RDi Assessment Agent

Run an AI and data governance assessment in your own environment, with RDi expert review at the points that need judgement.

Partner brief, version 2, 25 November 2026. Finance and tax reviewed (D132).

What it is

RDi Assessment Agent is a subscription product. It scores an organization against the RDi-Data Product Maturity Model (DPMM), using four domains and four maturity levels, then produces reports and a growth plan. Your team runs it as a Docker image on a laptop or in your own cloud account. Your data stays on your systems, and RDi consultants guide the work and review the results.

How it works

1. Register your Anthropic API key once, through the Setup screen or as a Docker secret. RDi never receives the key.
2. Read the data disclosure and record consent. No evidence leaves your environment before consent.
3. Answer RDi questions in sequence, grouped by domain. Each question shows the evidence it expects.
4. Upload or point to files, PDFs, reports and exported emails. Every answer links to its source passage.
5. Missing or weak evidence triggers a re-ask that names what would satisfy it, up to three asks per question.
6. An assessment lead can accept, reject or annotate any automatic result. Each override is logged with a reason.
7. The sponsor approves the scored assessment and growth plan. A quarterly health check and an annual refresh follow.

What you get

Output	What it contains	Who uses it
Completion, progress and maturity views	Share of questions answered, change since last assessment, and level per domain	Assessment lead and sponsor
Growth plan	Draft targets per domain, paced by capacity and bounded by cultural limits	Sponsor and contributors
PDF report	Branded from your branding pack; states the framework version and evidence references	Leadership and distribution
Data export	CSV or JSON for your own analysis	Analysts
Audit history	Every change, override and consultant action, logged	Auditors and governance teams

Who it is for

Audience	Profile	Tier
Small organizations and independent data leads	Up to 20 systems, one primary user	Individual
Midsize organizations	21 to 75 systems, a team of contributors and an assessment lead	Team
Large or multi-entity organizations	76 or more systems, or several entities, often in regulated sectors	Enterprise (quoted)
Consultants and partners	Advisers who run assessments for their clients	Partner programme: partner@rdi-data.com

Pricing (USD)

Tier	Scope	Licence, annual	Licence, monthly	Review, per year	Year one (licence and review)
Individual	1 user, up to 20 systems	$2,250	$243.75	17.5 hours, $4,375	$6,625
Team	Up to 10 users, 21 to 75 systems	$12,000	$1,300	80 hours, $20,000	$32,000
Enterprise	76 or more systems, or several entities	From $75,000 per year, quoted per customer	Not offered	Quoted per scope	Quoted

Multi-year terms (Team and Enterprise licences). Each additional contract year is 5% below the annual price, up to five years. For two years or more, 10% of the licence fee is opex and 90% is capex. Individual licences renew each year and have no multi-year terms or opex split.

Term	Team licence total	Team opex (10%)	Team capex (90%)	Enterprise total from	Enterprise opex (10%)	Enterprise capex (90%)
1 year	$12,000	not split	not split	$75,000	not split	not split
2 years	$23,400	$2,340	$21,060	$146,250	$14,625	$131,625
3 years	$34,200	$3,420	$30,780	$213,750	$21,375	$192,375
5 years	$54,000	$5,400	$48,600	$337,500	$33,750	$303,750

Multi-year totals are licence fees only. Review services are not included.

Consultant time is billed at a $250 blended hourly rate, with hour blocks available. Review hours are allotted each quarter, and unused hours expire at the end of the quarter. Enterprise review services are quoted. Prices are indicative in USD and subject to a signed agreement.

Models and key ownership

The assessment runs on Anthropic models, using your own Anthropic account. You pay Anthropic for model use and hold the key. Evidence and answers go to Anthropic under your account, and Anthropic's terms apply. Other model providers are not supported in this version. By default, scoring and grading are performed by you, aided by the model. The models apply the scoring criteria only to the evidence files you choose to analyze, and they do not examine the contents of your files unless you prompt them to.

Next step

To arrange a discovery call, visit rdi-data.com. Enterprise customers can request a quote through the RDi-Data customer portal. Interested partners can contact partner@rdi-data.com.