

Hello,

I am writing to you in response to the online consultation about harmful content.

My name is Drew Wilson. In 2004, I started taking an interest in digital rights and began following the developments as they happened in Canada. I followed the copyright, patent, and privacy debates. In 2005, I took up the act of journalism and began writing about the events, analyzing these issues on Slyck and, eventually, ZeroPaid. In 2013, I founded freezenet.ca to carry on this endeavour at my own pace and continue to write news articles impacting digital rights around the world, not just in Canada.

As such, I have witnessed first hand both the enormous benefits the Internet brings and the pitfalls that we encounter over a very long period of time.

It's one of these pitfalls that the online harms debate wishes to address. With the consolidation of people on larger platforms like Twitter and Facebook, what does one do about the problems of harmful, hateful content and misinformation? The negative impacts are obviously well documented. From the suicide of Amanda Todd to the January 6<sup>th</sup> insurrection of the US Capitol to the harassment of visible minorities and members of the LGBT community to the vaccine misinformation (which many have cited as a reason for the slow uptake on vaccination rates in the US), online content has an impact on the real world and real people.

Of course, I don't just offer platitudes of saying the right thing. I've personally been the victim of harassing messages, and misinformation over the years. While they do not happen that often, it has happened. As such, I have experienced the negative impacts first hand. It's never easy and, at times, it can be quite distressing. Still, I have, luckily, been able to push past these incidences and move on to try and do my part to make the Internet a better place.

With having seen others start up small websites and having gone through the process involved myself, I have a pretty good understanding what it's like to be an online startup myself. There are a lot of misconceptions of what the behind the scenes of a website is like.

Many believe that the web is just Facebook, Google, Twitter, and a handful of platforms out there. Therefore, if they know how those sites operate, they know everything there is to know about how the Internet works. That obviously is a really bad way to gauge how the Internet works today despite the popularity of the largest platforms.

Another stereotype often seen on TV is that a group of people in business suits rent out a whole floor in an open concept office plan. They are backed by hundreds of thousands in venture capital and the floor is filled with dual screen Mac computers with people clicking and typing away. A spokesperson is happily pointing out how they have finally been able to launch this amazing website and the future is so very bright for them.

My first hand experience is that neither of these scenarios is even close to the typical norm of what it's like to start a website for most people. For a portion of websites, they were started up because a business decided they just needed a web presence. Often, these sites are thrown together just for the purpose of ticking a box and is left as a side project to an employee who seems to know a few things about this whole Internet thing.

For many of the remainder websites, it's often anywhere between one and three people deciding they have an idea. They research hosting and domain name registration solutions, spend the hundred or so dollars needed, get those accounts running, and start largely learning on the fly. Some have a decent computer science degree or a design diploma while others are just plain learning from scratch. Will the website take off? Will it fail? Who really knows? After all, Google started as a couple of servers sitting in someones garage and now they are a multi-billion dollar giant. You never really know if an idea can take off or not. Some may not even start a website in the hopes of making it big. Rather, it's just a small project for a couple of friends or a group of people. Others are just putting together a site for portfolio purposes.

Put simply, the Internet is huge. The number of sites that exist so often measures in the billions and the number of active sites measures in the hundreds of millions. At best, any one person will have a cursory understanding of what the Internet is like. It is impossible to fully comprehend the full extent of what the Internet is today. You can only really take in a tiny patch of the digital space today.

It is with this in mind that when we talk about regulation on the Internet, there are a couple of fundamental questions we need to ask. This includes:

1. Is this regulation really feasible?
2. Who is this regulation targeting?
3. Will it harm the overall Internet ecosystem either directly or indirectly through unintended consequences?
4. What will the impact be on people?

Trying to implement good regulation on the Internet is notoriously difficult. Between the stakeholders with competing interests and dealing with different political ideologies and even the technical aspects of how technology really works presents an absolute minefield for where legislation can go wrong. Very few regulations have really worked well. Some success has been seen with network neutrality laws and even some privacy laws like those found in Europe. A vast majority of laws that have negatively impacted are often driven by political ideological reasons or by heavy lobbying from specific stakeholders. An example of the former is the US debate surrounding Section 230 reformation and an example of the latter is the numerous copyright reform laws and legislation.

