

Network-Centric Spam Control Solutions

(version 3)

Internet Research Group
July 2005

About The Internet Research Group

www.irg-intl.com

The Internet Research Group (IRG) provides market research and market strategy services to product and service vendors. IRG services combine the formidable and unique experience and perspective of the two principals: John Katsaros and Peter Christy, each an experienced industry veteran. The overarching mission of IRG is to help clients make faster and better decisions about product strategy, market entry, and market development. Katsaros and Christy recently published a book on high tech business strategy *Getting It Right the First Time* – Praeger, 2005 www.gettingitrightthefirsttime.com.

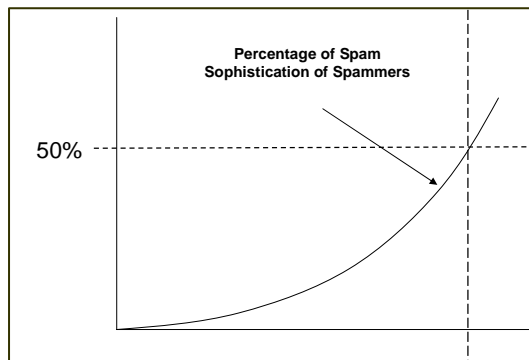
Table of Contents

1.	Network-Centric Spam Control Solutions	1
2.	So What? Now What?.....	2
3.	The Network-Centric Solution Framework.....	3
4.	Defining spam	4
5.	Authentication Technologies.....	4
6.	Reputation Authority Requirements.....	7
7.	Changing the spammer’s Cost	7
8.	Reputation Monitoring and Reporting	9
9.	Putting the Pieces Together.....	10
10.	Summary.....	12



1. Network-Centric Spam Control Solutions

In 2003 spam became an un-ignorable problem simply because the percentage of mail traffic that is spam grew beyond 50% of the total Internet traffic. Had spam stabilized at (say) 20% of mail, things would be very different. When spam percentages grew beyond 50% we saw a classic “tipping point” occur. Up to that moment, by-and-large spam was a mail subscriber annoyance. When the volume went beyond 50% all of a sudden the volume of spam dominated provisioning and operation of email as a whole – bandwidth, servers and storage now all had to be scaled because of the spam component.



Breaking through 50% served as a clarion call-to-action for the email community, and resulted, for the first time, in real community action. There was a collective realization that spam

control required that everyone do their part and also that there were interesting and powerful collective solutions.

Since 2003 the volumes of spam continued to grow reaching in excess of 80% of all Internet traffic by the beginning of 2005. The rapid growth of broadband connections and the use of virus distribution by SPAM created large collections of “Zombie” machines on broadband networks sitting ready to serve the needs of the spammer (creating large volumes of spam mail from previously unknown systems).

Despite the growing sophistication of the spammer and the growing volume of spam messages since 2003 we have made a great degree of progress minimizing the impact of spam on mail systems and mail users. The use of network-centric filtering and blocking methods was key to this turnaround.

When spam volumes exceeded 50% spam was unavoidably a “denial of service” problem. Processing and storing spam messages denied the use of mail server resources for legitimate mail processing. Volumes of spam in an end-user mailbox denied the use that storage for legitimate mail. Before the spam crisis, clever methods had been developed to deal with other forms of network based denial of service

attacks. Adaptations of these ideas and similar new ones turned out to be the most important single weapon in the war against spam. As is often the case in science, finding the right perspective in which to consider a problem leads more directly to the solution, and that certainly seems to be the case here. Our report focuses precisely on network-centric solutions to spam.

2. **So What? Now What?**

Spam control drove real spending, and as such became a focus for venture investment and entrepreneurial product development. The importance of IP reputation was a poorly understood aspect early that eventually had a large impact on product effectiveness.

- Spam control becomes more effective and much simpler. The value of current leading edge technology spam filtering solutions diminishes rapidly.
- Spam control is based much more on network actions and much less on content analysis magic. The window for bringing exotic content analysis technology to market at high value closes.
- The spam solution moves more into the network changing the value of being a network or IT ecology vendor.

Network-centric solutions diminish the amount of mail filtered (the filtering load) by almost a factor of 20, while reducing the amount of spam left in the in-box by an equivalent factor.



