

AI Use Policy

An adaptable policy aligned with CQIP, written to meet UK and EU law, with notes for the United States and Canada

H D Fraser & Associates · Current as of 30 September 2026

Free to use and adapt, as a guide only.

This is a guide, not legal advice. It is a template written to meet UK and EU law as of the date above. Laws change and apply differently by sector, size and location, so have a qualified lawyer review your adapted version before you adopt it. See [UK and EU legal basis](#) for the laws it maps to, and [Notes for US use](#) and [Notes for Canadian use](#) for use outside the UK and EU.

Contents

1. AI Use Policy
 - How to use this template
 1. Purpose and scope
 2. Principles
 3. Roles and responsibilities
 4. Approved tools
 5. What information may be used with AI
 6. Prohibited uses
 7. Human review and automated decisions
 8. Transparency and labeling
 9. Accuracy, authorship and intellectual property
 10. Fairness and higher-risk uses
 11. Security
 12. AI literacy and training
 13. Incidents, records and review
 - Appendix A: AI tool register
2. UK and EU legal basis
3. Notes for US use
4. Notes for Canadian use

AI Use Policy

How to use this template

Replace every item in square brackets, remove any section that does not apply, and keep the rest in your own words. Before adopting it:

- ☐ Replace [Organization], [AI Lead], [Data Protection Lead] and the other placeholders
- ☐ List the AI tools you actually use in Appendix A
- ☐ Check the jurisdiction notes for every country where you operate or have customers
- ☐ Have a qualified lawyer review the adapted policy
- ☐ Approve it, date it, and tell everyone it applies to

The policy follows CQIP, the Content Quality and Intelligence Policy: AI is a tool that drafts, a named person approves, and every output can be traced to its source.

1. Purpose and scope

[Organization] uses AI tools to work faster and better, while keeping control of its information, its decisions and its reputation. This policy sets out how.

It applies to everyone who works for or on behalf of [Organization], including employees, contractors, volunteers and suppliers, whenever they use an AI tool for [Organization]'s work. An AI tool is any software that generates text, images, audio, video or code, makes predictions or recommendations, or acts on a person's behalf. That includes chat assistants, AI features inside other software, AI agents in hosting or office tools, and systems [Organization] builds itself.

This policy works alongside [Organization]'s data protection, information security, acceptable use and equality policies. Where they conflict, the stricter rule applies.

2. Principles

Every use of AI at [Organization] follows these eight principles. The first five come from CQIP; the last three restate data protection and consumer law in plain terms.

Principle	What it means in practice
Draft, not decide	AI output is a draft. A named person reviews and approves it before it is used, sent or published.
A named person is accountable	The person who approves AI output owns it, as if they had written it themselves.
Air gap by default	No AI tool is connected to [Organization]'s systems, data or accounts unless it has been approved for that purpose.

Principle	What it means in practice
Connect only when it earns its place	A live connection needs a defined task and approval, and is removed when it is no longer needed.
Use once, keep the benefit	Where AI helps, keep the result in documents and processes [Organization] owns, so work does not stop if a tool changes or goes away.
Traceable to a source	Facts, figures and claims in AI output are checked against a reliable source before use.
The least information needed	Only the minimum information a task needs goes into an AI tool, and never more than section 5 allows.
Open about AI	People are told when they are dealing with AI, or with content AI has generated or significantly altered.

3. Roles and responsibilities

Role	Responsible for
[Senior Owner], a director or equivalent	Approving this policy, accepting residual risk, and signing off higher-risk uses (section 10)
[AI Lead]	Keeping the AI tool register, assessing new tools, running training, and reviewing this policy
[Data Protection Lead] or Data Protection Officer	Advising on personal data, completing data protection impact assessments, and handling data subject requests
Managers	Making sure their teams know the policy and use only approved tools, and approving AI use in their area
Everyone	Using AI only as this policy allows, reviewing what they approve, and reporting problems straight away

In a small organization one person may hold several roles, but no one should approve their own higher-risk use.

4. Approved tools

Only tools listed in the AI tool register (Appendix A) may be used for [Organization]'s work, and only for the uses listed there. Personal or free accounts may be used only with Public information (section 5).

To add a tool, ask the [AI Lead]. A tool is approved only when every answer below is known and acceptable:

1. **Account.** Is it a business account in [Organization]'s name, with access that can be removed when someone leaves?
2. **Training.** Does the provider use our inputs or outputs to train its models? If so, can that be switched off?

