

Public Submission regarding the provisions of the proposed
Communications Legislation Amendment (Content Services) Bill 2007

from
David T. Bath

Thankyou for the opportunity to comment on the proposed legislation and associated regulations and processes.

Overview of Submission

1. I applaud the intent of the proposed bill, to impose "*obligations on content providers that supply content services to ensure that they are provided in a manner which is not likely to result in children being exposed to material that would be likely to offend a reasonable adult*".
2. I agree that with the general principle that ephemeral content should have ratings systems and distribution restrictions similar to that now in place for stored internet content, films, television, radio and computer games.
3. Because of the convergence of technologies, and the aggregation of services, I see some risk that innocent services and users may be unduly harmed by some of the proposed control mechanisms, and it may be difficult to assign responsibilities.
4. I feel the model of industry self-regulation of traditional media (TV, radio, film and press) should be applied to internet, streaming and other media, except where technological considerations require different treatment.
5. I believe that other control mechanisms might help achieve the ends of the bill with greater efficacy, and hope that these will be included in the bill and/or accompanying regulations. These could include, among other things: (1) use of domain restrictions similar to the proposed "xxx" top-level domain with strong penalties for adult service providers who do not use the appropriate domain, (2) use of a "blacklist" maintained by an appropriate authority, (3) provision of software and/or firmware for common platforms that can use a "blacklist", (4) provision of guidance for amateur and volunteer service providers.
6. I strongly object to any introduction of a mandatory "whitelist" control mechanism that may be considered as part of associated regulations.
7. I do not believe content-filtering (by automated examination of data streams) is practical.
8. Just as movie-trailers in cinemas have classifications and restrictions, and television programs are preceded by a classification notice, advertisements for adult content services should have a classification system, with prominent display of the classification within the advertisement (for example on TV) during and/or prior to broadcast of the advertisement.

Some of these points will not be addressed in the main body of this document.

Content aggregation

The use of aggregated services (e.g. using feeds) can present **content from a number of sources**.

A good example of innocent use of aggregation is the feed aggregator service from OzPolitics (ozpolitics.info) that draws a summary of selected Australian blogs that concentrate on society and politics. Each of these blogs may have a loose set of contributors, who identify themselves to the blog administrator by the only practical means, an email address. One of these feeds is from the parliament, with the names and a brief outline of the purpose of bills becoming available for public comment. I personally contribute to another blog source included in the OzPolitics aggregator.

The OzPolitics site also provides a guide to the constitution, the Australian political system, and a summary of current opinion polling, making it a very useful resource not only for adults, but also for students of many age groups.

Despite the best efforts of administrators at each of many levels, **it is possible for objectionable content to enter the feed** (such as a graphic image from a battlefield) and be displayed or passed on to other feeds. It may take some **time to alert administrators at different levels**, and an administrator may be unable to act immediately.

The Wikipedia is another example of a useful service, maintained by loose group of volunteers, that can be polluted by malicious individual assuming a number of different identities.

Many companies are starting to provide “wikis” on their own websites, where consumers may comment on products, or give advice to other consumers on product use. Again, a malicious individual can pollute the service.

A relatively new media form, the “mash-up”, combines information from many other services. This is often a map with specific points of interest and annotations supplied by a wide variety of sources, some manual, some automated (such as a current temperature at a site).

This presents a challenge, as different types of providers may be unable to respond to a “take-down” notice within a set time-frame. **Regulators should consider the nature of the content provider** (e.g. commercial, volunteer service, or social), **the value of other innocent content** in the service, and the **technical resources** available to the provider when deciding on reasonable response-times to any requests.

It may also be **difficult to decide** which of the many service providers within the hierarchy of aggregated services is the **appropriate point of contact** to remove objectionable content, **or for prosecution** if violations continue.

The provision of technical assistance by the government, or a private company acting as an agency of government, would be useful to content providers with limited resources **when a notice is issued**, and allow for a quicker response. This could be **seen as a “help-desk”** that can guide content providers, allowing them to satisfy regulator's requests. over the phone, through quick-response email conversations, or through instant messaging. In my opinion, unless the content provider is a business, the agency of government should accept the costs of the help-desk service, and any communication charges (such as through a mobile phone).