3. The Network-Centric Solution Framework

The majority of the product development effort during the spam control boom was based on the clever application of mathematics, heuristics and statistics – developing innovative and adaptive programs that looked at the subtle characteristics of millions of legitimate and spam messages, and building an effective filter that correctly categorizes the next message. For a while this approach seemed like the best way to get the better of the spammers. Unfortunately, the spam business grew the volume of spam grew and so did the number of spammers, as did their sophistication and capability. The war against spammers looked more and more like an arms race, with the victory becoming less and less certain.¹

Fortunately, as spam became a big and inescapable problem, the community view of spam (legitimate mail senders and receivers have a shared interest in defeating spam) and network view of spam (we know about network denial of service already) suggested alternative solutions that don't suffer from the same arms race problems. We can now see the framework of how spam can and will be dealt with. There are three pillars:

- **Address Authentication:** The existing email exchange protocol (SMTP) was designed in a kinder and gentler day and age in which the idea that people would maliciously misuse the various mail message headers was clearly not considered. As a result, the fact that the headers say that mail comes from amazon.com means nothing for sure. Fortunately, authenticating key elements of a mail message doesn't require changing the basic protocols. In the relatively short term, it will be easy for mail receivers to authenticate that mail comes from where it claims. Knowing where mail originates, and then being able to query the reputation of that mail source makes a dramatic difference in spam control.
- **Reputation and Community Concern:** An important part of spam control is self policing. If reputable mail senders work to assure that their public reputation is being spam-free, then mail recipients can receive mail from those sources without concern about spam and without using complex filters with a risk of falsely blocking legitimate mail. Reputation-based mail processing depends on having convenient sources of reputation (authorities) and on the desire of mail senders to maintain a good reputation.

¹ Using the Internet and Web for effective real-time collaboration, spammers exchange tools and provide services that consolidate much of the spam community's knowledge about spam-blocking tools and services, making the spam/anti-spam war one that evolves remarkably rapidly.

- **Reputation Management:** Mail reputation is like consumer credit rating. To be viewed as credit worthy, it is necessary to have a real credit history and for that history to be reputable. Mail reputation requires a history of mail use and that history to be reputable. Tracking and characterizing mail sender behavior and identifying and flagging anomalous behavior is an important part of creating a mail reputation mechanism.

4. Defining spam

Ultimately spam is undesired email and as such only definable on an individual-by-individual basis. Fortunately, in order to fix today's spam problem, we don't have to worry about that degree of philosophical clarity. Instead we can create a simple and objective definition of spam and it will suffice.

Today's huge spam volumes can only be generated because

- Spammers make irresponsible use of email addresses (e.g. directory harvest attacks) and send mail to individuals who have never expressed any interest – direct or indirect – in that mail.
- Spammers blatantly ignore requests to be removed from lists.

For the sake of this discussion, we will define a spammer as someone who sends volumes of email based on irresponsible use of lists (lists for which the genealogy is unknown) and/or does not reliably respond to requests to be removed from a list. If we can stop these forms of spam we can again make spam a manageable problem².

5. Authentication Technologies

Our primary purpose here is to explain network-based spam control means and show how these are far more important than the exotic filtering ideas that drove much of the spam boom product investment. There are other important network-centric technologies worth mentioning. This section summarizes current efforts to improve

² There are lots of ways in which people will argue about what is and is not spam. There should be little debate about whether our definition is reasonable. The debate won't be whether people who violate these rules are spammers; it will be about (1) whether this definition is powerful enough to be useful and (2) whether there are reasonable ways of translating this definition into a formal spam reputation. Both issues will be addressed later.

In the discussions we have had about this work, some people have suggested that these schemes work even if we weaken the definition of a spammer to someone who does not respect opt-out requests.



message authentication.

Underlying much of the spam problem is the fact that the world of email today looks nothing like the world of email when the SMTP protocol was designed 21 years ago. At that time, a mail server was big and expensive to buy as were network connections. In the beginning, mail servers were costly to operate and rare. The creators of the protocol standards that define today's SMTP mail systems could not have reasonably anticipated a day when a \$500 mail order computer was thousands of times more powerful than those servers, when a broadband Internet connection was easily affordable by many teenagers out of their allowance and when nefarious organizations had created and could one demand exploit millions of Zombieized broadband computers.

