CPRS | SCRP

AI Verification & Accountability Toolkit

Practical tools for the responsible use of AI in communications and public relations

October 2026

BEFORE WE START

Acknowledgements and Disclosure

Led by the Communications and Public Relations Society of Canada (CPRS) Thought Leadership Council, this toolkit was inspired by the publication, AI and Digital Ethics for Communicators, produced by the Chartered Institute of Public Relations' and member contributors, Frank Dias (Lead Creator), Courtney Ellul, Hina Sharma, Jenni Field, Andrew Hesselden, Fa Mafi, Amber Brittain, Nafisa Shafiq, and Grace Park.

Published in 2026, the CIPR document was released just as the Thought Leadership Council was reviewing the results of our delegate engagement during the 2026 National Conference in St. Andrews, New Brunswick. The feedback from communications and public relations professionals was clear: we needed guidance to manage the rapid adoption of AI; we needed to understand disclosure and verification; and we needed practical tools rather than more governance discussions.

In the spirit of full disclosure, this toolkit was developed using AI, specifically Claude CoWork and ChatGPT. The lead contributor used AI to analyze the delegate engagement results against the CIPR publication, analyze recent studies and thought leadership (listed in Appendix 10), scan the Canadian legal environment, and organize the structure before drafting. AI was used to support multiple drafts of the document and to synthesize feedback from the reviewers. A final copy edit was conducted, and AI supported the translation of the document into French.

Claude Design was also used to guide the final design of the document in Microsoft word, providing instructions for the lead contributor (because hey, we're on a budget here).

Here's the human process:

CLASSIFY	This was classified as a Tier 3 highest oversight (see section 04). All engagement data was scrubbed of identifiers, emails and names and details before being analyzed. It was understood from the beginning this would require significant review and oversight to ensure the accuracy and quality of the toolkit.
PROMPT	Specific projects were developed within the AI platforms with very strict instructions and requirements. The document was built section-by-section with significant human feedback and corrections before moving to the next section.
VERIFY	All sources in this document were reviewed by the lead contributor. All statistics and information were checked by both AI platforms and then manually located and verified by the lead contributor. The document was reviewed by those noted in the table below. French translation was reviewed by Louis Payette of ChangeMakers.
DISCLOSE	All working drafts indicated that AI was used in the development, so reviewers understood the tools used and that they were reviewing AI-assisted and/or AI generated outputs. No content was provided to the reviewers without human review.

Special thanks to our contributors and reviewers

THOUGHT LEADERSHIP COUNCIL

Colleen Killingsworth, MCM, APR, FCPRS (TLC Chair)
Kim Blanchette, APR, Chart.PR, FCPRS (Lead Contributor)
Lisa Covens, MCPRS, MA, CAIP
Sara Rafuse, MBA, CPF
Alex Sevigny, Ph.D., APR, FCPRS
Christine Szustaczek, MCM, APR, SCMP, FCPRS
Amy Thurlow, Ph.D., APR, FCPRS
Martin Waxman, MCM, APR, MCPRS

EXTERNAL REVIEWERS

Andrew Blanchette, MCPRS, Senior Vice President, ChangeMakers
Andrea Sweet, Senior Vice President, Castlemain
Louis Payette, Senior Director, Castlemain

CONTENTS

Table of Contents

01 Purpose and relationship to CPRS Guidance

02 The Canadian legal landscape

03 Our role in AI ethics

04 Ethical checks

05 Verification and accountability tools

06 Disclosure decision tree

07 Dos and Don’ts by task type

08 Scenarios

09 Getting started checklist

10 Appendix – Resources

01

Purpose and relationship to CPRS AI Guidance

CPRS has chosen to approach this as a verification and accountability toolkit, not a governance toolkit.

That’s a deliberate choice.

01 | PURPOSE AND RELATIONSHIP TO CPRS AI GUIDANCE

This toolkit is meant to serve as a companion to CPRS's AI Guidance, approved by the CPRS National Board in August 2026. The Guidance sets seven core principles for responsible AI use: transparency, accuracy and integrity, privacy/security/consent, equity and inclusion, human oversight, accountability and advancement.

This toolkit shows how to apply those principles to everyday work and builds out three things the Guidance itself commits CPRS to developing as companion resources:

- sample disclosure language
- prompt guidance and
- a self-assessment tool for organizational AI governance.

CPRS has chosen to approach this as a verification and accountability toolkit, not a governance toolkit.

That's a deliberate choice.

In the battle over who "owns" AI governance between IT, legal, executives and communications and public relations, accountability can disappear.

Communications and public relations professionals are stewards of organizational and corporate reputation. Specific decisions in the AI workflow, such as what gets verified before it goes out, what gets disclosed and to whom, are ours to own. We have always applied rigorous review, verification and approval standards to our work, regardless of who (or what) drafted it.

The toolkit closes a gap we have seen emerge in the last 18 months. In 2025, *the CPRS State of Public Relations in Canada* research found that 90% of Canadian professionals believe AI must be used ethically, but only 27% report a well-established, implemented AI policy at their organization. Boston Consulting Group's (BCG) 2026 survey of Corporate Affairs and Communications leaders found the same pattern: 88% say they are not fully prepared to lead an AI transformation in their function.

Research from the Institute for Public Relations found that IT is typically the primary owner of generative AI adoption, but communications and public relations professionals act as the translator between technical teams and the rest of the organization. That translator role carries verification and disclosure responsibilities of its own.

> As one CPRS member said at the National Conference in Banff, “we’re running with AI before we’ve learned to walk with policy.”

This toolkit has been adapted, with permission, from the Chartered Institute of Public Relations’ (CIPR) Inside Group guide, AI and Digital Ethics for Communicators, and draws from other best practices and thought leadership from the Public Relations Society of America (PRSA), the Institute for Public Relations (IPR), among others. A full list of sources is available on page 42.

This toolkit gets us walking by providing the infrastructure and guardrails you need, whether you’re just getting started or are seeking to improve your existing process.

Things are moving quickly, likely faster than this toolkit can be updated. Communications and public relations professionals are encouraged to stay up to date on the latest trends and regulatory changes.