Technology convergence

The different media that distribute content are becoming increasingly intertwined. The internet is available on mobile phones. Traditional land-line telephones are integrated with the internet by VoIP. It is possible for a camera in a mobile phone to broadcast to the web.

This is not a new phenomenon. Since the early days of the usenet, when computers dialed each other up, email lists, newsgroups and mirror archives were distributing information across protocols and transport mechanisms. The WAIS (Wide Area Information Service) allowed content to “cross over”, and the HyperText Transport Protocol (HTTP) which is the mainstay of the modern World Wide Web merely built on these other transports and media types, and extended them.

Even then, when most of the people on the usenet were from academia, government and IT companies, well-respected services could be polluted through a combination of malicious individuals and novel ways of integrating information across transport mechanisms, despite our best efforts.

It is likely that even broadcast media (such as TV) might start to include feeds from other sources, and included on-screen like the “ticker” on news bulletins, or like subtitles.

Regulatory mechanisms should be able to address these new media types as they emerge, with different levels of vetting and licensing requirements depending on visibility of the service, the resources of the service provider, and the risk of pollution, while **maintaining an approach that is consistent yet flexible.**

For example, a commercial broadcaster might require a delay and filtering mechanism between receiving information from a feed and transmitting it, similar to the delay and “cut-off” switch on live radio and television to prevent broadcast of inappropriate language or identification of individuals.

As **new ways of providing content can appear “overnight”**, a regulatory system should be able to **respond to innovations in a timely manner**. This consideration re-inforces the need for a generic approach and consistent guidelines that can be adapted quickly to new situations.

Another convergence is the blurring between traditional service providers and individuals who provide a service (whether useful or not) to the public.

Some “services” that are theoretically public are in a relatively private area, with a limited audience (although not a closed group). It may be worth treating relatively untravelled areas of the net like we do people telling dirty jokes or swearing in an obscure corner of a park, while requiring better behaviour from the same people in a shopping centre at lunchtime.

Evolution of technical controls

The efficacy of different technical controls changes over time, and some highly desirable objectives can be technically impossible to implement, or once-workable solutions might become impractical.

It is probably best to illustrate some general principles using “old” problems and solutions that are easier to understand, and make analogies to more modern technologies that would require very detailed discussion. An advantage of this approach is that **we can learn from history about what did and did not work.**

Objectionable email, maillists and newsgroups

Whether through “spamming”, traffic congestion, or noxious content, early usenet services such as email, mailling lists or newsgroups were harmed by malicious individuals or companies.

One of the approaches the community used to control this was sometimes called “letter-bombing”, with everyone who received innappropriate content simply replying (with a very large payload or attachment) to the sender which would “blow up” the offender's mailbox if enough people had been offended, probably fill the offender's disk, and often **took the offending computer off-line**. This was punitive, and the effect had **many similarities to the proposed requirement for a carriage service provider to remove access** to services with prohibited or restricted material.

This gradually became impractical for a number of different reasons, including congestion of legitimate traffic of other users and computers along the same route, or inconvenience to legitimate users sharing the same departmental computer as the offender. This inconvenience might not be minor, and could take a whole department off-line. **Innocent people and organizations were unduly harmed.**

Instead, administrators communicated with each other and made best efforts to remove the offending items, and withdraw login or connection privileges from the offending individual. This might sometimes take a week or two if files had been automatically replicated across the usenet.

At the same time, newsgroups became organized into categories similar to modern domain names (such as “.com”, “.com.au”, “.org”, “.org.au”) except that the most significant part was at the front. Newsgroups starting with “comp.” (computing) were highly reputable, those starting with “alt.” (alternative) were less moderated, while those beginning with “alt.sex.” were considered the **appropriate location for adult material**, and would not be accepted by most computers if they useful fairly simple filtering rules.

While where someone put content in a newsgroup was voluntary and unenforceable (allowing computers without filters to discover new information sources as they were created), this was **remarkably effective at reducing harm**. This approach is **analogous to the “.xxx” domain** (and possibly an “.xxx.au” domain) **proposed for adult content**. I will address this proposed domain in more detail in a later section.

Another approach involved reputable sites providing a **list of addresses** that were considered **harmful** (a “**blacklist**”), or a list of **vetted addresses** (a “**whitelist**”). Email servers would check incoming mail against a blacklist and reject it, or only receive mail from a whitelisted address. In general whitelists were practical only within an organization, but impractical for use in the wider network.