The SMTP design assumed implicitly that mail server operators were generally known and responsible members of the community and could be trusted, a bad assumption today. As a result, SMTP was designed in what seems today a very naïve manner and is a very easy protocol to spoof. Any "script kiddy" can create mail that appears to come from an arbitrary sender, while sending the mail from an unknown and untraceable IP address. Spammers are bright and agile opponents; and today they are well-armed as well.

An academic would probably conclude that we must fix SMTP by redesigning it from the bottom up for today's use, and in fact it would be relatively straightforward to do so. However, if you consider the billions (literally) of computers with software based on the existing SMTP protocol, making any wholesale replacement and upgrade isn't feasible. Practically we need to find ways of making SMTP more secure without throwing it out.

Recognizing the problems and limitations of practical solutions, mail technologists have developed schemes to use side channel information (information not transmitted by SMTP) to validate elements of SMTP. Any mail sender needs to use DNS (Directory Name Services) services to convert a destination mail address (e.g. amazon.com) into an IP number for a mail receiver for that destination. The same DNS services can be used in reverse, to retrieve additional information about mail sending domains that can be used to validate the authenticity of mail header information.

In general it's not a question of these different technologies competing with one another but rather each providing complementary capabilities. IP-based reputation is a generally useful means of gaining useful contextual information. Two other important initiatives are sender authentication and message signing.



Sender authentication schemes (e.g., SenderID³ / SPF⁴) use the DNS database to distribute information that helps defeat email fraud. In the case of email abuse control we can extend the same scheme to distribute the IP numbers of mail servers that are authorized to send mail in the name of some domain (e.g. the mail servers that intel.com uses to send email). When mail is received that purports to be from intel.com (that's what the mail headers say) then we can check to see if the server sending it is authorized by Intel to do that.

There are many ways you can choose to look at spam control. From a network architecture perspective you would like to maintain the separation of duties of the various network layers in which case it makes sense to separate the authentication and reputation of mail senders from that of mail processors (the server systems that relay and transfer mail along the way – “MTA’s”; see the Certified Server Validation proposal -- CSV⁵).

Message signing (e.g. DKIM⁶) provides an alternative means of validating the authenticity of a mail message. Again it uses DNS to distribute information about a mail sender. In this case it distributes a public key that enables the validation of a digital signature that is embedded within a mail message. If the mail claims to be from yahoo.com and has an embedded digital signature then it should be the one that only yahoo.com can create. The use of public key encryption makes it impossible to forge Yahoo's signature and yet be able to identify the signature as authentic.

Like IP-based reputation, sender authentication and message signing are relatively easy to add into existing email systems because they don't require any changes to the mail protocols proper. Information is distributed via DNS (organizations directly control their own DNS entries). Implementing these schemes requires that the new information be published by the sender and subsequently checked by receiving mail servers. A remarkably large percentage of Internet mail is sent from a small number of addresses (e.g. aol.com, yahoo.com, msn.com, hotmail.com ...) and the same is true for mail destinations. Thus it's possible to make significant progress with a very small amount of work by getting a new scheme implemented by the community of high volume mail senders and receivers, an effort that is tiny compared to a scheme that would require protocol changes for all the systems sending or receiving email.

This is not to say that the email protocols can't and shouldn't change. Efforts to make fundamental improvements to email protocols are well

³ <http://www.microsoft.com/mscorp/safety/technologies/senderid/default.aspx>

⁴ <http://spf.pobox.com/>

⁵ <http://www.mipassoc.org/csv/>

⁶ <http://mipassoc.org/mass/specs/draft-allman-dkim-base-00-10dc.html>



underway; they'll just take a significantly long time to have practical impact

6. Reputation Authority Requirements

Reputation authorities are independent, arms-length operations. All of the large mail service providers could also provide these services but there are fundamental conflicts of interest given that they also sell diverse marketing and eCommerce services as well.

Initially, a reputation authority will succeed if it is able to do the following:

- Accurately reflect the mail behavior and spam characterization of mail senders (see enough traffic).
- Quickly inform mailers of potentially disruptive changes in their observed behavior.
- Respond quickly and reasonably to disputes about reputation.