3. **Location.** Where is data processed and stored? If outside the UK or EU, what transfer safeguard applies?
4. **Contract.** Are there written data processing terms covering confidentiality, security, sub-processors and deletion?
5. **Retention.** How long does the provider keep prompts, files and outputs, and can we delete them?
6. **Security.** Does it support single sign-on or multi-factor authentication, and keep access logs?
7. **Exit.** Can we disconnect it, remove its keys, and keep our work if we stop using it?
8. **Risk.** Will it process personal data, or support a higher-risk use under section 10? If so, a data protection impact assessment comes first.

Where a tool connects using [Organization]'s own provider account ("bring your own key"), the key is held by the [AI Lead], used only for the approved task, and removed when the task ends.

5. What information may be used with AI

Check the class of the information before it goes into any AI tool. If unsure, treat it as the higher class and ask the [Data Protection Lead].

Class	Examples	Personal or free tools	Approved tools
Public	Published web pages, press releases, public reports	Allowed	Allowed
Internal	Internal procedures, meeting notes, drafts with no personal data	Not allowed	Allowed
Confidential	Contracts, pricing, financials, source code, client material, anything under an NDA	Not allowed	Only tools approved for Confidential, with the information owner's agreement and where the contract allows
Personal data	Names, contact details, customer and staff records, recordings of people	Not allowed	Only tools approved for personal data, with a lawful basis, removing names wherever the task allows
Special category, criminal and children's data	Health, ethnicity, religion, sexual orientation, biometrics, criminal records, any data about children	Not allowed	Not allowed, unless the [Senior Owner] approves after an impact assessment and legal advice
Credentials	Passwords, access keys, bank or card details	Never	Never

6. Prohibited uses

No one may use AI at [Organization] to do any of the following. The first group mirrors practices banned outright under the EU AI Act.

Banned in the EU, and banned here

- Manipulating people's behavior with deceptive or subliminal techniques in ways that could cause significant harm
- Exploiting anyone's vulnerability due to age, disability, or social or economic situation
- Scoring people on their social behavior or personal traits in ways that lead to unfair treatment
- Predicting whether a person will commit a crime based only on profiling or personality traits
- Building facial recognition databases by scraping images from the internet or CCTV
- Inferring emotions of staff or students, except for medical or safety reasons
- Sorting people by biometric data to infer race, political opinions, trade union membership, religious beliefs, sex life or sexual orientation
- Generating intimate images of a real person without their consent, or any sexual imagery of children

Also not allowed at [Organization]

- Creating realistic images, audio or video of a real person without their consent, or passing off AI content as a real event
- Writing fake reviews, testimonials or endorsements, or presenting AI content as independent opinion
- Making a final decision about a person's job, pay, service, credit or access without the human review in section 7
- Monitoring staff with AI without a documented purpose, an impact assessment and notice to the people affected
- Putting information into an AI tool that section 5 does not allow
- Using AI to get around [Organization]'s security, or anyone else's

7. Human review and automated decisions

Every piece of AI output is reviewed by a named person before it is used. The more it matters, the deeper the review:

Output	Minimum review
Internal notes, summaries, ideas	The user reads it before relying on it
Anything sent outside [Organization] or published	The user checks facts, tone and sources, and approves it by name
Advice, contracts, regulated statements, prices	A second person with the right expertise approves it
Decisions about a person (below)	Meaningful human review, before the decision takes effect

Decisions about people. [Organization] does not let AI make a decision with a legal or similarly significant effect on a person, such as hiring, dismissal, pay, credit or access to a service, without meaningful human review. Meaningful means the reviewer understands how the output was produced, sees the relevant information, and

has the authority to change the result. The person affected is told that AI was involved, can ask for a human to reconsider, and can challenge the decision.

8. Transparency and labeling

[Organization] is open about where it uses AI:

- **Chat assistants.** Anyone talking to an AI system on [Organization]'s behalf is told so at the start, unless it is obvious.
- **Images, audio and video.** Realistic content that AI has generated or altered is labeled as such wherever it is published.
- **Published text.** Text that AI generated to inform the public is labeled, unless a named person has reviewed it and [Organization] takes editorial responsibility for it.
- **Privacy notices.** [Organization]'s privacy notices say where AI processes personal data, why, and what rights people have.
- **Recruitment.** Job applicants are told if AI is used to screen or assess them.

Labels are plain words, such as "This image was created with AI", not symbols alone.

