FEATURE 2

HAMFIGHTING – HOW ACCEPTABLE ARE FALSE POSITIVES?

David Harley
Small Blue-Green World, UK

There has been ongoing controversy over the last couple of years about the (allegedly) aggressive nature of *Verizon*'s anti-spam strategy. Complaints in various forums of poor email delivery service from the ISP seemed to be confirmed by claims from *Verizon* 'insiders' that a policy of rejecting mail by IP block resulted in the loss of all mail from large portions of Europe and Asia. This led to a much publicized class action, resulting in a settlement offer from *Verizon* to compensate customers who lost legitimate mail between October 2004 and May 2005.

But let's take a step back from the specifics of the *Verizon* case. After all, the settlement offer leaves the exact details of *Verizon*'s policy and actions unresolved and unconfirmed, and we can only speculate on the accuracy of the whistleblower comments posted in various forums and quoted widely elsewhere [1]. In any case, there are no obvious indications that *Verizon* was actually in breach of its terms of service, however aggressive its policies. The affair does, though, call into question some widely held assumptions.

100% SPAM BLOCKING

Can an anti-spam service stop all spam being delivered? Probably not, if only because there will always be a grey area where one man's spam is another man's ham. And, of course, some people have not yet developed sufficiently sensitive taste buds to distinguish between the two. Without getting tied up and bogged down in exact definitions, I guess that most of us would be happy to lose the unholy mixture of nuisances that assault our filters – kiddie porn, 419s, cheap c1al15 and v1ag4, phishes, pennystox bulletins, job opportunities in the burgeoning money-laundering market, OEM versions of our own software, tsunami victim hoaxes, religious epiphanies and all. However, there are two prevailing views of spam management:

- It isn't possible to stop all spam from being delivered.
- It may be possible to stop all spam, but only if you accept that some legitimate mail may be lost. A *Verizon* statement puts this view even more starkly: 'Any spam-blocking method will, inevitably, result in the blocking or delay of legitimate email.' [2]

In fact, these are not mutually exclusive philosophies. They're two points on a spam management continuum between allowing all mail and allowing none. Few businesses go so far as to reject all external email, though some public sector departments are absurdly bashful, and go to extreme lengths not to publicize their email addresses or web pages. One might almost suspect a conscious rate-limiting approach to workflow management.

However, it's not unusual for corporate administrators to refuse all mail from certain geographical regions or domains because of high volumes of spam, viruses, phishing mail, and so on. They may even reject email from any source not currently known to them, though in that case there's usually a mechanism by which outsiders can apply to be included on the corporate whitelist. They may add offending addresses to an in-house blacklist, or they may use one of the many open DNS blacklists. All these approaches have their downsides, in that they can entail a risk of losing mail traffic which may impact on business processes.

Institutions and individuals, however, may be able to afford these negatives (unless, of course, they are simply unaware of them). An enterprise can take decisions to block huge swathes of IP block, or all unsolicited mail, or all mail that doesn't come from a whitelisted source, on the basis of an informed risk analysis process. They may decide that the risk that they'll lose customer enquiries, mail relating to ongoing transactions, enquiries relating to freedom of information or data protection and so on, is acceptable given the nature of their spam problem.

However, that raises a question: is it acceptable for an ISP – or, indirectly, a third party such as a Domain Blacklist (DBL) maintainer – to impose that risk upon them without consultation? Is it acceptable to do so in the case of an individual rather than an enterprise?

OPEN AND SHUT CASES

Open blacklists are a mixed bag: they range from highly professional services, scrupulously maintained, to zealots who are happy enough to blacklist the whole of Germany in the hope of applying enough pressure to persuade a recalcitrant domain to conform to their particular ideal.

Where a blacklist is maintained responsively and responsibly, so that the risk of false positives is minimized by prompt weeding, a reasonable balance is maintained. That balance lies between the need to apply sanctions against those who (through malice, greed or simple ignorance) are abusing Internet mail mechanisms, and a wish to spare those whose only crime is to share a gateway or other IP space with an abuser, from being punished for the sins of others.