Over time we believe that reputation authorities will differentiate additionally in these ways:

- The timescale in which the authority responds to changes in behavior both to the better and to the worse (the time to resolve a grey area rating; the time it takes to identify significant behavior changes).
- The ability to differentiate reputation (consumer categories, differences in law and culture). What is considered to be spam will ultimately depend on many factors including the sensitivity of the recipient (consumer categories), and the legal and cultural environment (law and culture).
- The ability to differentially identify and service valuable sub segments of mail traffic (e.g., Law Firms, ...).
- The regions of the world-wide Internet that are adequately monitored and characterized (e.g. something more than a service for North American mail and mailers).

7. Changing the spammer's Cost

Spam economics are simple. It costs almost nothing to connect a



computer to the Internet and send mail. There is no formal registration of an email sender unless you want people to be able to send mail back to your address. Since spam is one way (typically giving a URL that directs the respondent to a Web site), and because SMTP doesn't assure that the mail sender listed in the message header is valid, a spammer can just turn on the computer and start sending mail without any worry about whether or not any of the mail header addresses are legitimate or work.

Generating volumes of spam from a fixed IP address will cause that IP number to be blacklisted in reasonably short order, but if the spammer is on a cable network or DSL line that assigns IP numbers dynamically, there is no long-lived IP address to worry about. Poorly managed mail servers that are configured to relay mail on demand are another source of delight to a spammer, since mail sent through an open relay comes from a known mail server.

Community action and a little technology change the spammer's economics radically. SMTP mail addresses and the IP address from which the mail originates can be authenticated with relatively straightforward schemes (some version of which is in use in many large mail processors already). Mail server operators have learned to configure their servers properly or risk showing up on an IP black list. Reputation services put back pressure on mail originating domains to suppress spam generators and make it difficult to exploit any open mail relay or use Zombie mail servers.

With these changes, the ease of originating mail changes dramatically. It no longer suffices to have a computer with a dynamic IP address or to work through an open mail relay. You're left with the choices of sending from a mail domain with an at best questionable reputation (the equivalent of pasting a big "Block Me!" sign on your messages) or going with a reputable mail domain and trying to work around their controls. Legitimate volume email marketers will move to use a small set of careful bulk mail processors, and those services will carefully maintain relationships with the large email receivers, but all at some price and certainly under careful mutual scrutiny. As this process continues, spammers are driven from legitimate mail domains to ones that are less diligent about controlling spam. This migration in turn worsens the reputation of the less diligent domain, so their mail is increasingly blocked, forcing them to choose whether they improve their outgoing spam control or suffer the consequences of not doing so. If they do improve outgoing spam control then the spammer must migrate again to an even less diligent domain. Relatively soon this process distills the world of email domains into white domains (spam free) and black domains (known spammer). It becomes very hard for a spammer to find a grey domain and lurk behind the uncertainty.



Spammers have tried to evade these controls by creating their own domain registries so that new domains can be created as needed. These tactics can be neutralized by monitoring the reputation of domain registrars as well as the domains they create.

Will all this actually happen? Will mail service providers police their own users to suppress spam? Will network service providers work with one another on these community schemes? We're pretty sure the answer is yes, and the motivation is simple economics again. If changes don't occur, spam is a dominating cost problem for all the legitimate players in the mail community. The amount of money each can save by solving the spam problem pays for a lot of cooperation.

8. Reputation Monitoring and Reporting

With authenticated mail addresses we have the beginning point for observing and reporting on the reputability of a mail source. If you can trust that mail is from `somedomain.com` (which you can't do without address authentication) then you can ask a few simple questions:

- Is `somedomain.com` a known sender of mail or did they just pop up out of nowhere?
- Is `somedomain.com`'s current mail sending activity expected or anomalous? Do they normally send 100 messages a day and they have already sent 1,000 today?

Additionally, if you have access to output from one of the major spam filtering systems, you can make meaningful comments on whether or not the mail sender is believed to send spam (keep in mind that we're going to use our very simple definition of spammer before you head down one of the rat holes). In almost all cases (given that self policing will occur) known mail senders will be viewed as spam-clean.

There are many ways in which a reputation reporting service could be constructed and operated as a business. It isn't our goal here to tell you precisely how that will happen, but rather to tell you that it will happen. It will happen because it is an important part of solving the spam problem, because it is quite feasible to do it, and because the cost of doing it is far less than the cost of the spam problem today. The key mail processors (very-large organizations and enterprises and leading service providers) will find high value in having one or more reputation authorities exist. Only a few are likely to run an authority themselves (although all of them have sufficient information on inauthentic or previously unseen mail domains and mailers with bad reputations to do so) because of the conflict of interest it would create since they also sell



networking marketing and commerce services. It will be better for all of them if they can support a shared community definition of reputation that they don't control.