9. Accuracy, authorship and intellectual property

AI tools can state wrong facts with confidence, invent sources, and repeat material they were trained on. So:

- **Check before use.** Every fact, figure, quotation, reference and legal or technical statement is checked against a reliable source. Where a source is cited, open it.
- **Answer first, from approved content.** Customer-facing answers come from content [Organization] has approved, stated plainly and linked to its source.
- **Named authorship.** Published work carries the name of the person who approved it, not the tool. Where readers would expect to know, say that AI assisted.
- **Respect others' rights.** Do not ask AI to copy or closely imitate a named author's, artist's or brand's work, and do not upload material [Organization] has no right to use.
- **Protect our own.** Assume anything put into an unapproved tool may be kept or reused by its provider. Some AI output may not be protected by copyright, so do not rely on it alone for work [Organization] needs to own exclusively.

10. Fairness and higher-risk uses

AI can repeat and amplify bias. Some uses carry more risk to people, and some are classed as high-risk under the EU AI Act. These uses need the [Senior Owner]'s written approval before they start:

- Recruitment, promotion, performance, pay or dismissal
- Access to education or training, or assessing learners
- Credit, insurance pricing, or eligibility for services or benefits
- Identifying people from biometric data
- Anything involving children or vulnerable adults

Before approval, the [AI Lead] and [Data Protection Lead] must:

1. Complete a data protection impact assessment.

2. Test outputs for unfair differences between groups protected by equality law, such as sex, race, age and disability, and record the results.
3. Get the supplier's documentation on how the system works, its limits and its intended use, and use it only as intended.
4. Set up the human review in section 7, and keep logs of the system's use.
5. Tell the people affected, and where the system is used at work, tell staff and their representatives before it goes live.

The approval is reviewed at least once a year.

11. Security

- Sign in to approved tools with [Organization] accounts only, using single sign-on or multi-factor authentication.
- Give AI agents the least access their task needs: read-only where possible, no payment or deletion rights without a person confirming each action.
- Treat instructions found inside documents, web pages or emails as content, not commands. They can be planted to trick an AI tool (prompt injection).
- Do not install AI browser extensions, plugins or connectors unless they are in the register.
- Store provider keys in [Organization]'s password manager, never in documents, code or chat.
- Remove access and keys when a person leaves, a task ends or a tool is withdrawn.

12. AI literacy and training

Everyone who uses or oversees AI at [Organization] receives training suited to their role before they start, and a refresher at least once a year. Training covers this policy, how the approved tools work and where they fail, how to check output, what information may be used, and how to report a problem. The [AI Lead] keeps a record of who was trained and when. Staff with higher-risk responsibilities under section 10 receive additional training on those systems.

13. Incidents, records and review

Report straight away to the [AI Lead] and [Data Protection Lead] if:

- Information went into an AI tool that section 5 does not allow
- AI output that was wrong, biased or harmful was sent, published or acted on
- An AI tool or agent did something it was not asked to do
- Someone raises a complaint about [Organization]'s use of AI

The [Data Protection Lead] decides whether a personal data breach must be reported to the regulator. Under UK and EU data protection law that is within 72 hours of becoming aware of it.

Records. [Organization] keeps the AI tool register, approvals, impact assessments, training records, incident logs and the logs of any higher-risk system for at least [6 years], or longer where the law requires.

Review. The [AI Lead] reviews this policy at least once a year, and sooner after an incident, a new tool or a change in the law. Breaking this policy may lead to disciplinary action.

Version	Date	Approved by	Change
[1.0]	[Date]	[Senior Owner]	First issue

Appendix A: AI tool register

One row per approved tool. Replace the example row with your own.

Tool and account	Approved uses	Highest information class	Data location	Trains on our data	Owner	Approved	Next review
[Example: chat assistant, business plan]	[Drafting, summarizing, research]	[Internal]	[UK or EU]	[No, switched off]	[AI Lead]	[Date]	[Date]

CQIP, the Content Quality and Intelligence Policy, is a methodology standard authored by H D Fraser MSc and published through The Content Framework, not a technology stack. It is implemented across static PHP, WordPress and other systems, and this policy applies whatever tools an organization uses.

UK and EU legal basis

This is a guide, not legal advice. It shows which laws the AI Use Policy is written to meet and where. Laws change and apply differently by sector and size, so take advice from a qualified lawyer before relying on it.

United Kingdom

The UK has no general AI law. The King's Speech of 19 May 2026 announced none, so AI use is governed by data protection, equality, consumer and copyright law.

