



MODULE 0

The Leadership AI Governance Toolkit

A governance position you can defend to your CEO, your board, or a journalist. Seven short artefacts built around four controls, each usable in under thirty minutes.

- ▶ **Disclosure** *when you tell people AI was involved*
- ▶ **Accountability** *who signs off, and who answers*
- ▶ **Risk and IP** *what never goes in, and who owns what comes out*
- ▶ **Capability** *whether your team can use it well*

Built by Michael MacLennan, twenty years in communications.
Free. Ten pages. Usable this week.

Leader Comms With AI · commswith.ai/leader

Version 1.0 · June 2026



START HERE

The gap, the four controls, the fortnight

Most communications teams adopted AI faster than they governed it. Drafting, research, summarising, monitoring: it is already in the workflow. What is usually missing is the part a leader is accountable for. When do we disclose it, who signs it off, what must never go into a tool, and who answers if it goes wrong.

Researchers tracking AI in PR since 2017 keep finding the same thing: teams believe AI policies have been asked for far more often than anyone reports putting them in place. Around half of in-house teams never ask their agencies about AI use at all, and agencies report the mirror image. The profession is using AI widely and governing it barely.

From the Comms With AI leader interview with Ben Verinder, May 2026.

Seven artefacts, four controls

Each artefact takes under thirty minutes. Worked through over a fortnight, they move you from “we’ll look into it” to “we have this in hand”.

	Artefact	Control	When
1	Where you stand: the readiness check	All four	Day 1
2	The disclosure decision guide	Disclosure	Days 2–3
3	The approval and accountability map	Accountability	Days 4–5
4	The claims and accuracy check	Accountability	Days 6–7
5	Confidentiality red lines and the IP note	Risk and IP	Days 8–9
6	The board-ready one-pager	All four	Days 10–12
7	The 30-day adoption path	Capability	Days 13–14

What this is not

This toolkit supports good governance and responsible adoption. It is not legal advice, and it does not claim to make you compliant with any specific regulation. What it gives you is the structure to get there with your own advisers. The bar is higher for regulated and public-sector teams, and this toolkit says so where it matters.



Where you stand

About 20 minutes

Twelve statements, three per control. Tick what is true today, not what is planned. The control with the most blanks is your priority for the fortnight.

Disclosure

- We have agreed when we tell audiences that AI was involved in our work.
- Our agencies and freelancers know what we expect them to disclose to us.
- We could answer a journalist asking “do you use AI?” without a scramble.

Accountability

- Every content type we publish has a named human approver.
- AI-assisted work goes through the same review as human-only work, or stricter.
- If an AI-assisted piece went wrong tomorrow, we know who answers for it.

Risk and IP

- Everyone on the team knows what must never be pasted into a public AI tool.
- We know which of our tools train on what we put in, and which do not.
- We understand that largely AI-generated work may carry no copyright at all.

Capability

- We know what the team is actually using AI for, not what we assume.
- Training has matched our real work rather than generic demonstrations.
- We run deliberate, bounded experiments rather than trying AI on everything.

Before you score it: run an amnesty

Get the team in a room that trusts each other and ask what people are actually using AI for. In three years of running these sessions, Ben Verinder reports there has not been one where the head of comms was not surprised. Usually the team is further ahead than anyone thought.



The disclosure decision guide

About 25 minutes

The default assumption has flipped. Increasingly, stakeholders assume you are using AI whether you say so or not. Being forthright about where you use it does a second job: it demarcates where you do not, and protects the premium on your human work. Disclosure alone is not enough, though. A University of Arizona study across thirteen experiments found that AI use stakeholders do not endorse carries a trust penalty even when disclosed. Telling people you use AI is the start; winning support for how you use it is the work.

Four questions for any piece of work

- ▶ Did AI materially shape the substance, beyond grammar, formatting or research?
- ▶ Would the audience feel misled if they discovered the role it played?
- ▶ Is this a high-trust category: health, finance, public policy, crisis, anything regulated?
- ▶ Has a named human verified the content and taken ownership of it?

A yes to questions 1 and 2 together, or to question 3 alone, means disclose. Question 4 must be a yes before anything ships, disclosed or not.

Three levels of disclosure, with lines you can lift

Level 1 · Assistive

Grammar, formatting, summarising your own notes. No per-item disclosure; covered by your published team position.

Level 2 · General

AI contributed to drafting or research; a human shaped and approved the result.

