



EQUAL Workshop on Improving the Quality of Email Services (Spam Suppression)

Wednesday 9th December 2009

University of Amsterdam

Amsterdam, the Netherlands

Minutes by: John DYER

TERENA

27 December 2009

1. Introductions

The workshop on improving the quality of email services was organised as a result of discussions in the TERENA taskforce on the Management of Service Portfolios (TF-MSP). Many members of the taskforce had indicated that their NREN had an interest in improving the quality of email services by actively working on the suppression of SPAM. Several NRENS have already put their own solutions in place. The purpose of the workshop was to discuss these solutions and explore the possibilities for future joint action.

2. SURF-mailfilter, Paul Dekkers, SURFnet

Paul Dekkers started his presentation by explaining the features of the SURF-mailfilter. The mailfilter is a hosted inline scalable system with redundant capability. It provides a high-quality filtering system that detects spam, viruses and phishing attacks. Fine control over the system is delegated to institutional administrators and end-users. The system was conceived in 2005 as a service to small institutions however it has grown to serve around 60 institutions hosting 1600 domains and currently has a throughput of 15-30 million messages per day. The system provides improvement over locally run services and brings costs savings for the institutions.

SURFnet compared internal development of a system with commercial solutions. The investigations found that the product CanIt from Roaring Penguin provides the power and flexibility needed. The feedback from the mailfilter users is that the service is easy to use and they appreciate its capabilities, performance and availability.

The system can be expanded considerably as the amount of spam increases. SURFnet will continue to develop the system to improve results. Paul explained that SURFnet is interested in working with others and possible collaboration.

Presentation Material is available from:

www.terena.org/activities/tf-msp/meetings/20091209-equal/equal-surfmailfilter02.pdf

Roaring Penguin site: www.roaringpenguin.com

3. RACE, Spanish academic mail network, Jesus Sanz de las Heras, RedIRIS

RedIRIS has been working for the last 15 years on coordinating the improvement of the IRIS-MAIL e-mail services for the Spanish academic community. This coordination undertaken with a working group has allowed RedIRIS to acquire knowledge and experience which has been used to underpin many initiatives and services. Current activities include reputation services, spamtrap networks and whitelist & blacklist services. Jesus explained the uniform strategy for RedIRIS mail community known as RACE. It includes evaluation, guidance and best practice. In addition, RACE provides an accreditation and certification service for universities. RACE uses 33 weighted technical criteria for accreditation which include items such as: Anti-relay & Logs policy; reverse DNS records; abuse and postmaster mailboxes. The accreditation process is undertaken by a team of ten volunteers under the co-ordination of RedIRIS.

Presentation: www.terena.org/activities/tf-msp/meetings/20091209-equal/RACEamst.ppt

RACE (In Spanish): <http://www.rediris.es/race/>

4. RedIRIS Blacklist, Whitelist & Spamtrap, Francisco Monserrat, RedIRIS

Francisco opened his presentation by explaining the problems associated with the use of blacklist for spam suppression. Whilst blacklists can be effective the processing load can be heavy, slowing down the delivery of mail. In addition it can take some time to remove addresses from the list in the case of false positives. The main problem for the users associated with the use of over-zealous blacklists is the non-delivery of mail. As a consequence RedIRIS has developed a service based on the RKS product from Sandvine. The RedIRIS service, known as IRISRBL integrates information from several sources: blacklists; whitelist & spamtraps. The system has a simple web-based interface to allow the postmaster to manage the lists ensuring that false positives can be dealt with quickly. More than 60% of RedIRIS constituency is using IRISRBL which results in a load of around 350 DNS queries per second. In order to avoid problems with blacklisting of SMTP servers a whitelisting system was developed. The whitelisting service makes SMTP server information available as a DNS based list or configuration files for the servers themselves. Francisco pointed out the use of whitelists can reduce problems caused by blacklists alone, but it does not provide assurance regarding the mail that gets through.

Presentation: www.terena.org/activities/tf-msp/meetings/20091209-equal/equal-fjmc.ppt

Sandvine site: <http://www.sandvine.com>

5. Reputation Systems, Jaime Pérez, RedIRIS

Jaime Pérez explained some work being undertaken by members of the TERENA Taskforce TF-EMC2. The group has been exploring the application of reputation systems. Reputation systems for fighting spam using blacklists and whitelists is clearly an interesting case for TF-EMC2 to explore. TF-EMC2 is also exploring IRISRBL plus Attribute Authority to store and share reputation (black/white lists).