CPRS is a member of the Global Alliance for Public Relations and Communication Management and supports the Global Alliance’s Venice Pledge for the responsible use of AI in public relations. This toolkit is built consistent with that commitment, putting the Pledge’s principles into day-to-day practice for CPRS members.

02

The Canadian legal landscape

The case for verification, accountability and clear rules around disclosure is growing, but not at the rate of adoption.

02 | THE CANADIAN LEGAL LANDSCAPE

It's important to understand what governs AI use for Canadian communications and public relations professionals in 2026, and how federal law intersects with provincial requirements, specifically in Quebec.

At the federal level, the Personal Information Protection and Electronic Documents Act (PIPEDA) itself does not reference AI. However, the Office of the Privacy Commissioner (OPC) has stepped in with joint principles for responsible, trustworthy and privacy-protective generative AI. The OPC has also issued findings directly concerning AI systems: a 2026 finding on OpenAI (#2026-002) and a 2026 finding against X Corp and xAI over Grok's deepfake image generation (#2026-004).

The OpenAI finding came out of a joint investigation with Quebec's, BC's and Alberta's privacy regulators. It found that OpenAI scraped personal information at a scale that went beyond what was appropriate, without valid consent, and had not tested whether the personal information its models generate about real people is even accurate.

OpenAI has since committed to a new filtering tool, clearer notices and better access and correction processes. The federal regulator considered its findings conditionally resolved; BC and Alberta concluded consent requirements still were not met. Their laws are more specific on consent, and both found the practice still non-compliant. For communications and public relations professionals, this has real impacts. A tool this widely used inside organizations, often without any procurement review at all, was found non-compliant on the two things practitioners rely on most, consent and accuracy.

The Grok finding is even more concerning. X and xAI launched an image-generation tool without a completed privacy impact assessment. In fact, the assessment wasn't finished until months after launch.

News reporting and outside researchers, cited in the Commissioner's finding, estimated Grok was used to generate and share about 1.8 million non-consensual sexualized deepfakes, including, per one estimate from the Center for Countering Digital Hate, more than 23,000 depicting children. The OPC has not independently verified these figures.

The Commissioner recommended suspending the tool until safeguards were proven. The companies declined, and say the monitoring they added instead cut incidents by half.

As regulatory scrutiny increases, there are considerable and growing resources that help communications and public relations professionals build their own approaches for responsible use of AI.

Specifically the Treasury Board of Canada's FASTER principles (Fair, Accountable, Secure, Transparent, Educated, Relevant) and its Directive on Automated Decision-Making provide effective tools for AI use in the public sector that can inform private-sector practice.

Quebec goes even further. Law 25's amendments to the Act Respecting the Protection of Personal Information in the Private Sector put a specific duty on automated decisions, requiring disclosure and explanation whenever an automated system decides about a person without meaningful human involvement. Quebec also requires a designated privacy officer and imposes a two-tier penalty structure, administrative monetary penalties up to $10 million (or 2% of worldwide turnover, if greater) and fines up to $25 million (or 4% of worldwide turnover, if greater).

The impact of AI without guardrails is about more than hallucinations or the reputational embarrassment that follows. IBM's Cost of a Data Breach Report 2025 found that, in Canada, breaches involving shadow AI, tools used without organizational approval or knowledge, carried an average of CA$308,000 in added cost, and 20% of organizations studied had a security incident involving shadow AI. A separate study commissioned by IBM found that while 79% of full-time office workers use AI at work, only 25% rely on enterprise-grade AI tools. The rest use a mix of personal and employer tools (33%) or rely entirely on personal apps (21%). Shadow AI was a top concern in roundtables at CPRS's 2026 National Conference in St. Andrews, New Brunswick, with unmanaged data exposure the core concern raised.

TELUS's 2026 AI Trust Atlas, produced with Leger, found 90% of Canadians want AI regulated, but only 27% are confident current laws are adequate. The case for verification, accountability and clear rules around disclosure is growing, but not at the rate of adoption.

03

Our role in AI ethics

When everyone owns AI governance, no one does.

03 | OUR ROLE IN AI ETHICS

When everyone owns AI governance, no one does.

That poses risks for organizations and their audiences. TELUS's 2026 AI Trust Atlas reports only 34% of Canadians trust companies that use AI. The same research shows what would change that number: 76% would trust a company more if it actively reviewed its AI systems for harm, 73% if it explained its AI use in plain language, and 69% if it gave the public a say in how AI gets used.

CPRS calls for the "early and ongoing involvement of communications professionals in AI design, deployment and governance to assess reputational risk, ethical considerations and public impact."

The five working roles below demonstrate how that involvement can happen day-to-day. In some cases, our role is to inquire and escalate rather than execute. What's important is that communications and public relations professionals establish working relationships with the functions that own or support each responsibility.

Translator/Educator

Owns

Works with: IT (source of technical explanation)

Trust Steward

Owns

Works with: Legal

Risk Spotter

Supports and escalates

Works with: IT, privacy officer, data scientists / ethicists for research-based AI work; data analysts, for statistical or methodological review of AI outputs

Accountability Enabler

Advocates and applies

Works with: Government, Legal and IT.

Change & Capability Enabler

Owns internally | advocates externally

Works with: HR, Learning and Development

Translator / Educator

Before an AI tool gets used or expanded inside your organization, communications and public relations should be able to explain it in plain language – what it does, what information it touches and what could go wrong – to executives, employees or the media.

If the organization can't provide that explanation, it's likely the tool isn't ready to use. This translator function is mirrored in the IPR research: *IT owns adoption; communicators explain it to everyone else.*

This role also needs to watch for AI tools already in use that nobody signed off on. Shadow AI creates real reputational risk on top of the security exposure.

That explanation rarely ends the job. Explaining a tool well enough for people to trust it tends to create additional needs: ongoing education on the tool itself, and on what it changes about how people work.

Trust Steward

Using the tools in this toolkit, communications and public relations professionals provide advice and counsel to their organizations about whether and how something gets disclosed as AI-assisted. In many cases that is a judgment call, based on risk and impact. But in Quebec, when an automated system decides about a person with no meaningful human involvement, the law requires disclosure.