After seeing countless laws being debated over the years around the world, the looming online harms legislation seems destined to fall within the former category. As someone who is hoping to make an impact on the Internet in a positive way, I find it also incredibly troubling despite the problems it hopes to solve being very real problems.

According to the technical paper, the government wants to tackle 5 forms of harmful content. This includes terrorist content or content that actively encourage terrorism, content that might incite violence and hate speech.

For terrorist content, the problem with that is that what constitutes terrorism is always changing. For instance, during the Harper government, there was a push for define some forms of environmental activism as "eco terrorism". Mercifully, that didn't come to full fruition, however, there was motivation to do so. Furthermore, there have been instances where private companies or individuals try to define the lawful conduct of a person or group as terrorism as well. So, this raises the question, "If the definition of 'terrorism' is constantly changing, how do you expect website's of all shapes and

sizes to really keep up?” After all, this legislative push wouldn’t just theoretically be here with the norms of the government of today, but all future governments as well.

For violence or the threat of inciting violence, this is also an extremely loose thing and one I actually witnessed on a website called Techdirt. There was an article about a controversial judge and the patent system. Apparently, that caused someone to comment with the following:

*“Hell, eventually somebody might decide that it’s cheaper to pay a hitman to just cut a brake line or something than go through discovery in that judge’s court. “*

While this is, indeed, a rather salty comment, it’s not necessarily advocating violence. However, it was enough for the US Marshals to demand that the comment be preserved internally along with all information held by the site. The investigation didn’t move forward and a gag order on what happened was released. This alone highlights why policing content advocating violence is not going to be easy by any means.

Hate speech, of course, is not going to be any better. If someone decides to simply offer a comment as an illustrative example, then does that constitute spreading of hate speech? The obvious answer is that it depends on the context. At that point, the question is, where does one draw the line on that?

To return to the four questions I listed above, is this really something that can be pushed feasibly? Even on a well moderated site, we are already, at minimum, on shaky grounds as it is.

Things start to get worse with the second question: who is the regulation targeting? The technical paper defines this with the following in the second and third paragraph:

“The Act should define the term Online Communication Service (OCS) as a service that is accessible to persons in Canada, the primary purpose of which is to enable users of the service to communicate with other users of the service, over the internet. It should exclude services that enable persons to engage only in private communications.

The Act should provide that the Governor in Council may, after consultation with the Digital Safety Commissioner, make regulations (a) excluding a category of services from the definition of OCS; (b) specifying a category of services that is to be included by regulations, notwithstanding that it does not meet the definition of OCS, if the Governor in Council is satisfied that there is a significant risk that harmful content is being communicated on the category of services or that specifying the category of services would further the objectives of this Act; and (c) respecting the meaning of the term private communications for the purposes of the definition of OCS.”

While there have been examples laid out that says that platforms like Facebook, TikTok, and Twitter would qualify, the paper is clear that it is far more broad than this. It’s basically any and every website online that supports comments. It is clear that any website that supports a web forum is under this legislation. Wordpress, which is a CMS used by a huge variety of sites, supports comments as well. The only kind of site I can see not falling into the category of sites that would be regulated might be static web 1.0 websites built entirely out of HTML and CSS. If you make a website that says “hello world”, you probably will be safe. For everyone else? As far as I can tell, you’ll probably be under this regulation sooner or later. This isn’t even getting to the really complex communication methods of utilizing a third party service like Disqus where I wouldn’t even begin to be able to figure out what the site has to do to be compliant with the law.

If there is any doubt about this interpretation, paragraph 6 removes this doubt:

“The Act should ensure that it applies to all regulated Online Communication Services (OCSs), and Online Communication Service Providers (OCSPs) that are the closest legal entity to a regulated OCS, that provide services to peoples in Canada. “

So, to answer the second question, the proposed law as described in the paper appears to be targeting almost everyone who operates a website. This adds to the tenuousness of the feasibility of what is being proposed here.