“Produced with AI assistance and reviewed by [name or team].”

Level 3 · Specific

AI generated substantial content: images, translations, data summaries, drafts published close to original form.

“This [report / image / translation] was generated using [tool] and verified by [role].”

Write your default position down once, in three places: your website or proposals, your client or stakeholder agreements, and your team handbook. Then it is a standard, not a scramble.



The approval and accountability map

About 30 minutes

AI changes nothing about accountability. The name in the sign-off column owns the output as if they had written every word. This map makes that explicit for your real content types. Five rows are worked examples; rewrite them for your own work, then add the rest.

Content type	AI's role	Reviewed by	Sign-off	Answers publicly
Press release	Drafting support	Comms manager	Head of comms	Head of comms
Social post	Drafting support	Peer review	Comms manager	Comms manager
Internal announcement	Drafting support	HR partner	Comms director	Comms director
Board paper	Research only	Comms director	Comms director	Comms director
Crisis statement	None on message	Crisis lead	Crisis lead	CEO / crisis lead

Three rules that keep the map honest

- ▶ AI's role is declared per content type, in plain words: none, research, drafting support, generation.
- ▶ Crisis and incident content: AI may support logistics, never the message itself. The stakes are too high and the source material too sensitive.
- ▶ The first time AI is used on a new content type, approval escalates one level up this map. After that it follows the standard row.



The claims and accuracy check

About 15 minutes

AI's failure mode is fluency: confident, plausible, wrong. Run this check on anything AI touched before it ships. It takes fifteen minutes and it is the cheapest insurance in this toolkit.

- Find every claim**
List each factual claim in the piece, including the implied ones: “the leading...”, “the first...”, “trusted by...” all count.
- Verify against a source you would cite to a journalist**
For each claim, name the source. If you cannot, the claim comes out or gets verified the long way.
- Check the details AI alters most**
Numbers, dates, names, titles, quotes and currency figures, each against the original document, not memory.
- Hunt for fabrication**
Links, citations, case studies and statistics AI tools invent with complete confidence. Click every link. Find every cited study.
- Test comparative claims against reality**
“Faster than”, “unlike competitors”: check against what competitors actually do, not their marketing.
- Confirm regulated claims with the owner**
Anything touching health, finance, environment or product performance goes to legal or compliance, with the evidence attached.
- Store the evidence**
Keep the sources with the published piece. If it is challenged in six months, the substantiation should take minutes to find, not days.

The standard, in one line

Would a regulator, a journalist, or your sharpest competitor find the substantiation adequate? If the honest answer is “probably”, the work is not finished.



Confidentiality red lines, and the IP note

About 20 minutes

One test covers most cases: if you would not email it to a journalist, do not paste it into a public AI tool. The list below should only ever enter an enterprise tool after your legal or compliance team has reviewed the terms and signed off.

Never into a public AI tool

- Embargoed announcements
- Pre-announcement financials
- Crisis communications content
- Medical or health information
- Personal data
- Legal documents and advice
- NDA-protected information
- Security incident details

Quick check before you paste

- Is this covered by an NDA or a confidentiality obligation?
- Does it contain personal data, even incidentally?
- Would this damage us, a client or a colleague if it appeared elsewhere?
- Does this tool train on what I put in? If I do not know, treat it as if it does.

The IP note, in plain English

- ▶ Work generated largely by AI, without sufficient human intervention, may attract no copyright at all. Nobody owns it, including you.
- ▶ An agency assigning you copyright in largely machine-made work may be assigning nothing. Asking your agencies how they use AI is not box-ticking; it goes to the heart of the commercial agreement. Around half of in-house teams have never asked.
- ▶ Your own materials can be used to train models under some tool terms and supplier contracts. Consider a clause preventing your work being used for AI training, in both directions.

Copyright in AI-assisted work is unsettled law in most jurisdictions. For anything commercially significant, take advice.

Public sector: Freedom of Information

Treat prompts and AI outputs as potentially disclosable under FOI. Assume anything you paste into a tool, and anything it returns, could appear in an FOI response, and check your records-management position on AI tool chat history.



The board-ready one-pager

About 30 minutes

AI adoption is a change programme, not an IT programme. IT can govern the tools; it cannot define the reputational judgement, stakeholder acceptability or tone of use. That is yours, which means briefing upward is yours too. This skeleton gets a board or CEO from zero to confident in three minutes. Fill it in, keep it to one page, date it.