For some markets like education where extreme freedom of speech is argued, having such an external source of authority is the only way that a mail control mechanism is likely to work. None of the university mail managers we talked to want to be personally responsible for what mail is or is not blocked, and would much rather point at an external authority shared by other like institutions. Many of the ISP's we have discussed spam with share the academic desire of having someone else responsible for making the mail blocking decisions. It is much easier for any business or service provider to make mail blocking decisions based on some broadly used service or definition than it is to have to justify the same decision make locally.

A critical requirement for any reputation authority is access to enough mail transaction information to form reliable measures of a sender's reputation and behavior. The obvious ways of building a reputation authority include doing something that incorporates information submitted by mail senders or deriving something from one of the mail filtering services that already see a meaningful percentage of email traffic in the US. AOL, Yahoo, and Microsoft all process enough mail to accomplish the task.

9. Putting the Pieces Together

Our solution depends on these pieces:

- Ability to recognize known bad IP numbers (most large mail processors already do this).
- Ability to authenticate that the sending mail domain is who they claim to be.
- Ability to verify that the mail sender is a known sender of mail.
- Ability to get an indication as to whether the sender's recent mail behavior is expected or anomalous.
- Ability to get a reliable rating of the sender's spam reputation.

We put these pieces together as shown in the following chart giving the processing path for incoming mail:



Authenticated mail sender				IP number of sending server not authentic or known bad
Known email behavior			Unknown email source	
spam clean reputation and normal behavior	Questionable spam Reputation, or anomalous behavior	Reputation as spammer		
<i>Categorically pass</i>	<i>Filter</i>	<i>Categorically block</i>		

Using Reputation

If the mail sender’s SMTP header and IP address don’t authentically match, or if the IP number of the sender is reliably known to be bad (e.g. an open mail relay) then we categorically reject the mail (as early in the process as possible, ideally at the network boundary).

If the mail sender is authentic, and the sending domain is known and behaving normally (the typical case) then the mail is accepted and not filtered (a user may choose to filter personal mail more stringently but this scheme should achieve our spam percentage goals easily).

If the mail source is unknown (has no history) or is a known source of spam then the mail will be categorically rejected.⁷

In the case where the mail-sending domain is known but either their spam reputation is unclean (but not “bad”) or their mail sending behavior anomalous, then traditional content filtering will be applied.

It seems too simple to work, doesn’t it? We’ve presented the idea to lots of experts, and so far no real pitfalls have been identified. The scheme clearly depends on

- Sending domain authentication.
- Responsible reputation authorities.
- Self-policing of sending domains.

Fixing today’s spam problems seems more than enough economic justification for these solutions to be put into place.

⁷ At this point in the discussion, expect lots of rat hole arguments about how under these conditions no new company could ever send mail. There are lots of ways around this conundrum ranging starting with using a reputable mailing service (e.g. what Microsoft BCentral could be) and leveraging their established reputation to assuming that you could provisionally register with a reputation service while your mail reputation developed.



10. **Summary**

The economic and practical crises that occurred when spam volumes exceeded 50% of all mail traffic led to consideration of spam as a network (denial-of-service) problem rather than an undesired content problem. It is now clear that it is much easier to identify and control spam by watching message transmission behavior than it is to accurately identify spam by the message content.

Since that time the eye of the spam storm has passed with the emergence and application of network-focused spam monitoring and control mechanisms.



References

- Sender ID:
<http://www.microsoft.com/mscorp/safety/technologies/senderid/default.aspx>
- SPF:
<http://spf.pobox.com/>
- Certified Server Validation (CSV):
<http://www.mipassoc.org/csv/>
- Domain Keys:
<http://antispam.yahoo.com/domainkeys>
- Yahoo, Cisco Merge Email Specs:
<http://www.internetnews.com/bus-news/article.php/3519066>
- Messaging Anti-Abuse Working Group (MAAWG): www.maawg.org
- Comments by MAAWG on Sender ID/SPF:
http://www.maawg.org/about/whitepapers/spf_sendID

