



AGENTFORCE DATA READINESS

Readiness

Find out whether your data can actually support an AI agent — before you build one.

TIER 1

Snapshot

\$4,500

1 week

TIER 2

Most popular

Assessment

\$7,500

2 weeks

TIER 3

Remediation

\$12,000

3 weeks

An honest answer, not a sales exercise

Agentforce is only as good as the data and content it reasons over. These engagements tell you where you actually stand, what would need to change, and what an implementation would realistically cost. A finding that you are not ready yet is a valid and valuable outcome — it is considerably cheaper than discovering it halfway through a build.

Brysa Limited

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

What you will have at the end

A Brysa Agentforce Data Readiness engagement examines the objects, fields and grounding sources an agent would depend on, and reports plainly on whether they can support reliable answers.

A readiness score

A clear rating across data, content and configuration.

A risk summary

The specific things that would cause an agent to answer badly.

A grounding review

Whether your knowledge and records are usable as sources.

A remediation backlog

Prioritised work needed before an agent is viable.

A recommended next step

Proceed, remediate first, or do not proceed yet.

An implementation estimate

A firm price for the build, if the answer is proceed.

Included in every tier

- Kick-off and confirmation of the use case under consideration
- Structured review against the Brysa readiness framework
- A written findings note in plain language, not a raw data dump
- A readiness score and prioritised risk summary
- A recommended next step and, where appropriate, a firm fixed price for implementation
- A findings walkthrough with your nominated stakeholders

These engagements never change your org

A readiness engagement is analysis. It does not cleanse data, rewrite content or reconfigure your org, with the single exception of the limited configuration corrections included in the Assessment and Remediation tiers, which are always agreed in writing first.

Choose your depth

Start with the tier that matches how much you already know. Snapshot answers whether it is worth looking further; Assessment is the standard engagement; Remediation adds hands-on correction and a retest.

CAPABILITY	TIER 1 · SNAPSHOT	TIER 2 · ASSESSMENT	TIER 3 · REMEDiation
Fixed price	\$4,500	\$7,500	\$12,000
Best fit	Deciding whether Agentforce is worth exploring	Committed to exploring, need to know where you stand	Know there are gaps, want them closed and retested
Use-case workshop	1 hour	Up to 3 hours	Up to 5 hours
Objects reviewed	Up to 3	Up to 5	Up to 8
Fields reviewed	Up to 15	Up to 25	Up to 40
Grounding source categories	Not included	Up to 3	Up to 5
Data quality profiling	Headline completeness only	Completeness, duplication and consistency	Full profiling with sample-level evidence
Readiness score	Included	Included	Included, plus post-remediation retest
Risk summary	Included	Included	Included
Remediation backlog	Not included	Prioritised backlog	Prioritised backlog with effort estimates
Configuration corrections	Not included	Up to 5 minor corrections	Up to 15 minor corrections
Data standards workshop	Not included	Not included	Included
Implementation estimate	Indicative range	Firm fixed price	Firm fixed price
Findings walkthrough	1 session	1 session	2 sessions
Follow-up period	5 working days	10 working days, up to 3 hours	20 working days, up to 5 hours
Typical duration	1 week	2 weeks	3 weeks

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

Audit is not cleansing

Configuration corrections mean small, agreed fixes such as a missing field description, an inconsistent picklist or an unpopulated help text. They do not mean bulk data cleansing, deduplication, enrichment or content rewriting. Where that work is needed, it is scoped as a separate engagement once the backlog is known.

TIER 1

Snapshot

\$4,500

A fast, focused look at whether your data could support the use case you have in mind, and whether a fuller assessment is worth doing.

Typical delivery: 1 week

WORKSTREAM	WHAT BRYSA WILL DELIVER	INCLUDED LIMIT
Use-case definition	— One-hour workshop to confirm the agent use case under consideration — Identify the objects and audiences the agent would need to serve — Agree what a good answer would look like	1 use case 1 hour
Data review	— Review of up to 3 objects and up to 15 fields relevant to the use case — Headline completeness profiling across those fields — Identification of obvious structural blockers	3 objects 15 fields
Findings	— Readiness score across data, content and configuration — Risk summary listing the specific issues that would degrade agent answers — Recommended next step and an indicative implementation cost range	1 findings note
Walkthrough and follow-up	— One findings walkthrough with your nominated stakeholders — Five working days of remote follow-up for questions arising from the findings	1 session 5 working days

Assessment

\$7,500

The standard engagement. A full readiness picture across data, grounding sources and configuration, ending in a firm price for the build.