This role balances two concerns. On one hand, under-disclosing AI use risks misleading an audience, but so does over-claiming it: calling basic automation or simple formatting tools "AI-powered" for the sake of sounding current erodes trust the same way hiding real AI use does.

Risk Spotter

Risk here goes beyond privacy, which is important. It includes bias, accuracy, security, reputational exposure, compliance and legal issues. Communications and public relations professionals can surface these risks early, asking the awkward internal question before someone outside the organization asks it.

Most privacy impact assessments get written by IT, security or a privacy officer and are not always reviewed by communications and public relations professionals. Here the role is to support and ask key questions:

- has the information been verified?
- does an impact assessment exist for it?
- would this create issues if it were disclosed after the fact?

The OPC's Grok finding illustrates the danger of completing a privacy impact assessment after launch, once harm has already occurred. Spotting risk and working with your organization to manage the mitigation takes this role beyond compliance to key factors impacting reputation and trust.

Where AI touches audience or participant research specifically, the Canadian Research Insights Council's Guiding Principles for AI use in Market Research recommend multidisciplinary oversight, data scientists, ethicists and legal expertise alongside communications.

Accountability Enabler

The Treasury Board's FASTER principles and the Directive on Automated Decision-Making are government frameworks, built by government, legal and IT functions working together. Communications and public relations professionals need to ensure that their work follows a clear assessment, verification, approval and disclosure process.

In practice, this means building the templates, checklists and review steps that make ethical AI use repeatable, rather than something that depends on the expertise of whomever receives the work.

If your organization has an existing policy that guides those actions, great. If it doesn't, this is where communications and public relations professionals are accountable to advocate for one and show leadership by establishing a AI policy for the communications function, regardless of what exists organization wide.

Change & Capability Enabler

AI adoption is an ongoing change journey. How it gets introduced, sequenced and implemented shapes whether people trust it or resist it. This is where communications and public relations professionals play a key role.

Two ideas from Canadian thought leadership reinforce the importance of the role.

Dr. Alex Sévigny's four-element model for communicator capability, data, then AI literacy, then business process, then change management, deliberately in that order, argues against starting with tools. His efficacy-effectiveness-efficiency model tests whether AI works in a controlled setting, then in the real world, before you optimize for cost – a useful check for both this role and the Risk Spotter role.

Communications and public relations professionals need to build the skills and acumen necessary to understand the tools, how they're being adopted within their organizations and mitigate risks. And we need to do that from a place of understanding the efficiency pressures and realities of the organizational environment. We shouldn't be throwing up roadblocks, but we can support transformation with sound counsel and practical risk management.

A note on cognitive offloading

People have always offloaded cognitive work. We don't do math using an abacus, and we wouldn't ask for a house plan drafted by hand. AI tools help offload more of the thinking through drafting, summarizing and, in some cases, making decisions. Michael Gerlich's 2025 study, AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking, found heavier reliance on AI predicts measurably lower critical-thinking scores, with the effect strongest among younger participants. For teams, that points to extra care with newer and less experienced staff.

04

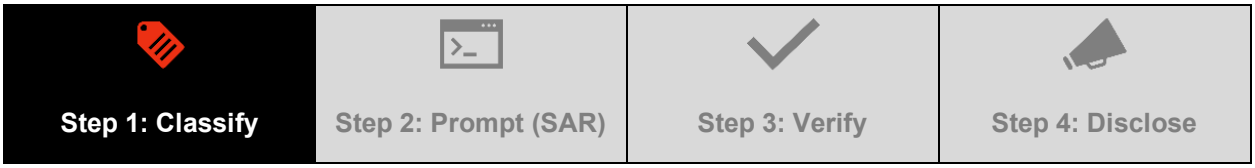
Ethical checks

Four integrated steps to working ethically with AI.

04 | ETHICAL CHECKS

There are four integrated steps to working ethically with AI: classify what you're doing, prompt it with discipline, verify according to what you classified, then disclose it.

Step 1: Classify



Classify before you ask AI for anything. Understand what you're about to create and the level of risk it poses for your organization.

> Everything in this section assumes the prompt itself doesn't contain personal identifiers or confidential information. If a task would put personal information, health information, or confidential organizational data directly into an AI tool, you need to pause unless you have a process in place to protect the information that has been approved by IT and Legal.

For example, analyzing feedback or consultation responses that include identifying details, you should de-identify or anonymize the data before it is put into the tool. That process needs to be documented and auditable.

Quebec's law is very clear about this. De-identified information no longer allows someone to be identified directly; anonymized information no longer allows them to be identified directly or indirectly, irreversibly. Confirm the specific requirement with legal counsel before relying on any anonymization approach for anything beyond routine internal use.

For everything else, setting a tiered approach helps you determine your approach. Six questions need to be considered to assess any piece of work:

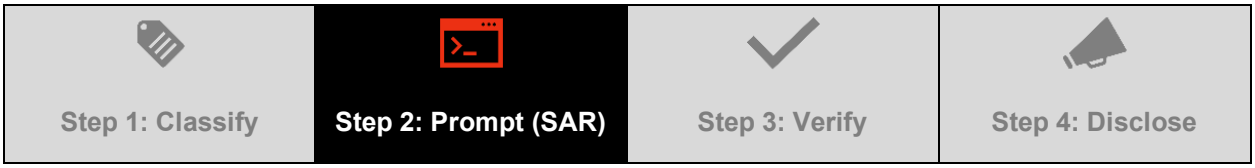
? Is it going public?	? Could it cause harm if it's wrong?
? Is it sensitive?	? Would it be hard to walk back?
? Does it involve facts or numbers someone will rely on?	? Decision-or-data trigger: does the work also rank, score or recommend something, inform or drive a decision, or contain numbers, statistics or data interpretation?

Start at Tier 1. Each “yes” moves the work up one tier, to a maximum of Tier 3. Two “maybe” answers count as one yes. A single yes, including “is it going public?” moves routine content to Tier 2.

