

Running the *Information Governance* Assessment

The assessment scores how deliberately a newsroom governs its information environment. This manual covers how to run it, how to read what it returns, and how every figure in the report is produced.

VERSION

1.0 · July 2026

APPLIES TO

Assessment v1.0 · 35 questions

AUDIENCE

Respondents · facilitators · boards

COMPANION

Information Governance position paper v1.0

The position paper makes the argument. The assessment measures it. This manual sits between them.

Sections 01–03 are enough to complete the assessment alone. Sections 04–06 are written for whoever runs it with a leadership team and carries the result into a meeting. The appendices document the scoring in full, so any figure in the report can be traced back to the answers that produced it.

Where this manual refers to **the paper**, it means **Governing the information environment of the AI-enabled newsroom** (v1.0, July 2026), which sets out the case and defines the seven domains this instrument measures.

CONTENTS	
01	What the assessment measures
02	Before you begin
03	Taking the assessment
04	Running it with a leadership team
05	Reading the report
06	After the report
A	Methodology & scoring
B	Domains, scales and levels
C	Questions that come up

01 What the assessment measures

It measures governance, not AI use.

Thirty-five questions across seven domains, five per domain, about five minutes per respondent. Each question asks about observable organisational behaviour — what exists, what is written down, who owns it, when it was last reviewed — rather than about intent, opinion or ambition.

That distinction sets the boundary. The assessment is not an AI-use policy check, not a compliance audit, and not a vendor evaluation. It does not measure how much AI a newsroom uses. A newsroom using very little AI can score well; an enthusiastic adopter can score badly. The question throughout is whether the information environment is mapped, owned and reviewed.

WHAT THE REPORT RETURNS

- A score of 0–100 for each of the seven domains, each mapped to a maturity level.
- An overall maturity score and level, weighted towards the weakest domain.
- Strengths, weaknesses and blind spots, named.
- Immediate priorities, a 90-day roadmap, a one-year roadmap and a five-item board agenda.
- Recommendations for all seven domains in three tiers: quick wins, medium-term, strategic.
- With more than one respondent: an Alignment Index measuring how far leadership agrees.

The seven domains and the five-level maturity model are taken from the paper unchanged. The paper's six questions for the board are the light entry point to the same framework; the 35 questions here are the full instrument, and they return a number the board can act on rather than a conversation it can defer.

It measures whether governance exists and is owned — not whether the newsroom is good at AI.

02 Before you begin

Answer for the organisation as it operates today, not as policy says it should.

Who answers

The tool offers six roles: editor-in-chief, publisher, CEO, CTO, editorial innovation lead, and board or supervisory member. Role is required; name and organisation are optional. The organisation field becomes the title of the report — the first respondent to fill it sets it for the whole session, and later respondents inherit it.

What is needed otherwise is a browser, roughly five minutes per respondent, and enough visibility of the organisation to answer on evidence. Where a respondent does not know, that is a finding rather than an obstacle: record the honest answer and let the Alignment Index surface it.

Language

The ENG / NL toggle sits at the top right and governs the landing page, the assessment and the paper alike. The choice travels in the address (?lang=nl) and is remembered by the browser. Switching language mid-assessment preserves every answer and re-renders the screen — including a report already on display, which is rebuilt in the new language.

NOTHING IS STORED — READ THIS BEFORE STARTING

Answers live in the browser tab and nowhere else. Nothing is transmitted, and nothing is saved to a server or to the device. Closing or reloading the tab discards every response; only the language preference persists. Complete a session in one sitting, and export the report before leaving the page.

The downloadable ZIP is the same instrument offline — open assessment.html in a browser. It behaves identically, including the absence of storage.

03 Taking the assessment

Seven screens, five questions each.

01	Respondent Set role, and optionally name and organisation, at the foot of the landing page.
02	Seven domain screens One domain at a time, five questions each. The progress strip shows how far the 35 have run.
03	Completion, then report Add another respondent, or open the report — and export it before the tab closes.

Three question types

Every question is scored 0–4. The type sets the scale and the weight.

TYPE	SCALE	WEIGHT
Agreement	Strongly disagree · Disagree · Neutral · Agree · Strongly agree	1.00
Evidence	No · Partly · Yes · Documented · Documented & regularly reviewed	1.25
Structural	Five options written for that question, ordered from absent to owned and rehearsed	1.25

Evidence and structural questions carry a quarter more weight than agreement questions. The instrument is deliberately biased towards what is written down and owned over what is believed.

