Australia's Stand on Spam

Jeremy Malcolm

Director, ISOC-AU

<jeremy@malcolm.id.au>

ABSTRACT

As we are all too painfully aware, the volume of spam in circulation has now soared to ridiculous levels. Record spam volumes are being recorded every month, with the most recent MessageLabs estimate putting the figure at 76% of incoming email. Over the past two years spam levels have increased by almost four times, and analysts are expecting spam volumes to peak at 90% of all email. Email is fast becoming unusable as a communications medium. Make no mistake, this is a war.

In this war, there are two fronts: a legal front and a technical front. This paper is primarily about the legal fight against spam, and specifically in Australia. However I will begin by saying a few words about the technical front in the war against spam, to offer some background for what follows, and because the two approaches do, as we will see, work hand in hand.

1. Technical measures for spam control

Various technical measures for spam control are either in use or are under being developed, but none of them is or is likely to be a silver bullet. In fact, we can rule some of them out altogether at the outset. These include requiring senders to make micro-payments for email, "challenge-response" mechanisms whereby non-whitelisted senders would be automatically challenged to double-confirm their emails, client-side digital signature validation, and the re-implementation of the SMTP protocol to provide intrinsic sender authentication.

There is nothing wrong with any of these proposals in principle, but the problem of course is that they are next to useless unless they are widely adopted, and none of them is sufficiently transparent, simple, backwards-compatible or cheap for the Internet community at large to be expected to adopt.

There are in my view only four technical anti-spam measures either currently in use or on the horizon that are likely to have a real impact on spam. These are prevention, SPF, blacklisting and filtering.

1.1 Prevention

By prevention, I mean techniques to stop spam from being sent in the first place, and there are a raft of measures that fall into this category, some of which are implementable by end users and some by ISPs. End users can be encouraged to patch their insecure Windows systems and open relays or proxies, and to install antivirus software. ISPs can do the same on their own networks, and can rate-limit suspicious streams of outgoing email coming from their users, and block the SMTP port of their users' own machines. But whilst prevention may be better than cure, it unfortunately isn't enough.

1.2 SPF

SPF, or Sender Policy Framework, is the only new technique that is likely to be widely adopted, given the support it has already received from the Anti-Spam Technical Alliance (whose members include Microsoft, Yahoo, Earthlink and AOL)[1]. SPF was developed last year, and in May this year was merged with Microsoft's very similar Caller-ID email scheme. In essence it is a mechanism to combat the forgery of email domains by spammers. When fully implemented, an mail server that receives an email claiming to be from a certain domain, will look up that domain to retrieve a list of authorised sending IP addresses stored in the domain name record. If the IP address that the email is being sent from doesn't match, it will be rejected. However, in the best case scenario even SPF will only stop spam with a forged From address.

1.3 Blacklisting

DNS-based blacklists are designed to block spam being received from known spammer-owned or compromised servers. Unfortunately

1. They have also recommended a digital content signing scheme, but this is still in early development and does not yet have nearly as much support from the Internet community as SPF. Yahoo's DomainKeys proposal (http://antispam.yahoo.com/ domainkeys) falls into the same category.

spammers also typically use prepaid "throw-away" Internet accounts from major Internet providers to send spam, which cannot be blacklisted without causing many false-positives.

To combat this, ISPs and email service providers have begun to clamp down on the ease with which users can register and obtain disposable dial-up accounts or email addresses. Even so, blacklisting is obviously an incomplete solution.

1.4 Filtering

One can be limitlessly inventive about the ways in which to filter one's email. Regrettably spammers also seem to be limitlessly inventive in the ways they seek to evade these filters. The four main methods of filtering may be classed as heuristic, bayesian, dictionary, and fingerprinting.

Heuristic scans create a profile of an email message from its headers and other core attributes to rate its likelihood of being spam. Bayesian filters use a statistical approach whereby the filtering system is trained to distinguish between spam and legitimate email using an algorithm. Dictionary scans are used to filter against particular words and phrases in the headers or body of an email. Finally, email fingerprinting is used to create a hash uniquely representing known spam messages, which is a reactive rather than a predictive technique.