Typical delivery: 2 weeks

WORKSTREAM	WHAT BRYSA WILL DELIVER	INCLUDED LIMIT
Use-case definition	— Workshops totalling up to 3 hours across the relevant business and technical stakeholders — Confirm the agent use case, audiences, escalation expectations and success measures — Agree the boundaries of what the agent should and should not attempt	1 use case 3 hours
Data review	— Review of up to 5 objects and up to 25 fields — Profiling for completeness, duplication and consistency — Assessment of ownership, record access and sharing implications for agent responses	5 objects 25 fields
Grounding source review	— Review of up to 3 grounding source categories, such as Knowledge, files or structured records — Assessment of coverage, currency, duplication and contradiction between sources — Identification of content that would produce unreliable or conflicting answers	3 source categories
Configuration corrections	— Up to 5 minor configuration corrections agreed in writing before being applied — Typically field descriptions, help text, picklist consistency or metadata labelling	5 corrections
Findings	— Readiness score and prioritised risk summary — Remediation backlog ordered by impact on answer quality — Firm fixed price for the Agentforce implementation where the recommendation is to proceed	1 findings note 1 backlog
Walkthrough and follow-up	— One findings walkthrough with your nominated stakeholders — Ten working days of remote follow-up, up to 3 hours	1 session 10 working days

Remediation

\$12,000

For organisations that already know there are gaps and want the highest-impact ones closed, then measured again.

Typical delivery: 3 weeks

WORKSTREAM	WHAT BRYSA WILL DELIVER	INCLUDED LIMIT
Use-case definition	— Workshops totalling up to 5 hours across business and technical stakeholders — Confirm the use case, audiences, escalation expectations and success measures — Agree remediation priorities and what will be corrected within the engagement	1 use case 5 hours
Data review	— Review of up to 8 objects and up to 40 fields — Full profiling with sample-level evidence for each finding — Assessment of ownership, sharing and access implications for agent responses	8 objects 40 fields
Grounding source review	— Review of up to 5 grounding source categories — Coverage, currency, duplication and contradiction analysis with worked examples — Recommendations on source precedence and retirement	5 source categories
Configuration corrections	— Up to 15 minor configuration corrections agreed in writing before being applied — Correction log recording what was changed, when and why	15 corrections
Data standards workshop	— One workshop establishing the standards needed to keep data agent-ready — Ownership, maintenance cadence and definition of done for each standard — Written standards summary for internal adoption	1 workshop
Findings and retest	— Prioritised remediation backlog with effort estimates — Updated readiness score following the corrections applied — Firm fixed price for the Agentforce implementation where the recommendation is to proceed	1 retest 1 updated score
Walkthrough and follow-up	— Two findings walkthroughs, including one with senior stakeholders — Twenty working days of remote follow-up, up to 5 hours	2 sessions 20 working days

Add-ons and ongoing care

Where the scope needs to stretch, 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 use case reviewed	\$1,900	+2 to 3 days
Additional grounding source category	\$1,300	+1 to 2 days
Additional 5 objects reviewed	\$1,000	+1 day
Additional 5 configuration corrections	\$1,000	+1 day
Executive briefing session	\$700	No change
Bulk data cleansing or content rewriting	Quoted separately	Scoped after findings

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

How the readiness engagement runs

Brysa follows a fixed sequence so the engagement stays within its scope and you receive a usable answer rather than an open-ended investigation.

1	2	3	4	5
Frame	Examine	Score	Report	Walk through
Confirm the use case and what a good answer looks like.	Review objects, fields and grounding sources against the framework.	Rate readiness and identify the risks to answer quality.	Produce findings, backlog and a recommended next step.	Present the findings and agree what happens next.

Typical timetable

PERIOD	SNAPSHOT	ASSESSMENT	REMEDIATION
Week 1	Frame, examine, score and report	Frame and examine	Frame, examine and profile
Week 2	Not applicable	Score, correct, report and walk through	Score, corrections and standards workshop
Week 3	Not applicable	Not applicable	Retest, report and walkthroughs
Follow-up	5 working days	10 working days	20 working days

When the clock starts

The engagement begins once the Statement of Work is signed, read access to the relevant Salesforce org and grounding sources is available, the nominated decision-maker is confirmed and the use case under consideration has been described in writing. Client delays move the dates accordingly.

