

THE DECISION INTELLIGENCE

Diligence Standard

Fifteen questions every enterprise and government buyer
should ask, of every vendor, including us.

How to use this standard

Every vendor in this category will tell you their AI is trusted, explainable, and enterprise-ready. Few can prove it. And the difference never shows up in a pitch. It shows up later, when a decision is questioned by a board, a regulator, an auditor, or mission oversight, and someone has to prove how it was reached.

These fifteen questions separate platforms that produce defensible decisions from platforms that produce confident answers.

Use it the right way: do not accept a description, require a demonstration. A real answer can be shown on your own data, in your own evaluation, not promised on a slide.

Evidence and Traceability

- 1. Can the platform produce a complete record of any decision, the evidence, the reasoning, and the rule that governed it, on your data?**

If what comes back is a dashboard or a summary, it is not a record.

- 2. Can every conclusion be traced to the exact place it came from, the document, the page, the field?**

"Grounded in your data" means nothing if you cannot follow the thread back.

- 3. Does the platform cite its sources, or ask you to trust its output?**

Trust is not a feature. Traceability is.

Reasoning and Explainability

- 4. Does it show how the answer was reached, the full reasoning chain, or only the answer?**

An unexplained recommendation quietly moves the risk onto you.

- 5. Does it reason across your connected data, or generate from a model's memory?**

Generation guesses. Reasoning traces.

- 6. Is the explainability method independently verifiable, patented, tested, documented, or simply asserted?**

Ask what actually backs the claim.

Conflict, Confidence, and Honesty

7. When two sources disagree, does the platform surface the conflict, or quietly pick one?

A conflict resolved in the dark is a conflict waiting to surface in an audit.

8. Is confidence calibrated to the evidence, with the basis shown, or is it an uncalibrated number?

A confidence score with no basis is decoration.

9. When the evidence is insufficient, does it say so, or answer anyway?

A system that always answers can never tell you when it does not know.

Governance, Audit, and Reproducibility

10. Are records immutable, and do corrections supersede rather than overwrite, with history preserved?

A decision you can quietly change is a decision you cannot defend.

11. Can a decision made today be reproduced and defended months or years from now, to an auditor, a regulator, or an inspector general?

Defensibility is not a moment. It is a permanent property.

12. Is the model and the data version recorded with every decision?

"The AI decided" is not an answer anyone will accept later.

Control, Deployment, and Accountability

13. Does your data stay in place under your governance, or must it be copied, moved, or sent to the vendor's cloud?

The safest data is the data that never leaves.

14. Can the platform deploy where you operate, on-premises, air-gapped, at the edge, or in your cloud, under zero-trust controls?

For classified and regulated work, this is not optional.

15. Where do people review and sign off, and is that human accountability captured in the record?

Decision Intelligence augments judgment. It does not replace accountability.

How to score it

A platform that can demonstrate all fifteen, on your data, in your environment, is a Decision Intelligence platform. A platform that can only describe them is a demo.

We publish this standard because reView is built to pass it. Run it on us, and on everyone else you are evaluating. The results will tell you who can defend a decision, and who can only produce one.

Run the Diligence Standard on reView. Request a meeting at data2.ai

The Scorecard

Score each vendor on your own data: **D** = Demonstrated · **S** = Only described · **N** = Not shown. A platform that demonstrates all fifteen is a Decision Intelligence platform.

#	What the vendor must demonstrate, on your data	V1	V2	V3
Evidence and Traceability				
1	Produce a full Decision Record on our data			
2	Trace every conclusion to its exact source			
3	Cite sources, not “trust us”			
Reasoning and Explainability				
4	Show the full reasoning chain, not just the answer			
5	Reason across our data, not generate from memory			
6	Explainability method is verifiable / patented			
Conflict, Confidence, and Honesty				
7	Surface conflicts, not silently resolve them			
8	Confidence calibrated, with its basis shown			
9	Abstains when the evidence is insufficient			
Governance, Audit, and Reproducibility				
10	Records immutable; corrections supersede			
11	Decision reproducible and defensible years later			
12	Model and data version recorded with the decision			
Control, Deployment, and Accountability				
13	Our data stays in place, under our governance			
14	Deploys on-prem, air-gapped, edge, cloud (zero-trust)			
15	Human review and sign-off captured in the record			