

Template · Free resource

AI Vendor & Tool Evaluation Scorecard

Demos are designed to impress, not to inform. This weighted scorecard compares AI vendors on the things that actually matter: data handling, security, residency, value and exit terms, so you can put two slick pitches...

Updated May 2026

<https://rangefrontlabs.com.au/resources/ai-vendor-evaluation-scorecard/>

Built in Toowoomba. Working across Australia and internationally.



Every AI vendor demos beautifully. That’s the job of a demo. The trouble starts later: when you discover the tool trains on your prompts, stores data in a region you can’t name, or quietly triples in price once you’re dependent on it.

This scorecard puts vendors side by side on the criteria that predict those problems, weighted by how much they actually matter. Score each option from 1 to 5, and the sheet gives you a single weighted number per vendor, so the comparison is about substance, not whoever gave the smoothest pitch.

This is the working tool that goes with our [Secure AI Buyer’s Guide](#). The guide explains the *why* behind each criterion; this scorecard is the *how*.

The criteria

Score each vendor 1 (poor) to 5 (excellent) on each line. The weights add to 100%, so the result lands back on a 1–5 scale you can compare directly.

#	Criterion	A “5” looks like	Weight
1	No training on your data	Contractual guarantee they never train on your inputs	12%
2	Retention & deletion control	You set retention, can purge to zero, deletion confirmed	10%
3	Data residency & sovereignty	Data stays in a region you choose, under laws you accept	12%
4	Sub-processor transparency	Full, current list of who else touches your data	6%
5	Security assurance	SOC 2 Type II or ISO 27001, recent pen test, encryption everywhere	12%
6	Access control & audit logs	SSO, role-based access, exportable audit trail	8%
7	Deployment fits your data	A model (SaaS / private / on-prem) that suits your most sensitive data	10%
8	Measurable value	A specific, provable outcome, not “boosts productivity”	12%
9	Total cost clarity	All-in pricing, no surprise usage cliffs	8%
10	Exit & portability	Clean export, confirmed deletion, no lock-in	10%

How the score works

The weighted total is a SUMPRODUCT of the weights and your 1–5 ratings:

$$\text{Weighted score} = \sum (\text{criterion weight} \times \text{your rating})$$

The downloaded sheet does this automatically for up to three vendors in adjacent columns. Type your ratings; the totals at the bottom update and rank themselves.

The dealbreakers

A weighted average is the right tool for comparing *broadly acceptable* options. It is the wrong tool for hiding a fatal flaw. Some answers are pass/fail no matter how good the rest of the score is:

- **It trains on your restricted or regulated data** with no opt-out. Stop.
- **It can’t tell you where the data is processed.** Stop, for anything sensitive.
- **You can’t get your data out or have it deleted.** Stop. That’s a hostage situation, not a vendor.

Mark these on the sheet. A 4.5 average means nothing if criterion 1 is a hard no for your data.

How to use it

1. **Score from evidence, not the demo.** A 5 on security needs the certificate and the pen-test summary, not a reassuring slide.
2. **Score per use case.** The right answer for drafting marketing copy and for processing health records is rarely the same tool. Run the sheet once per workload.
3. **Get the answers in writing.** If a vendor won't commit an answer to email or the contract, score it low: vagueness is a finding.
4. **Compare the totals, then sanity-check the dealbreakers.** Highest number wins only among options that clear every hard requirement.

A scorecard won't make the decision for you, and it shouldn't. What it does is make sure you decided on the things that matter rather than the things that demo well. When you want the awkward questions asked for you, that's what a [discovery call](#) is for.

Vendor scorecard

Criterion	Weight	Vendor A	Vendor B	Vendor C
1. No training on your data	0.12	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
2. Retention & deletion control	0.1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
3. Data residency & sovereignty	0.12	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
4. Sub-processor transparency	0.06	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
5. Security assurance (SOC2/ISO, pen test, encryption)	0.12	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6. Access control & audit logs	0.08	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
7. Deployment model fits your data sensitivity	0.1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
8. Measurable value / outcome	0.12	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
9. Total cost clarity (no hidden cliffs)	0.08	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
10. Exit & portability (export + deletion, no lock-in)	0.1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Dealbreaker check

- Trains on our restricted data with no opt-out?
- Cannot confirm where data is processed?
- Cannot export our data or have it deleted?

Need this adapted to your organisation, systems or data? Book a discovery call: <https://rangefrontlabs.com.au/contact/>