Decision and feedback cadence

- Brysa records all findings and evidence in the readiness workbook.
- Configuration corrections are agreed in writing before any change is applied.
- The client provides one consolidated response to the draft findings.
- Requests beyond the agreed scope are priced from the add-on menu where applicable.

Definitions and responsibilities

These definitions keep the fixed price clear and, just as importantly, make clear what an assessment is not.

TERM	WHAT IT MEANS IN THIS PACKAGE
Readiness score	A rating across data quality, grounding content and configuration, produced using the Brysa readiness framework and supported by evidence.
Grounding source category	One class of content an agent would reason over — for example Knowledge articles, uploaded files, or a defined set of structured records.
Configuration correction	One small, agreed metadata change such as a field description, help text, picklist value or label. It is not a data change.
Remediation backlog	A prioritised list of the work required to reach an acceptable readiness score, ordered by impact on answer quality.
Data cleansing	Bulk correction, deduplication or enrichment of record data. Never included in a readiness engagement at any tier.
Implementation estimate	A firm fixed price for a subsequent Agentforce implementation, valid for 60 days and based on the findings.
Follow-up period	A stated window after the walkthrough during which questions about the findings are answered remotely within the hour allowance.

What Brysa needs from you

- Read access to the relevant Salesforce org and grounding sources
- A nominated decision-maker empowered to approve scope and corrections
- A written description of the agent use case under consideration
- Availability of the business owner for the agreed workshops
- Access to a sandbox where configuration corrections can be applied first
- Written approval before any configuration correction is applied in production
- Named stakeholders to attend the findings walkthrough
- One consolidated response to the draft findings

What is not included

A readiness engagement is analysis. The following sits outside all three tiers unless expressly added to the Statement of Work.

- Agentforce, Salesforce, Data Cloud or third-party licences and subscription fees
- Any Agentforce implementation, agent build, topic or action configuration
- Bulk data cleansing, deduplication, enrichment or migration
- Knowledge article authoring, rewriting, consolidation or translation
- Prompt engineering, model tuning or evaluation harness construction
- Integration work, middleware or API development to make a source available
- Apex, Lightning Web Components or other bespoke code
- Security, sharing or permission redesign beyond identifying the risk
- Ongoing monitoring of data quality after the engagement ends
- Any guarantee of agent accuracy, which depends on factors beyond data readiness

Support, defects and additional requirements

AREA	HOW IT IS HANDLED
Included follow-up	Remote support during UK business hours for questions arising from the findings, limited to the stated allowance and period.
Clarification	A question about what the findings mean or how a recommendation was reached. Answered at no additional cost within the follow-up period.
Additional analysis	A request to review objects, sources or processes beyond the agreed scope. Priced from the add-on menu or quoted separately.
Remediation	Any change to configuration or data. Never included in an assessment; delivered as a separate engagement or through the Remediation tier.
Client delay or pause	Delivery dates move where access, information or stakeholder availability is delayed. An engagement paused for more than 20 working days may require re-planning.

Commercial terms and next steps

TERMS	
Fixed prices	Snapshot \$4,500 · Assessment \$7,500 · Remediation \$12,000
Credit on proceed	The fee is credited in full against an Agentforce implementation commenced with Brysa within 60 days of the findings being issued.
Payment	50% on commencement and 50% on issue of the findings note.
Deliverable	The deliverable is the findings note, readiness score and recommended next step. A finding that you are not yet ready is a complete deliverable.
Corrections	Configuration corrections are applied only with written approval and are recorded in a correction log.
Acceptance	The engagement is accepted on delivery of the findings walkthrough, or five working days after the findings note is issued if no material issue has been raised.
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

Snapshot Choose this when you are still deciding whether Agentforce is worth exploring at all.	Assessment Choose this when you are committed to exploring and need a firm implementation price.	Remediation Choose this when you already know there are gaps and want them closed and retested.
--	--	---

NEXT STEP

Let us confirm the right starting point

Every readiness engagement begins with a short, no-obligation call. It takes about thirty minutes and confirms three things: which tier fits how much you already know, which use case is worth assessing first, and what access we would need before the clock starts.

Next step

Brysa will hold a short suitability call to confirm the tier, the use case and the access required. We will then issue the Statement of Work and agree the start 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 Data Readiness packages. The signed Statement of Work governs the engagement.