These all suffer from similar shortcomings, in that they can often be overcome by the spammer adding legitimate text to the message, using random, mis-spelt or punctuated words, or simply including an image in the message in place of text. So once again, filtering is an ineffective solution to the spam problem, and we need to look elsewhere to find a long-term solution.

2. The legislative framework

This is where the law comes in. The Spam Act is the most obvious, but not the only piece of Australian legislation that can be used to combat the spam problem. I will briefly outline each of the major pieces of legislation that bear on spam, before returning to the Spam Act in more detail.

2.1 Spam Act 2003

Australia's Spam Act 2003 was passed in December last year, following a report written for the Government by the National Office for the Information Economy (NOIE) in April. The legislation fully took effect from 11 April 2004. Sending spam (or, in the legislation's terms, unsolicited commercial electronic messages) in breach of the requirements of the legislation now carries penalties of up to $220,000 per day, or up to $1.1 million for further infringements.

There is no specific minimum number of messages that must be sent before they are qualified as spam; a single message can be caught by the legislation. The Act also prohibits the use of address harvesting software or harvested address lists.

The Australian Communications Authority (ACA) has been designated as the enforcement authority for the legislation, which is to be reviewed after two years in order that it can be fine-tuned based on our early experiences of the new regime.

2.2 Trade Practices Act

Australian spammers who send spam that is misleading or deceptive, either in its body or in its headers, may be in breach of section 52 of the Trade Practices Act. Although there are so far no cases in which the Act has been applied against spam or spammers in Australia, the Federal Trade Commission has applied equivalent legislation against United States spammers.

2.3 Corporations Law

The Australian Securities and Investments Commission (ASIC) has taken action under the Corporations Law against a spammer who promoted an unlicensed investment scheme. In that case, the defendant pleaded guilty to interfering with, interrupting or obstructing the lawful use of a computer (for sending the spam through open relay mail servers) contrary to the Crimes Act, and to making or disseminating materially false or misleading statements or information that was likely to induce the purchase of securities, contrary to the Corporations Law. The United States Securities Exchange Commission also obtained judgment against him in the District Court of Colorado for about US$15,800 under broadly equivalent legislation.

2.4 Privacy Act

The intended effect of the Privacy Act in its current form is to preclude spammers from harvesting email addresses without the consent of their owners. The consent required must in general be consent to the use of the address for spam, rather than it being collected for some other purpose.

There are some loopholes to contend with, though. Spam can still be sent to email addresses that were collected for some other purpose if it is impracticable to obtain the

recipient's consent to the use of their address for spam. It will probably always be practicable to seek the recipient's consent by emailing them about it, but that somewhat defeats the purpose of using the Privacy Act to combat unwanted email.

Also, spammers can argue that unless an email address is able to be used to individually identify its owner, it is not personal information at all and so falls outside the scope of the Privacy Act, or alternatively that those who publish their email address publicly, for example on a Web site, are implicitly consenting to its collection. All of these things limit the usefulness of the Privacy Act as a weapon against spam.

2.5 Criminal Code Act

The expoitation of open relay mail servers by spammers is now in breach of the Criminal Code Act 1995 of the Commonwealth. It is likely even to be in breach of that provision if the mail server that is abused is overseas. In simplified terms, liability will attach if the offender institutes or assists in the institution of an unauthorised connection to an open relay server, provided that the connection is either initiated in Australia, or if results of the abuse occur in Australia, or if the offender is an Australian citizen.

Although the legislation is not unambiguous and these provisions have not been tested, it is certainly arguable that the mere sending to an Australian recipient of spam through a compromised open relay server overseas would constitute an offence by the sender, even if the sender had no other connection with Australia. The reverse, namely an Australian spammer sending to overseas addresses through an open relay, is even more likely to be caught by these provisions.

3. Consent, identify, subscribe

Returning to the Spam Act, the core of the Act is encapsulated in a mantra that the National Office for the Information Economy (NOIE) has come up with for the purpose of educating businesses about their obligations under the new legislation: “consent, identify and subscribe”.