Law	What it requires of a business using AI	In force	Policy section
UK GDPR and Data Protection Act 2018	A lawful basis, transparency, data minimization, security, written processor terms, transfer safeguards, impact assessments for high-risk processing, breach reports within 72 hours	Yes	4, 5, 8, 11, 13
UK GDPR Articles 22A to 22D, inserted by the Data (Use and Access) Act 2025, s.80	Solely automated decisions with legal or similarly significant effects are allowed only with safeguards: tell the person, let them make representations, give human intervention, let them contest. Stricter limits apply where special category data is used.	5 February 2026	7
Data (Use and Access) Act 2025, other changes	New recognized legitimate interests basis. A duty to handle data protection complaints.	5 February 2026; complaints from 19 June 2026	5, 13
Equality Act 2010	No direct or indirect discrimination, including through automated tools, and reasonable adjustments for disabled people	Yes	10
Digital Markets, Competition and Consumers Act 2024	Fake reviews and misleading practices are banned, with direct fines by the Competition and Markets Authority	6 April 2025	6, 8

Law	What it requires of a business using AI	In force	Policy section
Copyright, Designs and Patents Act 1988	Copying protected works without permission infringes. The government has said it will not reform copyright for AI until it is confident reforms meet its objectives (impact assessment, 18 March 2026).	Yes	9

The policy is stricter than UK law on automated decisions: it requires meaningful human review before any significant decision takes effect, which also meets the EU rule below. The ICO's guidance on AI and data protection is under review for the 2025 Act, and its draft guidance on automated decision-making, consulted on until 29 May 2026, will feed a statutory code.

European Union

The EU AI Act (Regulation (EU) 2024/1689) applies in stages. Its dates were changed by the Digital Omnibus on AI, Regulation (EU) 2026/1744, in force since July 2026. A business that uses AI is a "deployer"; most duties on model builders are not yours.

AI Act obligation	What a deployer must do	Applies from	Policy section
Prohibited practices (Article 5)	Never use AI for the banned practices, now including non-consensual intimate imagery and child sexual abuse material, added by the Omnibus	2 February 2025	6
AI literacy (Article 4)	Take measures to support AI literacy among staff who use AI. The Omnibus softened this from "ensure".	2 February 2025	12
Transparency (Article 50)	Label deepfakes. Label AI text published to inform the public on matters of public interest, unless a person reviewed it and someone holds editorial responsibility. Tell people when emotion recognition or biometric categorization is used on them.	2 August 2026	8
High-risk systems in Annex III, such as hiring, education, credit and essential services (Articles 26 and 27)	Use as instructed, assign human oversight, monitor, keep logs for at least 6 months, inform workers and their representatives before workplace use, inform people affected. Public bodies, and deployers doing credit scoring or life and health insurance pricing, must also complete a fundamental rights impact assessment.	2 December 2027	7, 10, 13

AI Act obligation	What a deployer must do	Applies from	Policy section
High-risk AI in regulated products (Annex I)	As above, for AI built into products such as machinery and medical devices	2 August 2028	10

Other EU law

Law	What it requires	Applies from	Policy section
EU GDPR, Article 22	Solely automated decisions with legal or similarly significant effects are allowed only by contract, law or explicit consent, with human intervention and a right to contest	Yes	7
EU GDPR, Article 35	An impact assessment before high-risk processing	Yes	4, 10
Product Liability Directive (EU) 2024/2853	Software, including AI, counts as a product for no-fault liability. This matters if you supply AI-enabled products or services.	Transposed by 9 December 2026	9, 10
EU adequacy decision for the UK	Personal data can flow from the EU to the UK without extra safeguards	Renewed 19 December 2025, until 27 December 2031	4

The European Commission published a voluntary Code of Practice on marking and labeling AI-generated content on 10 June 2026. Following it is a practical way to meet section 8.

Sources

- Data (Use and Access) Act 2025, section 80 and Commencement No. 6 Regulations, SI 2026/82
- ICO: artificial intelligence guidance hub
- ICO: consultation on draft automated decision-making guidance
- ICO: how do we ensure fairness in AI?
- Digital Markets, Competition and Consumers Act 2024 commencement, SI 2025/272
- Copyright and AI: impact assessment, March 2026
- House of Commons Library: King's Speech 2026, science and technology
- Digital Omnibus on AI, Regulation (EU) 2026/1744
- AI Act Service Desk: Article 26, Article 27, Article 50
- European Commission: Code of Practice on marking and labeling AI-generated content
- General Data Protection Regulation (EU) 2016/679
- Product Liability Directive (EU) 2024/2853
- ICO: receiving personal information from the EEA (UK adequacy)

Notes for US use

This is a guide, not legal advice. The United States has no single national AI law, and state laws are changing quickly. Take advice from a lawyer qualified in each state where you operate before adopting the policy.