Answering

- **Documented** means an artefact exists that another person could read without asking its author. If it lives in someone's head, the honest answer is lower.
- **Structural** options describe the arrangement as it stands, not the one being planned.
- If an answer could not be defended with evidence in a board meeting, choose the step below. A flattered score returns a flattering roadmap, which is worse than no roadmap.

Navigating and skipping

Back and Next move between domains, preserving answers in both directions; anything can be changed until Finish. A domain may be left partly answered — unanswered questions drop out of that domain's calculation rather than scoring zero. But a domain left entirely blank counts as 0, and because the overall score weights the weakest domain, one blank domain pulls the whole result down. Answer all 35.

04 Running it with a leadership team

One respondent produces a score. Several produce a governance signal.

When the editor-in-chief and the technology lead describe the same organisation differently, the gap is the finding: the same reality is being managed to different assumptions. That is a governance exposure in its own right, and one no single-respondent score can reveal. The Alignment Index measures it.

How to capture several respondents

On the completion screen, or from the report, choose **Add another respondent**. The landing page returns with the organisation pre-filled and the name cleared. Select the next role and run the 35 questions again. Repeat for each leader, then open the report once at the end.

THE ONE CONSTRAINT THAT MATTERS

All respondents must be captured in the same browser session, on the same device, without reloading the page. Sending the link to four people produces four separate single-respondent reports and no Alignment Index.

Two workable formats

- **Passing the laptop.** Each leader answers in turn on one machine, out of earshot of the others. Fastest, and keeps answers independent.
- **Facilitated round.** Project the questions and record each leader's answer in turn without discussion. Slower, but the divergence is visible live.

Do not discuss answers before they are recorded. Consensus reached in the room destroys the signal the Alignment Index exists to capture.

Three is the working minimum: the editor-in-chief, the publisher or CEO, and the technology lead. Add a supervisory board member where oversight itself is the question, and an editorial innovation lead where practice on the floor is likely to differ from policy above it.

05 Reading the report

The report is built to go into a board meeting unedited.

SECTION	WHAT TO READ IT FOR
Executive summary	Score, level, strongest and weakest domains, and whether leadership agrees. Two quotable paragraphs.
Overall maturity	The 0–100 figure against the five-level scale. Weighted towards the weakest domain — see Appendix A.
Domain scores	Radar and table. With several respondents the solid shape is the organisation mean; the dashed shapes are individuals.
Strengths, weaknesses & blind spots	Top two and bottom two domains. Blind spots flag a domain where leaders differ by 30 points or more, or a weakest domain below 40.
Alignment Index	Multi-respondent only. Read it before the score: where leaders disagree, reconcile before acting.
Immediate priorities	Quick wins for the two weakest domains. Actionable this quarter.
90-day and one-year roadmaps	Sequenced work. The 90-day plan addresses the two weakest domains; the one-year plan widens to four.
Suggested board agenda	Five items, written to be pasted into a meeting paper.
Benchmarking	Fixed illustrative peer medians, not live data. They frame the score; they do not validate it.
Recommendations by domain	All seven domains in three tiers. The weakest open expanded.

Exporting

Export / print report opens the browser print dialogue; choose *Save as PDF*. Enable background graphics so the charts and scales print. Interface controls are hidden from the printed version automatically. The report exists nowhere else — export it before the tab closes.

Live benchmarking activates once enough organisations have completed the assessment, with anonymous comparison by country, size, ownership and digital maturity. Until then the peer figures give the score a sense of scale and nothing more.

06 After the report

The test of the assessment is whether it changes a decision.

- **Name owners before the meeting ends.** The two weakest domains each need a named individual, not a committee that collectively owns nothing.
- **Put the 90-day priorities to a decision, not a discussion.** Each carries a resourcing implication; approve or decline it explicitly.
- **Set the reporting cadence.** The paper argues for information risk to reach leadership in the same rhythm, and with the same seriousness, as financial and cyber risk. A short standing item is enough.
- **Reconcile disagreement first.** Where the Alignment Index flags a domain, the gap in the leadership view is the first thing to close — acting on an averaged score that nobody holds is how governance becomes theatre.

Re-assessing

Run it again every six to twelve months, and after any material change: a change of primary provider, a merger, a newsroom reorganisation, a significant disruption. Keep the exported PDFs — the series, not the single score, is the evidence that governance is moving. Movement of five points or fewer between rounds is not significant; the instrument is self-reported, and small shifts reflect who answered and how they read a question as much as anything real. Read the level changes and the weakest domain instead.