TIER	WHEN IT APPLIES	WHAT IT MEANS
Tier 1: Light	Low-impact internal work with no “yes” answers	Self-check by the person who used the tool
Tier 2: Standard	Most day-to-day communications and public relations work with a single “yes” answer	An independent second reviewer checks the output
Tier 3: Enhanced	Work with two or more “yes” answers. For example, public and sensitive, or hard to walk back and feeding a decision	Intake, AI draft, review, specialist review, oversight, final QA and a validation log

CRIC’s Guiding Principles for AI use in Market Research use the same logic applied to research: ethical review should be proportional to scope, risk and potential impact.

Step 2: Prompt with SAR



Once you’ve established your tier, CIPR provides SAR, or Setup, Ask, Refinement as the discipline for the prompting itself, based on how much scrutiny that tier requires.

S **Setup:** privacy and data minimization. State the actual goal rather than a vague one: a specific brief produces a specific answer. Add context but stop at the privacy line. Beyond what’s been discussed about personal information, be sure to remove client names and commercially sensitive detail, based on your organization’s guidance.

A **Ask:** accuracy and scope control. Specify the format you want rather than leaving it open. Ground the request in your own sources: “use only the information below; if something’s missing, list it as a question” reduces invented detail more reliably than asking the tool to flag its own hallucinations, which it can’t reliably do. Set constraints

against jargon, invented statistics and marketing language, and break complex requests into steps rather than asking for everything at once.

R **Refinement:** bias, assumptions and impact. Treat the first response as a draft, not a finished answer. Ask what assumptions it made and what's missing. Ask it to challenge itself: where might this be biased, whose perspective is missing, what's the strongest counterargument? You may want to compare outputs across two tools or two attempts to understand where further verification is needed.

Step 3: Verify



Verification is guided by the tier assigned in Step 1.

- Tier 1 is a self-check by the person who used the tool.
- Tier 2 needs an independent second reviewer, someone other than the drafter.
- Tier 3 requires the full review sequence: intake, AI draft, review, specialist review, oversight, final QA and a validation log (see example below).

Where Step 1 flagged a decision-and-data trigger, go a step further to verify the underlying data before you use it, recompute the numbers and make sure a human interprets what the output means before anyone acts on it.

CPRS's AI Guidance gives every member a short checklist to run before using AI at all. The eight checks below take that same idea and apply it to a specific piece of work, in more depth. Each check is a question to ask yourself, both before and after prompting, and a warning sign for when you need to look closer.



Purpose & benefit

What benefit are we trying to deliver, and what harm could occur if this goes wrong?

Warning sign: you can explain the efficiency gain but not who benefits beyond your own team.



Data & privacy

What data is being used, does it include personal or sensitive information, and is there a documented reason for using it this way?

Warning sign: you can't quickly say what data went into the prompt or where it sits.



Accuracy & evidence

What's verified versus generated, and what needs fact-checking before publication?

Warning sign: a confident-sounding output is being trusted because it sounds right.



Bias & inclusion.

Who might be excluded or disadvantaged? Is there stereotyping risk in the data or output?

Warning sign: no one has questioned whose experience the output reflects. Where the work touches Indigenous data or communities, apply a distinctions-based approach and check against data-sovereignty principles such as OCAP® and CARE.



Transparency & disclosure.

When and how should AI involvement be disclosed, and what would the audience reasonably expect?

Warning sign: you'd be uncomfortable explaining your AI use to the audience after the fact.



Human oversight.

Who is accountable for the final output, and has a person reviewed it before publication?

Warning sign: you can't name who'd answer for it if it went wrong tomorrow. TELUS's AI Trust Atlas found 60% of Canadians say human oversight is essential, regardless of the task, and comfort with AI healthcare decisions rises from 39% to 63% once a human is confirmed to be supervising.



IP & rights.

Are copyright, licences and ownership respected, and could the output infringe someone else's IP?

Warning sign: no one has checked the tool's terms of use on who owns the output.



Security & misuse.

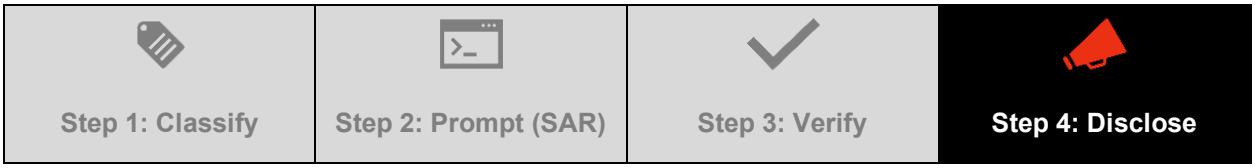
Could this be exploited for impersonation, phishing, data leakage or reputational harm?

Warning sign: no one has asked what happens if this ends up in the wrong hands.

Different tiers will need different checks. The table below provides some guidance based on your classification in Step 1.

TIER	WHEN IT APPLIES
Tier 1: Light	Accuracy, Disclosure, Human oversight, Data & privacy, as a base standard.
Tier 2: Standard	Accuracy, Disclosure, Human oversight, Data & privacy, plus whichever of the remaining four are relevant to the task.
Tier 3: Enhanced	All eight.

Step 4: Disclose



Verification confirms the work is correct. Disclosure confirms your audience knows how it was made and gives you a record to point to if anyone asks.

Every record should capture four things:

Business purpose. What the professional was trying to accomplish.	Inputs. What information and instructions the AI received: prompts, sources, other inputs.	Testing. How the output was tested and transformed before use.	Accountability. Who accepted responsibility for the final communication and when.

This is a basic and audit-worthy approach, which will keep evolving as disclosure norms mature.

Tracking: a simple log tracks every AI-assisted deliverable through its review lifecycle. A more detailed version can also track a hyperlink to each draft stage (AI draft, first human review, final review), an output-quality score, and whether the item has moved into a broader publishing or QA workflow. Add columns as your own tracking needs grow.

RECORD ID	OUTPUT	TOOL	DATE CREATED	CREATOR	REVIEWER	DATE REVIEWED	APPROVED	NOTES
1	[Deliverable name]	[AI tool]	[Date]	[Name/role]	[Name/role]	[Date]	[Yes/No/ Pending]	[Notes]

Already saying this takes too much time? When you consider the efficiency gains with these tools, a few moments tracking your outputs and review process for higher-risk items is a small investment in your reputation.

