



1455 3rd Street
San Francisco, CA 94158

February 26, 2026

From:
Ann M. O’Leary
Vice President of
Global Policy
OpenAI

To:
The Honourable
Evan Solomon, P.C.,
M.P.
Minister of Artificial
Intelligence and
Digital Innovation
House of Commons
Ottawa, Ontario
K1A 0A6
Canada

CC:
The Honourable
Gary
Anandasangaree,
P.C., M.P.
Minister of Public
Safety

The Honourable
Sean Fraser, P.C.,
M.P.
Minister of Justice
and Attorney
General of Canada

The Honourable
Marc Miller, P.C.,
M.P.
Minister of Identity
and Culture

Dear Minister Solomon,

The events in Tumbler Ridge are an unspeakable tragedy, and our hearts remain with the victims, their families, and the entire community.

Thank you for convening Tuesday’s meeting, and thank you for the frank discussion you led along with Ministers Miller, Anandasangaree, and Fraser, as well as Parliamentary Secretary Noormohamed, on how to help prevent tragedies like this in the future.

Today, we write to offer immediate steps OpenAI is taking in response to what we have learned in the wake of these events. We remain committed to cooperating with law enforcement authorities on the investigation into the Tumbler Ridge tragedy, and we are committed to an ongoing partnership with federal and provincial governments.

As we shared with you and with Canadian law enforcement, in June 2025, OpenAI made a decision to shut down a ChatGPT account of the perpetrator of the Tumbler Ridge tragedy after detecting a violation of our usage policy. OpenAI’s automated system detected the account, and it was subsequently sent to human review to determine whether our usage policies were violated and whether the account warranted referral to law enforcement. Based on what we could see at that time the account was banned in June 2025, we did not identify credible and imminent planning that met our threshold to refer the matter to law enforcement.

OpenAI works continuously to identify potential warning signals of serious violence, while also protecting the privacy and security of the vast majority of people who use our tools responsibly. We seek to identify and take action to address critical harm risks on our platform, which include threats to human life. We also must ensure our products are secure and private, and can be safely used by hundreds of millions of users across the world. OpenAI requires that all development and deployment of our models advance both human safety and human rights.

Over the past several months, we have taken steps to strengthen our safeguards and made changes to our law enforcement referral protocol for cases involving violent activities, but our conversation with you this week underscored that Canadians expect continued concrete action and we heard that message loud and clear. In our meeting, you and the other Ministers stressed that no community should have to face this tragedy. We agree.

OpenAI will:

- **Continue to Strengthen our Enhanced Law Enforcement Referral Protocol.** Referrals to law enforcement based on conversations with ChatGPT involve complex, challenging decisions as we strive to protect our users’ privacy while also taking action when needed for public safety. That’s why we continue to make ongoing improvements to our referral protocol. Several months ago, we partnered with mental health, behavioural, and law

enforcement experts to help us refine our criteria for when conversations cross the line into an imminent and credible risk, meriting a law enforcement referral.

Mental health and behavioural experts now help us assess difficult cases, and we have made our referral criteria more flexible to account for the fact that a user may not discuss the target, means, and timing of planned violence in a ChatGPT conversation but that there may be potential risk of imminent violence. With the benefit of our continued learnings, under our enhanced law enforcement referral protocol, we would refer the account banned in June 2025 to law enforcement if it were discovered today.

Today, we commit to working with the Canadian government and experts to continue strengthening our enhanced law enforcement referral criteria based on the Tumbler Ridge tragedy and the Canadian context. This will include continuing to analyze how imminent and credible risk is assessed and transparency regarding our reporting to law enforcement.

- **Develop a Direct Point of Contact with Canadian Law Enforcement.** Per the request of the Ministers, we will establish direct points of contact with Canadian law enforcement authorities to ensure that we provide information expeditiously to Canadian authorities in cases where we make a law enforcement referral based on the potential for real world violence.
- **Embed Country and Community Context in our De-escalation Work:** Our models should respond appropriately when users are in distress or pursuing prohibited behavior, with an emphasis on de-escalation and user safety. When ChatGPT recognizes that users need help, we want to help our users find localized support in their own communities. We will expand on our commitment to directing users to relevant support resources, such as local helplines when a user's location is known, ensuring assistance is specific to the user's country or region.
- **Enhance our System to Detect Repeat Policy Violators:** OpenAI has a system in place that seeks to identify repeat policy violators, including those who have had their ChatGPT accounts shut down for violating our violent activities policy, and then seek to create a new account. Despite this detection system, after the name of the Tumbler Ridge perpetrator was released publicly, we discovered that the perpetrator had used a second ChatGPT account. We shared the second account with law enforcement upon its discovery. We commit to strengthening our detection systems to better prevent attempts to evade our safeguards and prioritize identifying the highest risk offenders. We further commit to periodically assessing the thresholds used by our automated systems for detecting potential violent activities.



1455 3rd Street
San Francisco, CA 94158

These immediate commitments are only the first step in the work we must do in partnership with the Canadian government to improve AI safety. We seek continued dialogue and we would welcome working with the Canadian government to convene local stakeholders and industry to develop best practices for law enforcement referrals and AI model behavior in cases involving potential violence, including unique considerations for youth.

In the months ahead, OpenAI will also engage with federal and provincial governments, our industry peers, and local stakeholders from a range of disciplines and communities to ensure we are collectively meeting the needs of Canadians as we continue to improve our models and safety policies.

Thank you again for convening this important conversation. We know that together we must do all we can to strengthen AI safety.

Sincerely,

A handwritten signature in black ink, appearing to read "Ann M. O'Leary". The signature is fluid and cursive, with a long, sweeping tail that loops back under the name.

Ann M. O'Leary
Vice President of Global Policy
OpenAI