

Canada Bill C-27 (AIDA) - Proposed Amendments

February 2024

Below, Meta Platforms, Inc. has proposed important amendments to AIDA within C-27. Proposed amendments are represented in **red text**.

Priority Amendments

1. Removing Class 4 of Schedule 2 to ensure the “high-impact” categories enumerated are tailored to risk, and include only those AI use cases that present a uniform level of risk, such as those possibly producing legal or similarly significant effects.
2. Tailoring GPAI definitions to reflect the way AI systems operate in practice, and deferring further GPAI specific provisions to future regulations developed with multi-stakeholder collaboration where requirements can be shaped to be technically feasible and appropriately placed on the actor best positioned to comply with them.
3. Removing the AI and Data Commissioner’s audit and remote search authorities.

Additional Overarching Amendments

4. Providing additional clarity to the definition of Artificial Intelligence Systems
5. Adding an exception for research and open source systems and models.

Text of AIDA and ISED Minister’s Amendments with **Suggested Amendments in Red**

- 1. Removing Class 4 of Schedule 2 to ensure the “high-impact” categories enumerated are tailored to risk, and include only those uses that present a uniform level of risk, such as those possibly producing legal or similarly significant effects.**

Rationale

We welcome the introduction of more legal certainty on the types of use cases considered “high-impact” and agree that an exhaustive list should be incorporated toward that end. However, the proposed Schedule 2 list presents multiple concerning classifications that would steer the bill away from its initial risk-based foundation and lead to an overinclusion of AI systems with severe consequences for innovation, preventing Canada from fostering AI development, protecting expression, and maintaining competitiveness in the global marketplace.

Most significantly, the proposed classification of content moderation and prioritization systems as “high-impact” is misguided. The use of AI systems in these areas does not categorically rise to the same level of risk as the other areas contemplated in the high-impact list, such as employment, health care services, administrative decisions about an individual, and assistance of peace officers. Evidently, these are all scenarios in which the use of AI systems can give rise to effects of legal or similar significance. The same is not true for content prioritization.

At a minimum, equating content prioritization to actual high-impact scenarios will create disproportionate—and burdensome—obligations for a wide range of systems that are providing content that people enjoy, as well as helping to maintain the integrity of those systems. The consequences will range from harming the innovation ecosystem in Canada to less effective content moderation by platforms. At worst, this will inappropriately, and ineffectively, tackle content issues with obligations intended for AI systems.

Moreover, this approach will subject content policies to the broad regulatory and enforcement powers under the bill, with provisions such as expansive audits and the ceasing of system operations, which create particular concerns about potential government overreach and limits on speech when introduced in the context of content moderation. This proposed classification represents an outlier in the field of both AI and content regulation and would set Canada on a divergent track from other major markets, including key trading partners like the U.S. and the EU.

**Amendment
Schedule 2**

Text

High Impact Systems - Uses

~~(4) The use of an artificial intelligence system in matters relating to
(a) the moderation of content that is found on an online communications platform, including a search engine or social media service; or
(b) the prioritization of the presentation of such content.~~

- 2. Tailoring GPAI definitions to reflect the way AI systems operate in practice, and deferring further GPAI specific provisions to future regulations developed with multi-stakeholder collaboration where requirements can be shaped to be technically feasible and appropriately placed on the actor best positioned to comply with them.**

Rationale

The regime proposed in the amendments for the regulation of general purpose AI systems is vague and overreaching. We are specifically concerned about several provisions that set forth detailed, prescriptive obligations that are infeasible, not appropriately tailored, and misaligned in their objectives. Consensus is still developing related to the best way to approach technical requirements for GPAI, which could be significant to the creation of a GPAI framework in Canada that is interoperable with other jurisdictions. For this reason, regulation and codes of conduct developed in collaboration with industry and other stakeholders are a more desirable approach to defining the technical specifications of these requirements.

We understand that, globally, there has been an increase in interest with respect to the regulation of foundation models and GPAI. Indeed, there are other international regulatory frameworks that have considered rules for the providers of GPAI or foundation models, but there are significant differences in their approach as compared to the breadth being considered here. For example, discussions around GPAI in the context of the EU AI Act differentiated between whether systems are regular GPAs or GPAs carrying systemic risk. The debates also revealed that requirements for providers of GPAs (that is, the entities or people putting the models on the market) are mostly limited to information sharing with the downstream provider (that is, the person that builds upon the GPAI for their specific purpose), to empower them to mitigate the risks relevant to the use case they plan to adopt.