For the case-by-case judgment call on whether and how to disclose, see the Disclosure Decision Tree (Section 6).

What if...?

These steps can feel abstract until they're tied to a specific piece of work. A few scenarios communications and public relations professionals might experience:

You're summarizing consultation, survey or stakeholder feedback into a briefing note or executive summary. Trigger: Involves facts someone will rely on; may touch a sensitive topic **Do:** Apply the decision-and-data overlay: verify the source data before using it and confirm the privacy position if responses include personal information.	**You're pulling together campaign or media performance numbers, reach, engagement, coverage, click-through, for a client or leadership report.** Trigger: Involves facts or numbers someone will rely on **Do:** Recompute the numbers rather than take the AI's figures on faith. This is likely the most common Tier 2 scenario communicators will face day-to-day.
You ask AI to rank or recommend which content, message or channel performed best, to guide where budget or effort goes next. Trigger: Could cause harm if it's wrong; hard to walk back **Do:** Treat the ranking as decision support, not fact. A second reviewer checks the logic before it informs a decision.	**A brief or report you're producing could inform an eligibility, funding, policy or program decision.** Trigger: Could cause harm if it's wrong; hard to walk back **Do:** This moves to Tier 3: Enhanced. If the work is government or client-facing, check whether the Directive on Automated Decision-Making applies.

An AI-generated statistic or summary looks right, and there's no time to check the source before it goes into a briefing or public document.

Trigger: Involves facts or numbers someone will rely on

Do: This is exactly the scenario the decision-and-data overlay exists for. Verify before you rely on it: plausible is not the same as verified.

05

Verification and accountability tools

Things you can start using today.

05 | VERIFICATION AND ACCOUNTABILITY TOOLS

CPRS's AI Guidance already asks every organization to have four things in place:

An approved tool list	A review and sign-off process	An escalation path	A basic AI use policy

This section turns those four things into seven working parts you can use today. Parts one through five are set out below; parts six and seven continue in Section 6 (the Disclosure decision tree) and Section 7 (Dos and don'ts by task type). If you don't have the four basics yet, begin with those four basics. If you do, start with the first three parts below.

1 Disclosure guidance

Decide on your disclosure approach and your statement now. Build a small library of pre-approved disclosure lines for your team's repeat scenarios, internal updates, external content, client deliverables, so no one is drafting disclosure language from scratch, and everyone knows what disclosure means in practice.

> Worth watching: the European Union's Code of Practice on marking and labelling AI-generated content sets a standardized set of disclosure icons for generative AI, tied to Article 50(4) of the EU AI Act's transparency requirements for deepfakes and AI-generated content on matters of public interest. While not binding in Canada, a number of international organizations have signed on. The icon-based approach is an interesting alternative to text-only disclosure lines, especially when content is generated and not reviewed by a human.

2 Data handling rules

Use the Classify step and the guidelines on personal information to establish your ground rules. Personal data, client data and confidential or commercially sensitive information do not go into an AI prompt unless anonymized, or the tool is specifically approved for that data class.

3 Review and sign-off steps

Use the Tier approach to build out your sign-off steps. The simplest rule? No one approves their own AI output for Tier 2 and above.

4 Escalation rules

Named contacts in Privacy, Legal, HR, Security and Data/Analytics, with clear triggers for when and how to escalate, such as personal or sensitive data involved, reputational risk, external-facing content, statistical or methodological questions and anything in Tier 3.

5 Approved tool list and intended use

A simple table naming each approved tool, what it's approved for, and what data boundaries apply. Name an owner for this list and for the escalation list in item 4, so both stay current. Review and update at least quarterly, and whenever a tool, model, contract or data practice changes. For any third-party AI vendor, apply CRIC's due-diligence standard: review the vendor's security policies and compliance status before approval, and document the data-transmission protocol between your systems and theirs.

TOOL	APPROVED FOR	DATA BOUNDARIES
[Tool name]	[e.g. first drafts, internal summaries]	[e.g. no client or personal data]
[Tool name]	[e.g. approved for client-facing work]	[e.g. enterprise agreement in place]

06

Disclosure decision tree

Hint: when in doubt, disclose.

06 | DISCLOSURE DECISION TREE

Disclosure was the single most consistently named concern at the roundtable sessions during CPRS's 2026 National Conference.

CPRS's AI Guidance is clear: "When in doubt, disclose. Transparency is the starting position for AI use, not the fallback." It then splits disclosure into three contexts:

- **Internal use:** tell colleagues and managers when AI shaped a draft, analysis or decision they're relying on
- **Client work:** disclose when AI involvement is material to the client's decision-making or deliverable acceptance, and ask when in doubt and
- **Public-facing content:** disclose whenever AI involvement could reasonably affect how the message is received or trusted.

Guessing wrong in either direction has impacts. Over-disclosing brings chaos to routine work while under-disclosing risks misleading people.

The four-question Trust Test below helps you make those decisions in the moment.

1 **Is this being shared with anyone beyond you, even informally either internally, externally, with a customer, or publicly?**

If no, internal log only. No disclosure required.	**If yes, continue to step 2.**

2 **Did AI shape the content, analysis, or recommendation? Polish only (spell-check, formatting, or a minor suggestion you review and could easily reject) doesn't count.**

If no, no disclosure required. Document the AI's use internally.	**If yes, continue to step 3.**

3 **Is there a named human owner who has substantively reviewed the output and accepted accountability for it?**

Disclose	**If yes, continue to step 4.**

4 **Would a reasonable person in the audience feel misled if they later learned AI shaped this content?**

If no, disclosure not required. Document the reasoning behind the decision.	**Disclose**

In Quebec, when an automated system decides about a person based exclusively on automated processing, section 12.1 of the Act Respecting the Protection of Personal Information in the Private Sector requires disclosure and an explanation, regardless of what the four questions above conclude.

A usable disclosure statement covers four things: the tool used, what it was used for, how the output was verified and who reviewed the result.