Federal position

There is no federal AI statute. Existing federal law still applies to how a business uses AI:

- **Deceptive practices.** Section 5 of the FTC Act covers false claims about AI and deceptive uses of it. The FTC's rule on consumer reviews and testimonials (16 CFR Part 465, in effect since 21 October 2024) bans fake reviews, expressly including AI-generated ones.
- **Employment.** Title VII's disparate impact rule is written into the statute and applies to AI hiring tools, although the EEOC removed its AI guidance in January 2025 and federal agencies have been told to deprioritize disparate impact enforcement.
- **Copyright.** The US Copyright Office requires human authorship. Prompts alone are not enough, but human selection, arrangement or modification of AI output can be protected (report of 29 January 2025).

Federal stance on state AI laws. Executive Order 14365 of 11 December 2025 seeks a single national framework. It set up a Justice Department task force to challenge state AI laws and directed agencies to identify "onerous" ones. A White House framework (March 2026) and a draft preemption bill (June 2026) have followed, but no federal law has yet overridden state AI laws. Keep complying with state law, and watch for change.

State laws

The state laws most likely to affect a business using AI, by state:

State	Law	What a business using AI must do	In effect
California	Privacy Protection Agency regulations on automated decisionmaking technology	Risk assessments for covered processing; notice, opt-out and access rights where automated technology makes significant decisions about consumers	1 January 2026; automated decision duties 1 January 2027
California	Civil Rights Council employment regulations	Anti-discrimination law applies to automated decision systems in employment; keep related records for 4 years	1 October 2025
California	Bolstering Online Transparency (BOT) Act	Do not use a bot to mislead people about its artificial identity to sell something or influence a vote; clear disclosure is a defense	1 July 2019

State	Law	What a business using AI must do	In effect
Colorado	SB 26-189, which replaced the Colorado AI Act before it took effect	Notice when consumers interact with automated decision technology; a plain explanation within 30 days of an adverse decision; data correction; human review and reconsideration; records for 3 years	1 January 2027
Illinois	Human Rights Act, as amended by HB 3773	No discriminatory use of AI in employment decisions, no zip codes as a proxy for protected traits, and notice to employees and applicants	1 January 2026
Illinois	AI Video Interview Act	Explain and get consent before AI analyzes video interviews; delete videos within 30 days of a request	1 January 2020
New York City	Local Law 144	A bias audit of automated employment decision tools within the past year, published results, and notice to candidates 10 business days before use	5 July 2023
Texas	Responsible AI Governance Act	No intentionally harmful uses, such as encouraging self-harm or intentional unlawful discrimination	1 January 2026
Utah	Artificial Intelligence Policy Act	Say it is AI when a consumer asks, and up front in high-risk interactions in regulated occupations	Amended 7 May 2025; ends 1 July 2027
Virginia, Connecticut and others	Comprehensive privacy laws	Let consumers opt out of profiling used for decisions with legal or similarly significant effects	Varies by state

Several states have also passed laws on "companion" chatbots, mostly to protect minors. They generally exclude customer service bots, but check them if your chatbot offers personal or emotional support. California's AI Transparency Act (from 2 August 2026) applies to large generative AI providers, not to businesses that use their tools.

Adapting the policy for US use

The policy works in the US as written, because its UK and EU rules are mostly stricter. Make these changes:

1. **Section 4, question 3.** Replace the UK and EU transfer wording with the states whose residents' data the tool processes, and check each state's privacy law.
2. **Section 5.** Add "sensitive data" as defined in the state privacy laws that apply to you, and health information covered by HIPAA if you handle it.

3. **Section 7.** Keep the human review rule. It meets Colorado's reconsideration duty and the California and state privacy opt-out rights. Add a named contact who explains adverse decisions within 30 days.
4. **Section 8.** Keep chatbot disclosure on by default. It meets the California BOT Act and Utah law without further work.
5. **Section 10.** Add a bias audit before any automated hiring tool is used on New York City candidates, and notice to employees and applicants in Illinois.
6. **Section 13.** Replace the 72-hour breach rule with the state breach notification deadlines that apply, and keep employment decision records for at least 4 years to meet California.
7. **Laws named in sections 6 and 8.** Keep the EU-derived bans as company rules; they are good practice even where not legally required.