“Consent” means that there must be express or inferred consent to the receipt of the message. Consent can be inferred from an existing business relationship with the recipient, such as a customer who has provided their email address to a company on the presumed understanding that the company may send them email. However, there are limits to inferred consent. The fact that you have provided your email address to a company does not allow it to infer that you will accept email about unrelated products or services, or that it can pass your details on to anyone else.

“Identify” simply means that the sender must be identifiable and contactable, and its contact details have to remain accurate for at least 30 days after the message is sent.

“Subscribe” means that the recipient must be given the opportunity to unsubscribe from future communications, even if their consent had previously been given or inferred (although there may be contractual limitations on your ability to opt out of the receipt of some email). Again, the unsubscription details must remain accurate for at least 30 days. The Spam Regulations have clarified that the unsubscription facility must be offered using the same electronic mechanism that the original message was sent with, and without charging any extra fee or using a premium rate service.

4. The scope of exemptions

There are, however, some exemptions to the scope of the Act, and some of these could be seen as potential loopholes. First, there is an exemption in the case of “purely factual information”. This allows vendors, for example, to send out unsolicited newsletters, provided that the newsletter is factual in content rather than being a solicitation to buy. The dividing line, however, is surely very grey. Thankfully, even purely factual messages must still contain details allowing the recipient to unsubscribe, and any unsubcription request does have to be honoured.

Second, government bodies, political parties, religious organisations and charities are exempted altogether from the scope of the Act – even if they are sending a solicitation to buy goods or services. Educational institutions are also exempted, although at least the scope of that exemption is limited to emails sent to their present and former students.

Third, carriage service providers are exempted from the Act for their role in the delivery of spam (unless they actually authorised the spam to be sent, which is a separate offence). This is hardly so controversial, since normally ISPs are as much the victims of spammers as the end recipients, bearing as they do the cost of spam delivery and the brunt of any complaints.

Finally, the Spam Regulations can prescribe certain exemptions from the scope of the Act. To date, the Regulations have been used to exempt faxes from the Act's coverage, since faxes would otherwise fall within the definition of an

"electronic message" and therefore be caught by the Act. SMS and MMS messages sent over the mobile phone network however have not been exempted and are still caught by the legislation.

5. Enforcement

The Act will be enforced by the ACA by giving warnings, seeking enforceable undertakings, issuing infringement notices (which are analogous to speeding fines), or taking court proceedings. These steps will generally be taken in the order listed above, and generally only in response to complaints received from the public. In practice, the only complaints that will be processed are those about spam that originates from Australia, or perhaps from a country with which Australia has an collaborative agreement on spam (although spam sent from anywhere in the world to an Australian, or vice versa, falls within the Act's scope).

In addition to the ACA's own powers, once a spam complaint reaches the Federal Court, the Court has the power to impose penalties of up to $220k for a first offence of a corporate spammer, or $1.1m for continued offences. It can also make ancillary orders for compensation of affected parties, and disgorgement of profits earned. Search and seizure orders can be obtained. The court can also impose injunctions to prevent the spammer from continuing its practices, and unlike in the ordinary case the person seeking the injunction does not need to offer to compensate the spammer for its losses if the court later decides the injunction is unwarranted.

The ACA has recently brought online a report form[2] into which spam may be submitted for further investigation. Spam reports previously submitted to the ACA are currently being compiled with a view to evidence being taken in each State to support whatever action the ACA decides to take against the offenders. In May the ACA enlisted the assistance of the Australian High Tech Crime Centre to assist with these investigations. So far, no prosecutions have yet been made.

6. International measures

The Act recognises that the war on spam cannot be isolated to Australia, and so it empowers the ACA to work both at home and abroad in the coordination of education campaigns, research and liaison with other bodies in the fight against spam. To date this has resulted in Memoranda of Understanding being signed with South Korea, Thailand, the United Kingdom and the United States, by which the respective countries agree to exchange information about anti-spam policies and strategies, and security issues.