It should change decisions, not merely document them.

Where the report surfaces the argument, the paper carries it: its implementation section covers inventory, assigned responsibility and policy, and its appendix gives six questions a board can put to itself without any tooling at all.

A Methodology & scoring

Every figure in the report can be reconstructed from the answers.

Items and weights

Thirty-five items, five per domain. Each answer scores 0–4 by the position chosen, leftmost being 0. Agreement questions weigh 1.00; evidence and structural questions weigh 1.25.

Domain score

$$\text{domain} = \Sigma(\text{weight} \times \text{answer}) \div \Sigma(\text{weight} \times 4) \times 100$$

Calculated over answered items only, so skipping a question widens the influence of the rest rather than penalising the domain. A domain with no answers at all is treated as 0.

Overall maturity

$$\text{overall} = 0.75 \times \text{mean}(\text{seven domains}) + 0.25 \times \text{lowest domain}$$

The overall figure is imbalance-aware by design. A quarter of it is set by the weakest domain alone, so strength across six domains cannot conceal absence in the seventh. Levels are assigned by dividing the score into five equal bands of 20 points, listed in Appendix B, and apply to both the overall figure and each individual domain.

Several respondents

Each domain becomes the mean of the respondents' scores for that domain, and the overall figure is calculated from those means. Individual respondents remain visible as the dashed shapes on the radar.

$$\begin{aligned} \text{range} &= \text{highest respondent} - \text{lowest respondent, per domain} \\ \text{alignment} &= 100 - \text{mean}(\text{range across the seven domains}) \end{aligned}$$

An alignment of 100 means identical answers. Any domain with a range of 30 points or more is flagged individually and repeated as a blind spot. Blind spots also list either of the two weakest domains scoring below 40. Where neither condition holds, the report says so.

Limits worth stating

- Self-reported and unaudited: the instrument records what respondents believe to be true of their organisation.
- It measures whether governance exists, is owned and is reviewed — not whether it is good.
- A single respondent measures one person's visibility of the organisation, not the organisation.
- Scores are comparable over time within one organisation; comparison against the illustrative benchmarks is indicative only.
- Scores are held to one decimal and displayed whole, so a domain table may appear not to average exactly to the overall figure. That is rounding, not error.

B Domains, scales and levels

Seven domains, taken from the paper unchanged.

DOMAIN		ASSESSMENT DIMENSIONS
01	Digital Sovereignty	Dependency awareness · vendor concentration · procurement · exit strategy · resilience
02	Information Intelligence	Investigative support · discovery · synthesis · editorial insight · analytical capability
03	Safeguarding Plurality	Source diversity · framing diversity · narrative convergence · originality · monitoring
04	Model Governance	Model evaluation · explainability · review cycles · transparency · auditability
05	Knowledge Stewardship	Prompt libraries · documentation · institutional knowledge · sharing · preservation
06	Strategic Information Autonomy	Continuity · fallback · resilience testing · operational independence · contingency
07	Executive Responsibility	Board oversight · reporting · ownership · governance cadence · strategic decisions

	SCORE	LEVEL	MEANING
1	0–19	Awareness	Dependence is felt but unmapped. Governance is absent or accidental.
2	20–39	Experimentation	Tools are adopted case by case. Practice exists in pockets, neither owned nor recorded.
3	40–59	Structured Governance	Policies, owners and review cycles are documented, though not yet integrated with wider risk.

	SCORE	LEVEL	MEANING
4	60–79	Integrated Governance	Information risk is reported and managed alongside financial and cyber risk.
5	80–100	Strategic Leadership	Governance shapes investment and editorial strategy. The newsroom moves deliberately on what it knows.

C Questions that come up

The answers disappeared.

Responses live in the browser tab. Reloading, navigating away or closing it clears them, by design — nothing is written to a server or to the device. Export the report before leaving the page.

Can respondents fill it in separately and combine the results?

Not in this version. The Alignment Index requires all respondents in one browser session on one device. Separate sessions produce separate single-respondent reports.

Can an answer be changed?

Yes, at any point before Finish — Back and Next preserve everything. Once a respondent is recorded, their answers are fixed; Start over clears the session and begins again.

The score looks low for a newsroom that does a lot well.

Check the weakest domain: a quarter of the overall figure comes from it alone. Then check for a domain left unanswered, which counts as 0.

Which language does the report use?

Whichever is active when it is generated. Switching the toggle rebuilds the report in the other language without losing any answer.