Many blacklists are maintained by volunteers with a sincere conviction that everyone will benefit in the long term from

the punishment of transgressors. If they think about it at all, they may assume that incidental damage to innocent parties is a regrettable, but acceptable risk. They may even count consciously on such parties to apply referred upward pressure to the service providers and administrators who cling to discredited practices. Certainly, many of us will sympathise with the discouragement of unsecured relays, indiscriminate distribution of misdirected virus notifications, poor Non-Delivery Report practice, and so on.

Essentially, most blacklists prioritise what are perceived as the interests of the majority over the convenience of the individual. This Utilitarian philosophy of the greatest good for the greatest number is in some senses laudable – indeed, you could argue that the societies that most of us live in are to some degree founded on it – but it sits uneasily in the commercial context of a mutually agreed contract.

It's probably no coincidence that those lists that come closest to striking an acceptable balance between the blocking of bad mail and the free passage of good mail are those that have at least one foot in the commercial spam-management sector. Volunteer blacklists are not always updated promptly, or responsive to pleas from those who suffer collateral damage that they aren't responsible for someone else's mismanagement. They can usually fall back on the argument that 'We don't block anyone, we just publish a list'. Those that supply a contracted service, however, have a commercial interest in maintaining good and responsible relations with customers *and* potential customers, and certainly can't afford to keep failing to meet service level agreements.

INDUSTRY PRACTICE OR BEST PRACTICE?

It's not unknown for ISPs to risk blocking some legitimate mail, in the hope of reducing spam received by their customers to zero, or as near to zero as possible. Verizon, for example, has been quoted as saying that it 'block[s] narrowly' using 'methods that are consistent with industry practices'.

But is it *best* practice? Many organizations cannot accept the risk of life- or business-critical services being disrupted by false positives, and the better anti-spam services generally try to accommodate that need by taking strenuous measures to try to ensure that no legitimate mail is lost.

Specialist reputation services try to avoid blocking mail from non-spammers who happen to share IP space with spammers, and they apply fallback mechanisms such as quarantining suspect mail. This not only allows the customer some means of monitoring performance, but lets them retrieve mail that has been incorrectly classified as spam.

Home users don't usually do that sort of risk assessment, and often have unrealistic expectations that anti-spam services will not only block all their spam, but also give them access to all their legitimate mail – two expectations that are not necessarily compatible.

However, the reported reactions [3] of *Verizon* customers to this class action and the debate around it suggest that some home users need and expect reliable mail delivery. Some of them use their connection for business, and need reliable mail delivery just as much as *Fortune 100* companies do. Furthermore, it seems that even some recreational users, if forced to think about it, prefer to receive some spam rather than risk not receiving mail from family or friends.

CONCLUSION

Home and small business users are unlikely to demand the same guaranteed levels of service delivery that are built into major corporate contracts. In fact, they just want the provider to take care of it so that they don't have to think about it. They may not bother to read their provider's Terms of Service. But they do consider their ISP accountable for the safe delivery of legitimate mail, even for a basic service.

Smaller customers are starting to realize that they may need more flexibility (even if they have to pay extra for guaranteed delivery) and to know more about how the service they're receiving works.

If ISPs want to maintain a one-size-fits-all spam-filtering service, they need, as a minimum, to make clear what users can expect of the basic service, what optional extras are available, and what your money gets you in each case. Expectation management is key: it's no longer enough to say 'we block spam; how we do it isn't important'. In order to avoid legal action and maintain market share, ISPs need to realize that spam management is a balancing act. It's also an exercise in PR.

If ISPs really want to maintain a draconian level of spam filtering, they may want to consider ensuring that they whitelist organs like *The Register* with a reputation for voicing the concerns of the end-user.

REFERENCES

[1] http://www.pcworld.com/resource/article/0,aid,119717,00.asp.

[2] http://www.spamhelp.co.uk/2005/01/verizon-spam-filtering-statement.html.

[3] http://www.theregister.co.uk/2005/01/21/verizon_class_action/.