Of these four countries, Australia's stand is the strongest. The United States is the only other of the signatories with strong national anti-spam legislation, which for example requires spammers to provide their street address in any communications they send. Yet its Can Spam Act, passed by the US Congress last year, is weaker than Australia's Act, in that it allows spam is allowed to be sent in the first instance so long as an "opt-out" facility is provided.

The International agreements signed to date are only a starting point. Until China and Russia also join the fight, the impact on spam that the above four countries alone can achieve is probably going to be fairly limited. To this end the International Telecommunication Union (ITU) met in July to discuss the spam crisis, and the steps it takes amongst its member nations will be instrumental in bringing rogue spammer states into line.

7. Co-regulatory codes

There is one final weapon on Australia's legal armoury against spam that has not yet been covered – because it is not yet quite in place. This is the development of co-regulatory codes of practice. Co-regulation is a process by which industry-drafted codes of practice are registered with a government authority (in this case the ACA). Once registered the code applies to the entire industry sector in question, so that even those who are not signatories to it can be directed by the authority to comply with it.

Examples of co-regulatory regimes already in place include the various Codes on topics such as billing and customer complaints that the Australian Communications Industry Forum (ACIF) has developed and registered with the ACA, and also the Internet content regulation regime established under the Broadcasting Services Act, in which a code drafted by the Internet Industry Association (IIA) has been registered with the Australian Broadcasting Authority (ABA).

The Telecommunications Act was amended at the same time as the Spam Act to provide for the development of co-regulatory codes on spam and their registration with the Australian Communications Authority. Two such codes are currently under development, and by the time this paper is presented both should have been released for public comment as a prerequisite of their registration.

2. https://www.aca.gov.au/secure/complaint_form.htm

7.1 Australian Direct Marking Association E-Marketers Code

The Australian Direct Marketing Association is developing a code to bind the email and mobile marketing industry. This code will clarify certain aspects of the Act that are left ambiguous by the Act, in part by restating them and giving examples, and in part by adding new provisions to amplify the intent of the Act.

Amongst the expected inclusions are provisions outlining the circumstances in which consent to receive email can be inferred, how that consent should be recorded, the circumstances in which third-party contact lists may be used, standards that marketers should abide by in identifying themselves, and what unsubscription procedures are acceptable.

The code is also expected to raise the bar set by the Act in certain areas such as marketing to children, and viral marketing.

7.2 Internet Industry Association Internet Industry Code

The Internet Industry Association is collaborating with other Internet industry stakeholders in developing a code of practice for which will bind ISPs and Email Service Providers (ESPs).

This code will contain provisions that require ISPs and ESPs to educate their subscribers about their obligations under the Act, and to make spam filters or spam-filtering services available to their subscribers (although not necessarily without charge).

The code will also set minimum standards for ISPs to cooperate with law enforcement authorities in dealing with spam, and require them to take certain technical steps to minimise the potential for abuse of email services by spammers, such as closing its their own open relays, and reserving the power under their Acceptable Use Policies to close or suspend accounts of users who operate open relays.

The code will also specify how ISPs and ESPs are to receive reports and complaints about spam, and how those reports and complaints should be dealt with, depending on whether they relate to the ISP's own conduct, conduct of its subscribers, or simply relate to spam emanating from third parties.

8. Conclusion

Australia has taken a stand on spam, and the rest of the world is following suit. The aim of this paper has been to present an outline of the current state of play on the technical and legal battlefields, with particular reference to the effect of Australia's new Spam Act and Spam Regulations, and the likely effect of the two upcoming Spam Codes.

Although the measures currently under discussion are unlikely ever to result in the elimination of spam, the aim is to return it to a manageable level. When we reach the point that the majority of spam is being filtered out before it reaches the end user, it seems plausible that the return on investment for spammers will reduce to an extent that will discourage them from continuing the practice.

In the shorter term, Australia's tough stand on spam will also ensure that Australia is not regarded as a safe haven for spammers, so that international pressure will be able to be focused on those countries that are.