These are important distinctions that are critical to a truly risk-based approach and to ensure the requirements are effective and technically feasible. Critically, the EU AI Act also recognized that the law ought to account for how open source models differ, and for the benefits they bring. In the U.S. the Executive Order has tasked various agencies with responsibilities concerning foundation models and generative AI. Notably, NIST is working on a companion for their Risk Management Framework to help define global standards. These processes are already underway and it would be desirable to wait for their conclusion so that Canada can leverage and build upon them.

Determining the right obligations proportionate to the risks being addressed, and at which point those will be most effective, is an important aspect of AI governance. Unfortunately the GPAI regime proposed in the amendments does not reflect this approach as incorporated by other international regulatory frameworks. Our revisions propose addressing this gap through robust tailoring of definitions to consider the different roles of actors in the AI value chain and create a truly risk-based approach to regulating GPAI.

The revisions also address confusion in the amendments between ‘general-purpose systems’ and ‘general-purpose models’—while the amendments rely on the former, the latter term should be separately defined. ‘General-purpose systems’ are built on models and each have distinct roles that should be clearly defined to avoid placing obligations on actors unable to meet them. The amendments refer to ‘general-purpose systems,’ but what the defined term actually describes are ‘general-purpose models’ that can be adapted for use. Our proposed revisions create clearer distinctions between the two terms and defer specific obligations for GPAI to future regulation where technical requirements and other developing international frameworks can be properly considered.

<p>Amendment Section 5</p>	<p><u>Text</u> Definitions general-purpose model system means an artificial intelligence system that is designed for use, or that is designed to be adapted for use, in many fields and for many purposes and activities, including fields, purposes and activities not contemplated during the system’s development and which requires further technical implementation by the deployer before its use. General purpose AI model does not include models that are released under a free and open source license, or a similarly permissive license, that allows deployers to use, reproduce, distribute, copy, create derivative works of, and make modifications to, the software, or models intended for the sole purpose of scientific research and development.</p> <p>general-purpose system means an AI system which is based on a general-purpose AI model, that has the capability to serve a variety of purposes, both for direct use as well as for integration in other AI systems.</p>
<p>Amendment Sections 7-8</p>	<p><u>Text</u> Strike the following sections (as proposed in the amendments), and replace with “general-purpose system requirements” section below.</p> <p>General-purpose system — first time 7(1)</p> <p>Records (2)</p> <p>Existing-system (3)</p> <p>General-purpose system — making available 8(1)</p> <p>Existing-system (2)</p>

	<p>Changes – general purpose system (8.1)</p> <p>Managing operations — general purpose system (8.2)</p> <p>General-purpose system requirements</p> <p>7 (1) A person who designs, develops, or makes available for use a general-purpose AI model or general-purpose AI system shall establish measures, in accordance with the regulations, to provide the deployer with technical and other documentation necessary for deploying the model for the specific use case.</p> <p>(2) Obligations established in regulations shall not apply to a person who designs, develops, or makes available general-purpose AI models that are made accessible to the public under a free and open license, or similarly permissive license, that allows for the access, usage, modification, and distribution of the model, and whose parameters, including the weights, the information on the model architecture, and the information on model usage, are made publicly available.</p>
--	---

3. Remove the AI and Data Commissioner’s audit and remote search authorities.

Rationale

The audit powers, which include physical and remote access searches, would establish a regulatory regime that is poorly tailored to the context of AI technology. The AIDC audit powers, including in particular the remote access authority proposed in the amendments should be excluded from the bill altogether. The audit provisions are an unprecedented measure that are not necessary or proportionate for the purposes of this legislation. While the sharing of information with regulators is a common part of assessing compliance, it's critical that such sharing adheres to the customs and norms of audit regimes (which do not include remote or physical searches, preserve privacy and protect sensitive business information). The audit powers proposed in the amendments fail to strike that balance. In particular, the remote access authority proposed for the AIDC presents a concerning breadth of authority with risks to the privacy of user data, as well as security vulnerabilities and possible exposure of proprietary information. It may also create conflicts of law with respect to privacy and online safety obligations providers are subject to in other jurisdictions. These authorities are misguided and would institute a highly intrusive and costly enforcement regime.