Blacklists were very effective, but required good administration. Innocent people could be put on a blacklist in error. Sometimes the innocent person was reported by a malicious individual. Sometimes an innocent people would lose services because they came from the same department or organization as a malicious content provider. It could take some time to be taken off a blacklist, but **each blacklist provider had a well-advertised process that allowed innocent people to be removed from the blacklist.**

Technical expertise is falling

In the early days of networking, almost everybody who provided a service had enough technical expertise to manage their service and remove objectionable content.

These days, when almost anyone can provide a service, most do not have the expertise to remove content that may have been placed on their service, do not have the resources to examine all the areas within their service and may be unaware of good protective measures.

The regulator should be required to provide guidance for inexperienced service providers from a well-advertised location. This guidance should include

1. An outline of content classification principles
2. An outline of measures the service provider can take to limit the chance of their service being used by others to distribute inappropriate content.
3. A skeleton “terms of use” that small service providers *may* use as a notice when others become loosely co-operating members of that service provider, or for blogs, a skeleton “comment policy” they *may* choose to use.

Ideally the regulator can provide logos similar to standard MA and R notices used on films, games and hardcopy media, to indicate the rating was assigned by the service provider. This would allow the service provider to optionally indicate the nature of the content before a reader gets too deep into the site.

For example:

S-MA : *This site may contain adult concepts, occasional coarse language and occasional violent images.*

This might be appropriate for a service provided by an independent journalist reporting from a conflict zone, that includes comments from other sources.

Similar mechanisms are available to allow other software systems to identify the nature of externally sourced content. The regulator can produce guides designed for small service providers the detail how to make best use of such software systems, and encourage this practice.

Blacklists and whitelists in the modern world

Blacklists can work in the modern world, with the **address lists (and/or directories under a domain) well-maintained by an appropriate agency, and provision of software to parents, schools, and other users**, allowing their computers to lookup (or copy) the blacklist just as virus scanners regularly update a list of viruses.

To ensure that all platforms (computers, phones, etc) can make use of the blacklist, the format of the blacklist, and source code to use it, should be publically available.

The actual contents of the list could be maintained co-operatively with agencies from the governments of other nations, and perhaps with the co-operation of international bodies such as the United Nations and appropriate expert technical organizations (e.g. ICANN, which assigns names and numbers).

The blacklist and software could be distributed using the model of the document management system (known commercially as “MySource Matrix”) that was designed to meet all Australian laws, developed as NOIE (National Office of the Information Economy) was dissolved and functions moved to the Department of Finance as AGIMO (Australian Government Information Management Office).

In this model, the software and add-on modules is made available at no charge to agencies of government. The software and add-on modules have been “whitebranded” (i.e. the same software, but with government logos removed), all available at no charge to non-profit and charitable organizations, while only the core software is freely available to businesses, who must pay for the add-on modules.

I believe that the bill and associated regulations should allow for provision of a blacklist services and software, distributed in binary and source forms, at no charge.

Dynamic numerical addressing (rather than just using the domain names) can present a problem, as these numbers might only be leased for a few minutes, for the length of a single connection. Similar considerations apply to a broadcast from someone with a USB stick at an internet cafe or kiosk.

Another danger, especially where blacklists are maintained with an automated process is the possibility of a service provider (an individual or an organization) is host to a “bot” that takes control of part of their system, sends offensive content to the outside world and is placed on a blacklist. This is commonly used for sending spam and for sharing of files that break copyright. The technology used for “bots” has legitimate use for peer-to-peer networking, but is also used to transfer unclassifiable material.

Whitelists can be very useful within a particular organization that can afford to insulate itself from a large part of the “outside world” (an may be particularly appropriate for kindergartens), but whitelists for general use within a nation are only enforced by governments that overly restrict free speech and access to information. China is the most notable example of a country that enforces whitelists.

I strongly urge that the whitelist approach is optional rather than enforced, and cannot see it ever being practical except for very limited applications.