[Tool Name] was used to draft [what], verified by [method], and reviewed by [name/role] before publication.

Keep it specific. A vague acknowledgment that "AI assisted with this" leaves too much in question and removes the human review and verification.

A complication worth discussing

One piece of research complicates the "when in doubt, disclose" approach.

Schilke and Reimann's 2025 study ran 13 experiments, testing how people judge supervisors, subordinates, professors, analysts, creatives and investment funds, and found disclosing AI use lowers trust compared to not disclosing or to saying a human did the work, regardless of how the disclosure is worded, whether it's voluntary or required, or whether the evaluator already knew AI was involved.

Schilke and Reimann's research suggests being exposed for hiding AI use hurts trust even more than disclosing it upfront. The practical choice is less about a "smaller hit now" and more about matching disclosure to audience expectations, materiality and any policy or legal obligations that apply. The penalty is somewhat smaller among people who already have positive attitudes toward technology, but it isn't softened by how familiar someone is with AI or by their own AI use.

That risk may fade over time. Compared to the cost of being exposed later, the case for disclosing now still holds, and disclosure paired with genuinely responsible use should earn back trust as AI use becomes normal.

How you disclose matters as much as whether you disclose. TELUS's 2026 AI Trust Atlas found 73% of Canadians are more willing to trust a company that explains its AI use in terms they understand.

07

Dos and Don’ts by task type

What this looks like in practice.

07 | DOS AND DON'TS BY TASK TYPE

CPRS's AI Guidance sorts AI use into three areas:

- appropriate use
- use with caution and
- never use.

The six task types below show what that looks like in the specific formats communications and public relations professionals produce every day.

Drafting communications

Emails, social posts, intranet articles, scripts and other written content.

Quick tip: add a standard line to your prompt: "Verify facts, names and numbers, and cite sources where possible."

DO

- Use approved tools and approved data sources only.
- Remove sensitive details before prompting: personal data, client names, internal financials, commercially sensitive information.
- Be specific about what data the AI should and shouldn't use.
- Have a second person review and fact-check every Tier 2 and Tier 3 deliverable.
- Run a plain-language and inclusion check on anything high-visibility.
- Where content is bilingual, run a French-language equivalence check, not just a translation pass. A technically accurate translation can still miss tone, idiom or meaning.

DON'T

- Paste confidential documents into free or non-approved AI tools without confirming enterprise-grade protections tied to your organization's contract and data-handling review; if only a consumer version is available, check its data and privacy settings first.
- Publish AI-generated content without verifying facts, statistics, policy references or quotes.
- Use personal AI tools for work content. Follow your organization's acceptable-use policy.

Summarizing meetings and transcripts

Quick tip: if you share AI-generated notes, add a line: “AI-assisted summary. The meeting lead should confirm decisions against the original recording or notes.”

DO

- Confirm your workplace or client relationship permits recording and transcription before you start.
- Treat AI-generated summaries as drafts, reviewed by default, not finished records.
- Confirm key decisions and actions with the meeting owner and keep your own notes alongside the AI summary.
- Store summaries with the right access permissions, especially for client-facing meetings.

DON’T

- Share restricted summaries broadly “for convenience.”
- Treat an AI summary as the definitive record without human review.

Audience listening, sentiment, and pulse analysis

Quick tip: add a note wherever this data is shared: “Indicative themes, not a definitive measure of individual intent.”

DO

- Explain what’s being analyzed and why.
- Aggregate data where possible and minimize personal information in outputs.
- Check for bias: who’s over- or under-represented in what you’re analyzing?
- If AI directly interacts with people, chatbots, adaptive surveys, tell them they’re talking to AI, not a person, and disclose any synthetic data or synthetic personas used and to what extent.

DON’T

- Use AI to infer sensitive personal attributes, such as health, political views, or ethnicity, without consent.
- Present AI output as “truth” rather than “signal.” It still needs human judgment and verification.
- Let someone reasonably believe an AI system is a human interviewer or moderator.
- Present findings drawn from synthetic personas or synthetic data as though they came from real respondents.

Manager toolkits, scripts and FAQs

Quick tip: build a one-page manager FAQ with four sections: key messages, likely questions, escalation paths and the manager's role.

DO

- Give managers what they need: what's known, what isn't, what to say, where to escalate and their specific role.
- Time releases to reduce noise and keep one source of truth.
- Test messaging with a small group before a broad cascade, particularly across regions or language groups.

DON'T

- Flood channels with frequent updates that add anxiety rather than clarity.
- Assume a message will cascade accurately without clear guidance and an escalation route.
- Sit on insights about how messages are landing. Feed manager and employee feedback back into the next round.

AI team charter and "how we use AI here" narratives

Quick tip: as a team, write a short "what we will and won't do with AI" list. It's a faster confidence-builder than a formal policy.

DO

- Write for trust: clear accountability, clear boundaries, clear support routes.
- Publish a simple "AI use in comms" statement covering what you do and don't do with AI.
- Include a redress route: what someone can do if they disagree with an AI-influenced outcome.
- For agency teams, confirm whose AI policy governs client work, yours or the client's, before the charter goes out. Client contracts increasingly address this directly, so check the current agreement rather than assume.

DON'T

- Overpromise certainty ("AI will solve everything").
- Hide AI involvement where the audience would reasonably expect disclosure.
- Dismiss people's concerns about job security, environmental impact or other AI worries. Bring in someone who can answer them.

Automating routine comms workflows

Recurring tasks where AI can save real time: drafting status updates, scheduling and tailoring social posts, assembling newsletters, triaging inbox queries, first-pass meeting recaps, tagging and routing media mentions.

Quick tip: for every automation, answer the following: who owns it, what triggers human review and how would you know if it broke?

DO

- Start with high-volume, low-risk, repeatable tasks where a human still reviews the output before it goes anywhere. This adapts Treasury Board's low-risk framing to your own workflows.
- Map the workflow end-to-end before automating any step.
- Keep a named human owner for every automated workflow, even when AI is doing most of the work.
- Build a stop trigger into each workflow: a signal (sensitive topic, complaint, anomaly) that pulls a human back in.
- Re-review automations quarterly. Tools, prompts and source material all drift.