Sources

Some entries rely on secondary sources and should be confirmed before relying on them, in particular the Colorado signing date and the Utah amendments.

- Executive Order 14365, Federal Register
- FTC: final rule banning fake reviews and testimonials
- FTC: proposed policy statement on AI accuracy, July 2026
- US Copyright Office: copyrightability report, January 2025
- California Privacy Protection Agency: CCPA regulation updates
- California Civil Rights Council: automated decision systems regulations
- Colorado General Assembly: SB 26-189 and Seyfarth summary
- Illinois AI Video Interview Act
- New York City: automated employment decision tools
- Texas HB 149, enrolled text
- Davis Polk: Utah AI Policy Act amendments
- Code of Virginia, § 59.1-577

Notes for Canadian use

This is a guide, not legal advice. Canada has no federal AI statute in force, and provincial rules differ. Take advice from a lawyer qualified in each province where you operate before adopting the policy.

Federal position

Canada has no AI statute in force. The proposed Artificial Intelligence and Data Act died with Bill C-27 in January 2025.

- **PIPEDA** remains the federal private-sector privacy law. Its principles of accountability, identified purposes, consent, limited collection and use, accuracy, safeguards and openness all apply to personal information used with AI.
- **Bill C-36, the Protecting Privacy and Consumer Data Act**, was introduced on 15 June 2026 and has had first reading only. It would replace PIPEDA's privacy rules and require organizations to disclose when they use automated decision systems that could have a legal or similarly significant effect. It is not law yet.
- **Guidance, not law.** The Privacy Commissioner's principles for generative AI (7 December 2023) and the voluntary Code of Conduct on advanced generative AI (September 2023) set expectations. Canada's National AI Strategy was announced on 5 June 2026.
- **Competition Act.** False or misleading representations, including AI-generated claims, are prohibited.
- **Human rights.** Federal and provincial human rights laws apply to discrimination whatever its cause, including automated systems.

The federal Directive on Automated Decision-Making applies only to federal government institutions.

Provincial laws

Province	Law	What a business using AI must do	In effect
Quebec	Private sector privacy act (Law 25), s.12.1	Tell a person when a decision about them is based exclusively on automated processing; on request, explain the information, reasons and main factors used; let them have it reviewed by a person	22 September 2023
Quebec	Same act, ss.3.3 and 17	A privacy impact assessment before acquiring or developing a system that processes personal information, and before sending personal information outside Quebec, with a written agreement	22 September 2023

Province	Law	What a business using AI must do	In effect
Ontario	Employment Standards Act, 2000 and O. Reg. 476/24	Employers with 25 or more employees must say in public job postings if AI is used to screen, assess or select applicants	Postings from 1 January 2026
Ontario	Human Rights Code	Applies to AI used by private organizations; the Human Rights Commission offers an AI impact assessment tool	Yes
Alberta and British Columbia	Personal Information Protection Acts	No AI-specific rules in force; general privacy duties apply	Yes

Manitoba and Ontario have AI laws for the public sector only. Nova Scotia has proposed job posting disclosure like Ontario's.

Adapting the policy for Canadian use

The policy meets most Canadian requirements as written. Make these changes:

- Section 4, question 3.** Name where data is stored, and for Quebec residents complete a privacy impact assessment and written agreement before data leaves Quebec.
- Section 5.** Treat all personal information as needing meaningful consent for the AI purpose, in line with PIPEDA.
- Section 7.** Keep the human review rule. It already meets Quebec's s.12.1 and anticipates Bill C-36. Add a named person who explains the main factors behind a decision when asked.
- Section 8.** In Ontario, add AI disclosure to public job postings if you have 25 or more employees and AI is used on applicants.
- Section 10.** Add a privacy impact assessment for any new system that processes personal information in Quebec, not only higher-risk ones.
- Section 13.** Replace the 72-hour rule with PIPEDA's duty to report breaches creating a real risk of significant harm "as soon as feasible", and Quebec's equivalent.
- Language.** In Quebec, make the policy and notices available in French.

Sources

- LEGISinfo: Bill C-36 (45-1)
- PIPEDA, full text
- Privacy Commissioner: principles for generative AI
- Voluntary Code of Conduct on advanced generative AI
- Canada's National AI Strategy announcement
- Quebec private sector privacy act, CQLR c. P-39.1 (French)
- Ontario: AI disclosure in publicly advertised job postings
- Ontario Human Rights Commission: AI impact assessment
- Alberta PIPA overview
- Competition Bureau: false or misleading representations
- Treasury Board: Directive on Automated Decision-Making