We next find ourselves moving to our third question: “Will it harm the overall Internet ecosystem either directly or indirectly through unintended consequences?”

This nicely align with what we see next in the technical paper. Paragraph 10 states:

“The Act should provide that an OCSP must take all reasonable measures, which can include the use of automated systems, to identify harmful content that is communicated on its OCS and that is accessible to persons in Canada, and to make that harmful content inaccessible to persons in Canada, as may be prescribed through regulations by the Digital Safety Commissioner, on approval by the Governor in Council.

- a) The Act should provide that an OCSP must take measures to ensure that the implementation and operation of the procedures, practices, rules and systems, including any automated decision making, put in place for the purpose of moderating harmful content that is communicated on its OCS and that is accessible to persons in Canada, do not result in differential treatment of any group based on a prohibited ground of discrimination within the meaning of the [Canadian Human Rights Act](#) and in accordance with regulations. “

“[A] The Act should provide that an OCSP must address all content that is flagged by any person in Canada as harmful content, expeditiously after the content has been flagged.

- a) [B] The Act should provide that for part [A], “expeditiously” is to be defined as twenty-four (24) hours from the content being flagged, or such other period of time as may be prescribed by the Governor in Council through regulations.”

While this is already an extremely high bar, paragraph 11 makes this additional stipulation:

“[A] The Act should provide that an OCSP must address all content that is flagged by any person in Canada as harmful content, expeditiously after the content has been flagged.

- a) [B] The Act should provide that for part [A], “expeditiously” is to be defined as twenty-four (24) hours from the content being flagged, or such other period of time as may be prescribed by the Governor in Council through regulations.”

So, anyone at any time can make a complaint. This already adds an incredible burden on website owners as it is. What’s more is that we see this for paragraph 12:

“[C] The Act should provide that an OCSF must institute internal procedural safeguards providing users of the service in Canada with the following, as may be prescribed through regulations by the Digital Safety Commissioner, with the approval of the Governor in Council:

- a) accessible and easy-to-use flagging mechanisms for harmful content;
- b) notice of the OCSF’s content moderation decision within twenty-four (24) hours of the content being flagged, unless the timeframe is changed by the Governor in Council;
- c) the accessible and easy-to-use opportunity to make representations, and compel an OCSF to promptly review and reconsider its decision; and
- d) notice of the OCSF’s decision upon reconsideration, which must be provided without delay, including a notice of the recourse available to the Digital Recourse Council of Canada. “

I wouldn’t even know where to begin with trying to be in compliance with this. As a result, I find myself wondering if my site has a future under these heavy regulations. This further raises the question, “If someone who actually follows these issues can’t even begin to figure out how to be in compliance, what about the millions of others who don’t even have my level of experience with these issues?”

Where things really start flying off the rails, however, is paragraph 14:

“The Act should provide that an OCSF must generate and provide reports on a scheduled basis to the Digital Safety Commissioner on Canada-specific data about:

- a) the volume and type of harmful content on their OCS;
- b) the volume and type of content that was accessible to persons in Canada in violation of their community guidelines;
- c) the volume and type of content moderated;
- d) resources and personnel allocated to their content moderation activities;
- e) their content moderation procedures, practices, rules, systems and activities, including automated decisions and community guidelines;”

That is combined with paragraph 15:

“The Act should provide that an OCSF must maintain records as necessary for the proper administration of the Act, in accordance with the requirements set out in the Act or prescribed through regulations by the Digital Safety Commissioner, or as otherwise required by law. “

To say that any website can comply with 14. (a) in this paper is extremely unconvincing. What qualifies as harmful content and what doesn’t qualify as harmful content may differ from person to person. The ask is to quantify content that is subjective. As far as I’m concerned, no website in existence today is adequately capable of producing this. In short, the government is asking the practically impossible.