DON'T

- Automate anything touching a regulated decision, a dismissal, a complaint or an emotionally sensitive topic without an explicit human checkpoint.
- Chain tools together (AI drafts, AI posts, AI replies) without a human checkpoint somewhere in the chain.
- Let an automation run unattended after the person who built it leaves. Orphaned automations are a common source of quiet failures.

08

Scenarios

Use these for team discussions, training sessions or a personal check.

08 | SCENARIOS

Use these for team discussions, training sessions or a personal check when something similar comes up.

AI wrote my CEO message

SITUATION

A draft of the CEO's quarterly all-staff message was generated by AI and lightly edited. It reads well. Sign-off is requested.

CHECKS

Are the claims, numbers and commitments verified against source material? Does it sound like this leader, not a generic executive voice? Has the leader authorized AI drafting for messages issued in their name? Would the audience feel misled if they knew AI drafted it?

ACTIONS

Verify facts and named commitments before sign-off, name one person who owns both the words and the commitments in them and retain the draft and sign-off record under your organization's records-retention policy.

PRINCIPLE

Trust drops when people feel messaging is inauthentic, regardless of where the first draft comes from.

We want AI to analyze open-text survey comments

SITUATION

Leadership wants AI to run thematic analysis on an engagement survey's open-text responses. It can be done in minutes, but speed isn't the same as accuracy.

CHECKS

Is personal data present, and if so, what is the lawful basis for using it and what is the retention period? Which groups are over- or under-represented in the dataset? Have respondents been told, in plain language, that AI will analyze their comments? Do they have a way to opt out of AI-analyzed feedback or request a human review of their comments?

ACTIONS

Aggregate by default and minimize personal data in outputs, publish a short explainer of what's being analyzed and why and state limitations in any reporting rather than presenting themes as fact. Where comments touch sensitive or distressing topics, route them to a named human reviewer rather than AI alone and confirm who handles a participant complaint if one comes up.

PRINCIPLE

AI can spot patterns, but it is your insight that decides if that pattern is real and if it matters.

We need to explain our AI use to the public

SITUATION

A customer-facing service uses AI and you've been asked to write the public messaging.

CHECKS

Can a non-expert read it back to you accurately? Are you describing what the AI does, specifically? If the decision is automated, what happens? Will the company be accountable?

ACTIONS

Draft a "how this works" explainer covering what the AI does, what data it uses, the human's role and the solution if something goes wrong, and test it on a non-expert before publishing.

PRINCIPLE

If you can't explain it in plain language, you're not ready to deploy it.

The chatbot that made a promise

SITUATION

Your chatbot gives a customer the wrong information and they act on it.

CHECKS

"The chatbot said it" is not a defence – a BC tribunal held Air Canada liable on exactly this fact pattern after its chatbot misstated a bereavement-fare policy (Moffatt v. Air Canada, 2024 BCCRT 149).

ACTIONS

Was the chatbot tested against actual policy, including sensitive topics? Was a human escalation route built in for high-stakes questions? Did the customer know they were speaking to AI?

PRINCIPLE

Treat your chatbot like every other spokesperson.

A closer look: chatbots as AI agents

The scenario above treats a chatbot as a single wrong answer. In practice, many customer-facing chatbots now qualify as AI agents: they retain memory across a conversation, decide what to ask next, and sometimes take actions (booking, refunding, escalating) without a person approving each step. Apply the same checks here that you'd apply to any AI agent: a named human owner, a tested escalation path for sensitive topics and disclosure that the audience is talking to AI, not a person.

We've been using it quietly

SITUATION

A team has used AI across drafting, research and analysis for months. Nothing was hidden, but nothing was disclosed either and leadership now wants a public commitment to responsible AI use.

CHECKS

Can you describe current use accurately to leadership? How will the team react if disclosure happens without their input?

ACTIONS

Map current use before it's shown to anyone, disclose to internal stakeholders before external ones and build governance forward rather than retroactively.

PRINCIPLE

Quiet adoption is still adoption. The longer it goes unnamed, the more there is to explain later.

Our agency used AI on our copy

SITUATION

You discover a long-standing agency has been using AI in work invoiced as bespoke. The contract never addressed it either way.

CHECKS

What does your agreement say, or not say, about AI use and accountability? Was the work good, and if so, is the concern one of principle or of cost?

ACTIONS

Add an AI-use clause to supplier contracts at next renewal specifying what must be disclosed and when, then audit your top suppliers now even if you don't have every answer yet.

PRINCIPLE

AI use by your suppliers and agencies is still your accountability.

Can I paste this into AI?

SITUATION

It's late in the day, you have three things to draft or summarize fast and the free AI tool in your browser could do all of them in minutes.

CHECKS

What class of data is this, public, internal, confidential or regulated? Which tool tier are you using, approved enterprise, reviewed business subscription or free consumer tool? If this leaked tomorrow, could you undo it?

ACTIONS

Default yes for public data in approved tools; for internal data, first confirm it isn't confidential, sensitive or contractually restricted. Default no for confidential or regulated data in anything non-approved, and when in doubt, follow the de-identification or anonymization process in Section 4 rather than simply removing names and numbers.

PRINCIPLE

The question isn't whether AI is safe. It's whether this data, in this tool, for this task, is safe.

Our AI went off script

SITUATION

A marketing tool with agentic features, one that can search, draft, post and reply on its own, has been running unattended for weeks. It's doing good work, until it responds to a customer complaint in a tone nobody approved and nobody notices for two days.

CHECKS

Who is the named human owner of this agent, and how often do they review its output? Can you explain, in plain language, what decisions it's allowed to make without a person in the loop? If it coordinates with other automated tools, could their combined activity look organized or intentional to an outside observer, even if no one designed it that way?

ACTIONS

Map every AI agent operating with any autonomy: what it can do without approval, who owns it and how someone would notice if it went wrong. Build in a human checkpoint proportional to what the agent can access and do in practice, including foreseeable failure modes, not just what it was designed to do.

PRINCIPLE

An agent that acts without a name attached to it is still your accountability.

