

Fighting the Battle Against Spam

Find out what you can do and what WebFarm are doing to help reduce the amount of Spam you receive...

Unless you are one of a lucky few Internet users, you will no doubt be inundated every day with massive amounts of UBE (unsolicited bulk email), or UCE (unsolicited commercial email), otherwise known as Spam. Spam is an ever increasing nuisance (estimated at over 90% of all email traffic), which can often be malicious in nature, that is choking mail servers and networks throughout the Internet, not to mention causing a loss of productivity in every workplace that uses electronic communications. What the Government are doingSeveral countries around the world have outlawed spam, and now impose penalties against anybody found to have been a source of Spam. The New Zealand Government has recently brought in the "Unsolicited Electronic Messages Act 2007", which is legislation that prohibits any person from sending unsolicited commercial email, or using address-harvesting software or a harvested-address list in connection with the sending of unsolicited commercial electronic messages. This is a step in the right direction, but unfortunately a very small percentage of the Spam we receive originates in New Zealand. What the Internet Community as a whole is doingMost ISPs hate Spam, and work together to fight against "spam-friendly" providers who allow users to send out Spam from their networks without consequence. Most ISPs use several forms of "black-list", which basically is a method of identifying spam-friendly providers and rejecting any mail received from them. What you can do as an Email UserThe obvious way to reduce the amount of Spam you receive is to make sure that spammers don't have your email address! Spammers get hold of email addresses in five main ways:

- They pick them up when they're used publicly on the Internet, e.g. in a newsgroup posting or on a webpage. This is by far the most common way, and is known as "harvesting".
 - They buy a CD or List of email addresses. These addresses were most likely harvested in the same way described above, and can often be years out-of-date.
 - They guess them. For example, it's a fair bet that "bob@example.com" could be a valid email address, although there's no way of knowing for sure. But it takes so little effort and cost to send Spam, they just try sometimes several hundred possibilities @example.com.
 - Some unscrupulous ISPs (and other organisations) sell spammers our email addresses, although this is quite rare.
 - We give them away. Always carefully read the privacy policy of any website before you give up your email address, as sometimes email addresses are passed on or used for purposes other than those we intended. Below are some simple things you can do to reduce the amount of Spam you receive: Choose a non-obvious Email Address As spammers can guess email addresses, use something that can't be easily guessed. For example, instead of bob@example.com, use bob34z@example.com. Be careful with your Email Address Never give your email address to a company you do not trust entirely. If in doubt, open a free email account with a provider like hotmail.com or yahoo.com, and use that address for communicating with the less trusted organisations. That way, if they do spam, you can close the account and you've only lost a free email account you weren't using for anything else.
- Don't post to usenet using an unmunged email address you care about. Use a throw-away address from a free email provider or munge your email address as described below. Never allow your email address to appear on a website in its true form, including on web-based discussion boards. Instead use Contact Forms or munge your address. Address Munging "Munging" is the act of mangling your email address so that it can still be read by a human but cannot be automatically harvested by spammers. For example, the email address bob@bobscompany.co.nz could be munged into any of the following: bob<at>bobscompany<dot>co<dot>nzbob@bobscompany.co.nz.REMOVETHISTOSENDEMAILbob@REMOVE-CAPS-AND–INVALID.bobscompany.co.nz.invalidbob@NOSPAM.bobscompany.co.nz.NOSPAM Whitelisting You set up your mail account such that some given word or string of characters must be in the subject line for any mail to be accepted, and then you explain this in any newsgroup postings and webpages containing your address. This way people can respond to you, but spam will be deleted from the server without you having to spend time downloading and reading it. This works especially well with webpages, e.g. use:

<AHREF="mailto: bob@bobscompany.co.nz?Subject=FRIENDLYMAIL: Comments about my webpage">Send me email!Then just "auto-delete" any mail that doesn't have FRIENDLYMAIL: in the subject line and download the rest. Challenge-Response Tools Challenge-response systems, also known as "Reverse Whitelisting" or "Permission-based" filtering, take a different approach to traditional spam-filters. Whereas traditional filters start from a stand-point that all mail is good then try to detect the spam, Challenge-Response systems start by assuming all mail is spam then only letting through people on a "whitelist". If the user receives mail from someone not on a whitelist, the system "holds up" the mail and sends a "challenge" message to the sender. If sender replies ("responds") to the "challenge" message, the original message is "released" and allowed into the user's mailbox, and the sender is "whitelisted" so any future emails will be allowed through without this rigmarole. The theory here is that the spammers won't bother to reply to the "challenge" - most of them are using forged email addresses so they won't even receive the "challenge". Put like that, it sounds like quite a good idea. But the simplicity of the solution doesn't reflect the complexity of the real world, and challenge-response has a number of problems:

- Mailing lists, especially discussion lists. If a message is sent to a mailing list with 1000 subscribers, would you receive - and have to respond to - 1000 challenge messages? Many Challenge-Response systems allow the user to whitelist a mailing list automatically, but this can be unreliable.
- Automated mailings - generated by a computer with no human intervention - have no human sender who can respond to the challenge message. This immediately breaks things like password reminder messages, confirmed opt-in mailing

lists, Cron job notifications and so forth. Again, these things could be whitelisted manually - but you would have to guess the email addresses they will be sent from, which would be difficult.

- Forged sender addresses. Spammers often forge the addresses of enemies or just random individuals as the senders of their spam - if a spammer forges the sender address of a 1,000,000 spam emails, the poor person whose email address was used could receive a "challenge" message from each and every victim!