<p>Amendment Section 15</p>	<p><u>Text</u></p> <p>Strike the following sections (as proposed in the amendments)</p> <ul style="list-style-type: none"> • Audit 15(1) Qualifications (2) Powers of Commissioner (2.1)
------------------------------------	---

	<p>Remote access (2.2)</p> <p>Limitation on remote access (2.3)</p> <p>Assistance (3)</p> <p>Copy to person (3.1)</p> <p>Copy to Commissioner (4)</p> <p>Cost (5)</p> <p>Delegation (6)</p> <p>Minister may request audit 15.1</p> <p>For greater certainty 15.2</p>
--	--

Additional Overarching Amendments

4. Providing additional clarity to the definition of Artificial Intelligence Systems.

Rationale

While we welcome the improvements made with the latest round of amendments, which have significantly restricted the scope of the definition by removing systems that are only partially automated from its scope, the definition remains overbroad and susceptible to capturing any software. We suggest adopting a definition that revolves around software that a) is machine based and b) is able to learn over time. The amendment we propose aligns with globally accepted definitions that are already applied across industry. In particular, we use as a starting point the definition provided by the Expert Group on AI at the OECD, which captures the distinction between complex AI systems and general logic-based algorithms, and we have adapted it to better fit into the framework of this Act. We believe that the amended definition is able to target the novel risks that AI systems uniquely present, while avoiding an overly broad scope that may impose regulatory burdens on technologies that do not present the same challenges that the Act is intended to tackle.

<p>Amendment Section 5</p>	<p>Definitions</p> <p>Artificial intelligence system means a technological system that, using a model, makes inferences in order to generate output, including predictions, recommendations or decisions.</p> <p>Artificial intelligence system means a machine-based system that, given a set of human-defined objectives,</p> <ul style="list-style-type: none"> (i) receives machine and/or human-based data and inputs; (ii) perceives real and/or virtual environments; (iii) infers those perceptions into models; (iv) generates outputs options by interpreting the models; (v) makes predictions, recommendations or decisions, which influence the environments it interacts with; and (vi) uses those inputs, predictions, recommendations or decisions to iteratively improve its model performance.
-----------------------------------	--

5. Adding an exception for research and open-source systems and models.

Rationale

It is important to clarify that the Act is intended to be applied to commercially available products. This appears to be the intent, as specified by **4(a)** (*The purposes of the Act are to regulate international and interprovincial trade 30 and commerce in artificial intelligence systems....*) and as it is explicitly connected to the Consumer Privacy Protection Act. Nonetheless, to avoid any confusion, and to ensure development and innovation in the AI field continue to be fostered without deterrence, we recommend specifying that systems and models exclusively developed for research and scientific development are not scoped into the Act.

Similarly, AIDA should incentivize approaches that support Canada’s goals of fostering AI innovation. As a result, AIDA should include an exemption for AI systems that are released under an open source or similarly permissive license (e.g. a license where information such as machine-learning model code, model weights, fine-tuning enabling code etc. are released), in recognition of the critical role that open source development plays in driving innovation and delivering economic benefits from new technologies. In the coming years, access to AI technology, such as general purpose and foundation models, will play a similarly crucial role in driving AI research, development, innovation and adoption. It is essential, therefore, that AIDA facilitate widespread access to, and innovation in foundation models.

To do so, providers of AI systems and models should be exempt from the requirements of AIDA whenever they decide to make their models available under open source or similarly permissive licenses, as highlighted above. An approach of this type, which can be described as open innovation, would not only allow Canadian researchers, developers, and citizens to benefit from advances in AI, it will also contribute to the creation of higher-performing, safer, and more secure systems and models as a broad community is able to test, scrutinize and improve openly available models.

<p>Add under “Non-Application” Section 3(4)</p>	<p>This Act does not apply with respect to AI systems and models, including their output, specifically developed and put into service for the sole purpose of scientific research and development. It also does not apply to AI systems and models, including their output, released under a free and open source license, or a similarly permissive license.</p>
--	---