Proposed “.XXX” domain

A modern version of the “alt.sex.” newsgroup tree is the **proposed “.xxx” domain, which is designed as the appropriate place for adult content**, just as other domains are readily identifiable, eg:

Educational: “.edu” (USA), “.edu.au” (Australia), “.ac.uk” (Britain)

Commercial: “.com” (USA), “.com.au” (Australia), “.co.uk” (Britain)

Details of the proposal, and discussion of it's strengths and weakness is available at the authority that assigns internet domain names and numbers using the following URL:

<http://www.icann.org/tlds/stld-apps-19mar04/stld-public-comments.htm>

A similar philosophy can be applied to content service over mobile phones, with an **easily recognizable prefix to phone numbers** that transmit adult content. Such phone services should always transmit Caller-ID.

These XXX domains are **designed to be monitored** by an appropriate organization, and **contain only adult, not unclassifiable, content**.

Adult service providers would register under this domain (or the equivalent for other technologies), and this domain (and/or country-specific equivalents) is readily blocked by simple technical controls.

It would be fairly easy to create regulations that encourage (or create incentives for) providers of adult content services to register in such a domain, and **treat any strong primarily-pornographic service content that is outside this domain as illegal**. The application of such regulations would be similar to the distinction between registered/legal brothels versus illegal brothels. The threshold should be similar to other R-rated matter.

Service providers should have a mechanism to argue that their service should not be included in the XXX domain, or have previously accepted XXX domain status lifted, by a consultative or arbitration approach in the first instance or a court if necessary. Such appeal mechanisms might avoid valid concerns about free speech that we expect in a modern secular democracy.

Because services can be provided from anywhere in the world (just like phone sex services), a top-level domain (just the “.xxx”) would be useful, not just a national one (“.xxx.au”). This would need **co-operation with other governments** and international technical organizations (such as ICANN).

Such domains, if history is a guide, can avoid many technical difficulties, while allowing registered businesses to carry out legitimate activities, and allow relatively easy prosecution of unregistered services whenever they are recognized.

Subject to privacy provisions, providers of adult content to mobile phones could automatically contact an agency, sending the phone number they want to contact. The agency could provide a negative acknowledgement if (1) the phone was registered to a minor, or (2) the number was given to the agency by an adult that wants to be on the equivalent of a “do-not-call” list. (This is an extension of the blacklist approach, but relates to sending classified material, rather than the receiver blocking it).

I strongly urge the bill and associated regulations to use the XXX domain concept (whether or not the actual letters are “xxx” or not) and equivalent constructs for other technologies as one of the control mechanisms to achieve the bill's aims, and **undertake necessary international negotiations**.

Content-filtering

Content-filtering, using **automated processes to examine a datastream** and guess the nature of the contents has always been problematic, generating many false positives and false negatives.

Filtering email messages that only contain text can be unreliable, and can prevent legitimate information transfer (for example, it is not uncommon for legitimate emails to be incorrectly identified as spam and lost).

Filtering images or sound is even more difficult.

In all media, legitimate content can be difficult to separate from improper content except with the human eye or ear. We'll know when content filtering is practical when we've have completely solved the computer virus, spyware and spam problems which are simpler examples of the same generic task.

Legitimate health-related content might contain explicit terms, and may use terms that a child might know from the schoolground to ensure the child understands it, rather than explain health issues using technical latin or greek terms, just as there is no point talking to five-year olds about cerumen when you want them to understand about ear-wax.

A photograph of a person showing lesions over a body to educate children about the dangers of sunburn or playing with matches is very similar to explicit pornography. It would be very difficult for software to distinguish a still from the "All Creatures Great And Small" TV series about a country vet from something extremely objectionable.

There will always be ways for technically-minded individuals to subvert a content filter.

It is a relatively trivial (if boring) task to get an image past an email filter that strips all attachments and removes any scripts simply by creating a HTML table with lots of small cells, each having a random letter in a small font with the foreground and background colors the same. Different colored cells combine to create the picture. To software (or even most humans examining the raw message), it might look like a very large crossword puzzle or sudoku, while appearing onscreen, it is seen as a relatively detailed image.

Cryptographic approaches (especially steganography) can permit transmission of unclassifiable material that will not be recognized until an innocent-looking image passes through a decryption program. At least such methods require action by the user, and are unlikely to be seen "accidentally" unless the target device has been infected with a bot or virus.

I do not believe that content-filtering by inspecting the datastream is practical, and urge regulators to reject such approaches.