription fees
- Bulk data cleansing, deduplication, enrichment or migration
- Knowledge article authoring, rewriting, consolidation or translation
- External system integrations, middleware or API development for actions
- Apex, Lightning Web Components or other bespoke code
- Any guarantee of agent accuracy, response quality or deflection rate
- Model selection, fine-tuning or custom evaluation harness construction
- Voice channels, telephony or third-party messaging platforms requiring separate licensing
- Ongoing monitoring, tuning or content maintenance after hypercare ends
- Full change management, adoption programmes or additional UAT cycles

Support, defects and additional requirements

AREA	HOW IT IS HANDLED
Included hypercare	Remote support during UK business hours for questions, agreed-scope defects and minor corrections, limited to the package allowance and period.
Defect	An agreed item does not operate in accordance with the approved configuration workbook or acceptance criteria. Defects are corrected at no additional cost.
Enhancement	A new feature, changed process, additional report, new automation or other requirement not approved in the package scope.
Additional scope	Brysa will explain the impact and provide a fixed price from the add-on menu or a separate estimate. No chargeable work begins without written approval.
Client delay or pause	Delivery dates move where information, access, data, testing or decisions are delayed. A project paused for more than 20 working days may require re-planning.

Commercial terms and next steps

TERMS	
Fixed package prices	Pilot \$5,000 · Operational \$13,500 · Advanced from \$22,000
Data Readiness	\$7,500, credited in full against any Agentforce package commenced within 60 days. Required before Advanced.
Payment	50% on commencement and 50% on UAT sign-off. Add-ons are invoiced with the second instalment.
Care Plans	Billed monthly in advance from the day hypercare ends. Rolling term with 30 days' notice.
Accuracy	Brysa commits to defined scope, tested behaviour against agreed scenarios and a working escalation path. Brysa does not warrant the accuracy of generative responses.
Acceptance	The solution is accepted when the agreed test scenarios pass and material in-scope defects are resolved, or five working days after deployment if no material defect has been reported.
Validity	All prices are in US dollars and are the total professional services fee for the agreed scope. Valid for 60 days from the date of issue.

Select the right starting point

Pilot

Choose this when you want to prove one use case with a controlled group first.

Operational

Choose this when a proven use case needs grounding, monitoring and daily operation.

Advanced

Choose this when the agent will face customers or cover a wider set of topics.

Let us confirm the right starting point

Every Agentforce engagement begins with a short, no-obligation call. It takes about thirty minutes and confirms three things: which use case is worth doing first, whether a Data Readiness assessment is needed before pricing, and what we would each need to have ready before the clock starts.

Next step

Brysa will hold a short suitability call to confirm the use case, whether readiness assessment is required and the access needed. We will then issue the Statement of Work and agree the kick-off date.

satish.thiagarajan@brysa.ai · [+44 \(0\) 7867 865 102](tel:+44207867865102) · brysa.ai



Satish Thiagarajan

Chief Executive Officer, Brysa Limited

This document describes Brysa's standard Agentforce Quick Start packages. The signed Statement of Work governs the engagement.