Of course, the harm extends beyond just website operators. Paragraph 14 (c) combined with paragraph 15 suggests that all comments be preserved in the event that the Digital Safety Commissioner comes knocking. As anyone who operates a website in any reasonable amount of time knows, even the government actually does not want that. The simple reason is in one word: spam. Does the government really want the records of, for me personally, approximately 2.3 million spam comments?

I find that highly unlikely. That would do neither side any good unless web administrators want to utilize this as a form of protest.

Further, paragraph 26 suggests that when a website administrator is forced to send information to the RCMP, not to disclose this report:

“The Act should provide that an OCSP must not disclose that it has (a) issued a notification to the RCMP or (b) issued a report to law enforcement and CSIS or disclose the contents of (a) a notification or (b) a report, if the disclosure could prejudice a criminal investigation, whether or not a criminal investigation has begun. “

This opens up the possibility that criminal records are made of people without their knowledge. As awareness is raised about a theoretical law that requires this, it only serves to encourage anonymous communications. For most rational people, if they have a choice between using the TOR network or a VPN service versus unknowingly getting a criminal record, they will choose the TOR network or a VPN service. While I don't know much about how CSIS operates, I'm pretty sure that if they had a choice between a simple communication and peeling open the layers of the Onion network (TOR) for that same message, they would rather choose the former for resource purposes alone.

To answer to the third question, as a result of all of this is, yes, it will harm the Internet ecosystem both directly and indirectly. It also answers the fourth question, “What will the impact be on people?”. The answer is, “substantially bad”.

First of all, everything about this strikes me personally as overly burdensome. For a lot of this, I can't figure out what technical solution would even come close to allowing my website to comply with these regulations. Quite frankly, I can't even begin to fathom a solution that would be capable of complying with something this subjective. I'm only one person. When I see paragraph 119 talk about \$20 - \$25 million fines, I don't even honestly know if it's even possible for me to maintain my website. For anyone who has less technical expertise than me, they probably don't even stand a chance staying in business or keeping up their site. The threat of fines like this will not only deter people of today to continue operating websites, but will also deter people from making new online startups in the future. It is not in the interest of Canada to block the starting of a Canada made tech giant of tomorrow. Further, it is not in the interest of Canada to send a message that Internet innovation is not welcome in this country – which is a message that is made so loudly and clearly in this technical paper.

For smaller players like me, the only viable option I see at this stage is to close up shop. If smaller players can't even have a hope of starting up something, the Canadian government will have effectively banned entrepreneurship not backed by significant sums of money from the outset.

For larger players all the way up to the tech giant's, the scale immediately becomes the problem. Sooner or later, there will be a slip-up. The multi-million dollar fines will immediately bankrupt the medium players easily.

For larger players, a serious question will be asked, “Is it worth it to risk regular fines?” Eventually, the answer will be “no”. It will be cheaper just to geo-block Canada than to comply with regulations this hazy.

What does this leave us? Canada effectively shutting down the entire Canadian Internet. I don't think I even need to explain the devastating economic impact that would have on Canada. It is self-evident. If

I were a visible minority or the subject of online hate for, say, sexual preference, I would find the idea shutting down the whole Internet in my name infuriatingly insulting. It is the equivalent of stopping road rage by destroying every road in the country. Does it solve the road rage problem? Well, you can't have road rage without roads. It's a solution that harms everyone and takes the approach of using a sledge hammer to squash a mosquito.

In conclusion, this whole paper is a terrible idea. It basically envisions that web administrators can wave a magic wand and magically make "harmful content" magically disappear. It would be incredibly misguided to think that regulation will somehow spawn innovation out of thin air in this context. When a reasonable solution isn't available, it simply isn't available. Should this paper become legislation and move forward and become law, we are only going to see mass closures of any business that relies on web infrastructure. Thanks to COVID-19, that is going to be a lot. So, for the sake of me and everyone else hoping to get a business start on the Internet, please do the right thing and toss this whole thing in the trash where it belongs.

Thank you,  
- Drew Wilson  
Founder of Freezenet.ca