Using a tool to send fake "bounce messages" When you send an email message to an address that doesn't exist, you receive a "bounce message" back. There's a school of thought that says that if you send fake "bounce messages" in response to the spam you receive, spammers will remove you from their mailing lists and you'll get less spam in the future. To this end, there are various tools - the most well-known being MailWasher - that will generate such "fake" bounce messages.

The general consensus is that this is a bad idea. Here's a few reasons why:

- There is lots of anecdotal evidence that suggests spammers as a rule are not interested in removing dead email addresses from their lists.
- The return address in almost all spam messages these days is forged, probably because the spammer knows his mailing list has lots of bad addresses and he doesn't want the bounce messages to fill up his own mailbox. So any "fake bounce" you generate probably won't reach the spammer anyway. So at best, your "fake bounce" will hop around between mailservers consuming resources before being quietly dropped. However, a lot of spammers forge their spam to look like it came from the email address of a real person - either someone who's annoyed them (e.g. an anti-spammer) or just some poor soul picked at random. So your fake bounce message - together with those of everyone else that uses such a tool - would end up in the mailbox of this entirely innocent third-party.
- By examination of the headers and included information in a bounce message, it's possible to make a reasonable inference as to whether it is real or fake. So even if your bounce message did somehow reach the spammer, his systems may well be able to figure out that it's fake and ignore it anyway.

What WebFarm are Doing… To Stop Spammers using our Mail Servers to send Spam It can be very difficult to effectively stop any Spam being sent out from our servers, however these are the things that we find make a difference:

- We have built up over the years a reputation for dealing with spammers quickly and decisively, so many spammers avoid us.
- We have an Acceptable Use Policy which contains a clause stating that if you are found to be Spamming, we will immediately and without warning disable your domain and email service.
- We have implemented a mail system on all of our servers which requires each user to log in and check for incoming mail before mail can be sent out – this effectively means that a Spammer cannot use our mail servers as an open relay to send out Spam.
- We monitor the email traffic generated by our customers. If a customer who hadn't previously sent more than three or four emails a day suddenly sends a hundred thousand messages for example, this would increase the load on our server and be picked up by our monitoring tools and dealt with swiftly.
- We monitor orders we receive for Web Hosting and Email Accounts, looking for suspicious looking fake addresses, and dodgy sounding Domain Names, and we don't even give these people the opportunity to get access to our mail servers – we refund their payment and send them on their way.

To Reduce the Amount of Spam our Customers Receive Trying to reduce the Spam received by our customers whilst still accepting all legitimate email is very difficult, as if we are too harsh on Spam we run the risk of trashing

“good” mail that our customers need to receive. We have adopted the following Spam fighting policies and techniques... Spam Filtering One tactic we use to cut down on Spam is filtering. Our Mail Transfer Agents (MTAs) receive all incoming mail and scan each individual email, tagging any mail that matches the pattern of a known piece of spam. Mail is graded on a points scale, where certain “spammy” characteristics score varying points. When the points accumulated by an email go over the Spam threshold score, the mail is classed as being Spam and is tagged as Spam before being delivered to the recipients mailbox. The recipient just needs to set up filtering in their email program to put the tagged spam into a separate folder, or alternatively they can choose in their Spam Filter settings to automatically delete any mail which is tagged as being Spam. Bayesian Filtering Bayesian Probability Filtering is a spam-filtering technique which has been integrated into our mail scanning tools. The idea is that you "train" the filter to recognise spam from non-spam, by telling it whenever it makes a mistake. This has been quite successful because everyone's spam is different and the types of legitimate mail everyone gets is different. Spamhaus SBL The Spamhaus Block List is a realtime database of IP addresses of verified spam sources and spam operations (including spammers, spam gangs and spam support services), maintained by the Spamhaus Project team to help email administrators better manage incoming email streams.

Our mail servers at WebFarm use this list to identify and block incoming mail from IP Addresses which Spamhaus has deemed to be involved in the sending of Unsolicited Bulk Email (Spam). Realtime Block/Blackhole Lists (RBLs) We use a variety of these Block/Blackhole Lists, which contain IP Addresses of known Spam sources, to “blacklist” mail from machines that have a reputation for sending a disproportionate amount of spam. Greylisting The “greylisting” technique involves deferring an incoming email (responding to the senders mail server advising that the mail can't be delivered at that time), and then accepting the second attempt to send the email. The reason why this is effective against Spam is because most spamming mail servers don't wait to see if the recipients mail server is going to accept the mail or not, they just bombard it with mail. If it's a legitimate email sent from a legitimate mail server, the sending mail server will try once, then try again a short time later when the first attempt is not successful. Once an email has been determined to be from a “good” source, the sending

mail server's IP Address is 'whitelisted' for 7 days – this means that all mail from that mail server will be received without being deferred first. Directory Harvesting Attack (DHA)Our mail servers are set up to detect DHA attacks and block these appropriately. DHA attacks involve a spammer (or a virus) trying to send to multiple email addresses at the same Domain Name e.g. bob@example.com, jane@example.com, info@example.com, sales@example.com. If our mail servers detect that these 'recipient failed' errors are occurring too often from the same sending mail server, we will defer all mail from that mail server until an hour has passed where no 'recipient failed' errors have occurred. Sender VerifyOur mail servers will not accept mail that does not come from a verified 'From' address. Basically, what happens when an email is coming in, our mail server responds to the sender's mail server and asks it 'does the 'from' address in this email actually exist?'. If the sending mail server replies 'yes' then we accept the mail, but if it cannot verify the sender's address, the mail is then deferred and will continue to be deferred until such time as we are able to verify the sender's address.