Adapted from PRSA's Ethical Standard Advisory #22, "AI Agents in Public Relations."

09

Getting started checklist

Whether you’re just getting started or refining your process.

09 | GETTING STARTED CHECKLIST

CPRS's AI Guidance already tells you how to build the basic infrastructure:

an approved tool list	a review and sign-off process	an escalation path and	a basic AI use policy.

If none of that exists yet, start there.

There's real demand for this: TELUS's 2026 AI Trust Atlas found 59% of Canadians want to learn more about AI. What follows are five practical actions that will help you get started in day-to-day communications and public relations tasks.

☑	1. Assess your outputs against a risk structure. Use the questions in Section 4 to set a tier for anything AI touches.	SECTION 04
☑	2. Develop your own verification model. Build from Section 4's tiered review chain and the SAR framework rather than starting from a blank page. Adapt it to your team's actual workload.	SECTION 04
☑	3. No one approves their own AI output for Tier 2 and above. Section 5's review and sign-off matrix names who checks what. Keep that separation even on a small team, where one person may hold more than one check but not all of them on the same item.	SECTION 05
☑	4. Disclose your use and your verification process. Section 6's Trust Test and model disclosure statement cover this. Include the verification and human review in your disclosure.	SECTION 06
☑	5. Build and test a crisis response plan for AI. Section 5's escalation route is the starting point. Test it before you need it: run a tabletop exercise using one of Section 8's scenarios.	SECTION 05/08

10

Resources

Some of the studies and thought leadership behind this toolkit.

10 | RESOURCES AND LINKS

These resources were used in the development of the toolkit. They are listed in order of first appearance in the document.

#	SOURCE	LINK	APPEARS
1	**CPRS AI Guidance**	Approved by CPRS National Board, August 2026. https://www.cprs.ca/Advocacy/AIinPRPanel-(1)	Section 1
2	**CPRS 2025 State of Public Relations in Canada**	CPRS, October 2025. cprs.ca (PDF)	Section 1
3	**BCG — GenAI Adoption in Corporate Affairs and Communications**	Boston Consulting Group, 2026. bcg.com	Section 1
4	**IPR — The Communicator's Role in Driving Generative AI Adoption**	Institute for Public Relations, February 2026. instituteforpr.org (PDF)	Section 1 (repeated in Section 3)
5	**CIPR Inside — AI and Digital Ethics for Communicators**	CIPR Inside Committee, May 2026. cipr.co.uk / newsroom.cipr.co.uk	Section 1 (this toolkit's source guide)
6	**Global Alliance for PR and Communication Management — Venice Pledge / Responsible AI Principles**	globalalliancepr.org	Section 1
7	**OPC — Principles for responsible, trustworthy and privacy-protective generative AI**	Office of the Privacy Commissioner of Canada.	Section 2
8	**OPC — PIPEDA Finding #2026-002 (OpenAI)**	Office of the Privacy Commissioner of Canada. priv.gc.ca	Section 2
9	**OPC — PIPEDA Finding #2026-004 (X Corp / xAI, Grok)**	Office of the Privacy Commissioner of Canada, June 2026. priv.gc.ca	Section 2 (repeated in Section 3)

#	SOURCE	LINK	APPEARS
10	**Center for Countering Digital Hate — Grok floods X with sexualized images**	Center for Countering Digital Hate. counterhate.com	Section 2
11	**TBS — Guide on the Use of Generative AI (FASTER principles)**	Treasury Board of Canada Secretariat. canada.ca	Section 2 (repeated in Section 3)
12	**TBS — Directive on Automated Decision-Making**	Treasury Board of Canada Secretariat. canada.ca	Section 2 (repeated in Sections 3 and 9)
13	**Act Respecting the Protection of Personal Information in the Private Sector (Quebec P-39.1 / Law 25)**	CQLR c P-39.1, consolidated to April 1, 2026. legisquebec.gouv.qc.ca	Section 2, 4 and 6
14	**IBM Security / Ponemon — Cost of a Data Breach Report 2025: The AI Oversight Gap**	IBM Security / Ponemon Institute, 2025. ibm.com | https://canada.newsroom.ibm.com/2025-09-03-IBM-R-Study-Shadow-AI-Use-Surges-as-Canadian-Workers-Outpace-Employers-in-AI-Adoption	Section 2
15	**TELUS — 2026 AI Trust Atlas**	TELUS, in partnership with Leger, January 2026. assets.ctfassets.net (PDF)	Sections 2, 3, 4, 6 twice, and 9
16	**CRIC — Guiding Principles for AI use in Market Research**	Canadian Research Insights Council, AI Ethics Sub-committee, July 2026. canadianresearchinsightscouncil.ca	Section 3, 4 and 5
17	**Dr. Alex Sévigny — 4 leadership capabilities for communicators in the AI age**	alexsevigny.ca, July 2026. alexsevigny.ca	Section 3
18	**Dr. Alex Sévigny — Efficacy vs effectiveness vs efficiency in communications and marketing**	alexsevigny.ca, October 2025. alexsevigny.ca	Section 3

#	SOURCE	LINK	APPEARS
19	**Gerlich, M. (2025) — AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking**	https://doi.org/10.3390/soc15010006	Section 3
20	**European Commission — Code of Practice on marking and labelling AI-generated content**	digital-strategy.ec.europa.eu	Section 5
21	**Schilke, O. and Reimann, M. (2025) — The transparency dilemma: How AI disclosure erodes trust**	Organizational Behavior and Human Decision Processes (OBHDP), 188, 104405. sciencedirect.com	Section 6
22	**Moffatt v. Air Canada**	2024 BCCRT 149 https://decisions.civilresolutionbc.ca/crt/crtd/en/item/525448/index.do?q=moffat	Section 8
23	**PRSA Board of Ethics and Professional Standards — Ethical Standard Advisory #22: AI Agents in Public Relations**	August 2026 https://www.prsa.org/docs/default-source/about/ethics/esa-22.pdf?sfvrsn=867ef5c2_1	Section 8

CPRS | SCRP

Communications and Public Relations Society of Canada | 2026

www.cprs.ca