1 · Our position

Two sentences. “We use AI where it strengthens our work, under four controls: disclosure, accountability, risk and IP, and capability. A named person approves and answers for every piece of work, with or without AI.”

2 · Where we use it today

Three bullets, specific: research and summarising; first drafts of [content types]; monitoring and analysis. Name the tools.

3 · Where we do not, and why

Two bullets. Crisis and incident messaging; anything carrying personal, embargoed or privileged material. One line on why.

4 · The controls in place

One line per control, pointing at the artefacts in this toolkit: our disclosure standard, our approval map, our red lines, our training plan.

5 · The risks we are watching

Accuracy in public content; confidentiality in tool use; the trust penalty of undisclosed or unendorsed use; copyright in AI-assisted work.

6 · What we need from you

A decision, a budget line, or nothing but awareness. Boards respond better to a specific ask than a status update.

Brief leadership before they ask. A board that first hears about your AI use from a supplier, a journalist or a regulator is a board you have already lost.



The 30-day adoption path

About 20 minutes

Governance on paper is not governance. This path turns the artefacts into team practice in four weeks, without a big-bang rollout. Two or three bounded experiments, not more: throw AI at everything and the impact becomes impossible to measure.

Week 1 · Find the truth

- ▶ Run the amnesty, then the readiness check (artefact 1).
- ▶ Agree the red lines as a team and put them where people work (artefact 5).

Week 2 · Set the standards

- ▶ Adopt your disclosure default (artefact 2) and the approval map (artefact 3).
- ▶ Choose two or three bounded experiments tied to real problems, not tools looking for a use.

Week 3 · Run the experiments

- ▶ Apply the claims and accuracy check to everything the experiments produce (artefact 4).
- ▶ Build prompt skill alongside: context, objective, audience, constraints, evidence, tone, review criteria. It is the briefing craft you already practise, made explicit.

Week 4 · Evaluate and brief upward

- ▶ A short, honest huddle: what worked, what did not, what gets adopted.
- ▶ Take the board one-pager upward (artefact 6) and set a quarterly review date.

Choose your champion deliberately

The person who volunteers as AI lead may not be the one you would choose. Pick for judgement and credibility with the team, not enthusiasm.

A dated policy beats a perfect one

A basic team policy now is worth more than a perfect organisational one later; any policy is right for about five seconds before the technology moves. Date it, review it quarterly.

Public sector: scope to approved tools

Before you run experiments, scope them to your approved-tool list and procurement sign-off. Use what is already sanctioned, and treat anything new as a procurement step rather than a quick trial.



MODULE 0 IS THE START

Honest limits, and what comes next

What this toolkit cannot do

- ▶ It cannot make AI tools accurate. The claims check catches errors; it does not prevent them.
- ▶ It cannot set your risk appetite. The artefacts give you the structure; the judgement stays with you.
- ▶ It is not legal advice, and completing it does not make you compliant with any regulation.
- ▶ It cannot govern for you. A toolkit nobody opens after the download governs nothing.

The evidence behind it

This module draws on the Comms With AI leader interview with Ben Verinder (May 2026), co-editor of AI for Public Relations (Kogan Page, 2026) and author of the CIPR's first papers on AI and data ethics, alongside the research discussed there, including King's College London's survey of attitudes to AI, the University of Arizona's thirteen-experiment study on disclosure and trust, and Global Alliance's 2025 work on transparency. Read the full interview at appliedcomms.ai/ben-verinder-ai-pr.

The rest of the series, all free

MODULE 1 · AVAILABLE

Capability and team readiness

MODULE 2 · AVAILABLE

Leadership-level strategy with AI

MODULE 3 · AVAILABLE

Board and stakeholder confidence

MODULE 4 · AVAILABLE

Measurement and assurance

All four follow-on modules are live now and free at commswith.ai/leader, alongside this one. The Govern phase of the free library has nine deeper templates behind artefacts 3 and 4, at commswith.ai/library/governance. And if you would rather work this through with someone who has done it before, Consult Comms With AI is a 45-minute session with a written advisory note, at commswith.ai/consult.



The Leadership AI Governance Toolkit, Module 0, version 1.0, June 2026. Built by Michael MacLennan for Leader Comms With AI. Produced with AI assistance, reviewed and approved by a human. © Comms With AI 2026.