Jaime encouraged work in this area to be coordinated with the work of TF-EMC2. He invited EQUAL delegates to participate in the ongoing work on reputation at the wiki on the RedIRIS site (address given below). The wiki is eduGAIN enabled with anonymous read, however the write access policy will change in the near future.

Presentation available from:

www.terena.org/activities/tf-msp/meetings/20091209-equal/reputation.pdf

Related links:

wiki.rediris.es/reputation

www.terena.org/activities/tf-emc2

6. RENATER anti-spam-service, Jerome Durand RENATER

In 2008 RENATER undertook a survey of anti-spam solution deployed around the RENATER community. As a consequence RENATER created a working group of experts looking at the topic of anti-spam measures. The result of activities in the working group was that RENATER undertook a pilot, issued a tender for procurement and deployed a production service before the end of 2009. The service runs on two Bizanga appliances (MTAs), with the components: Vade-Retro anti-spam; Vade-Retro anti-virus and Doctor Web anti-virus. RENATER expect to be able to accommodate up to 2 million mailboxes using this architecture. Currently they have licences for 500,000 e-mail boxes

The service processes ingress SMTP using: protocol filtering; RBL; SPF rejects; LDAP checks; blacklists; whitelists and content filtering. Antivirus is available as an option. The service provides both transparency and flexibility and there are plans to develop a web portal for anti-spam configuration. Further development plans include developing a RENATER whitelist / blacklist; Ipv6 support; outgoing mail filtering; DKIM & domain reputation services.

Presentation:

www.terena.org/activities/tf-msp/meetings/20091209-equal/SSU-ANTISPAM-TERENA-EQUAL-20091208.pdf

Bizanga Press Release:

http://www.bizanga.com/company/press_room/bizanga_powers_renater_antispan_service.php

Suppliers Websites:

<http://www.vade-retro.com/>

<http://www.drweb-online.com/en/index.asp>

7. EDU whitelist for e-mail, JP Velders, UvA & SURFcert

The EDU whitelist for e-mail is a DNS based list with the primary goals of reducing the negative impact of antispam whilst allowing research and education e-mail services to remain viable. The project aims to demonstrate that such a service is viable for a large community. A secondary goal is to integrate the results into large non-edu initiatives.

The EDU whitelist employs a policy, which constituencies must enforce (through their CSIRT) on their constituents in order to have their outgoing SMTP mailservers listed. The goal of the policy is to ensure that people using the EDU whitelist on their mailservers know what to expect of server addresses listed on the whitelist, and thus set a common standard.

So far the work of EDU whitelist has provided limited proof of concept. JP Velders invited delegates and their communities to participate in the initiative to provide feedback so that it may reach a critical mass. He also suggested a multi-NREN/TERENA infrastructure.

Presentation:

www.terena.org/activities/tf-msp/meetings/20091209-equal/20090925-TF-Equal-JPV-EDUwhitelist.ppt.odp.pdf

Edu whitelist site

<http://test.eduwhitelist.net/>

8. Important and relevant for e-mail systems, Dick Visser, TERENA

Dick Visser explained that as postmaster for TERENA he has to deal with suppression of spam for: regular e-mail; e-mail distribution lists and web-applications that send e-mail.

E-mail list authentication is trivial/non-existent. Not only are technical issues to be solved, but sometimes management reactions to receiving spam can lead to instructions to put in place over-zealous filters which in turn can cripple e-mail services to the general user. He explained that in his opinion solutions that rely on a binary (YES/NO) choice are problematic as they cannot deal adequately with questionable cases where MAYBE would be a more appropriate response.

Weighted score based systems such as those in products such as SpamAssassin are better. He went on to say that in cases where the sender domain has some way of controlling the decisions made by receiving servers binary decisions can be good. This however requires that manager of the sending MTA to properly configure its DNS entry attributes for parameters such as: SPF; DKIM etc.

Dick closed by saying that postmasters should be able to see spamtrap rejections in their logs and blacklist the offending sites.

Presentation:

www.terena.org/activities/tf-msp/meetings/20091209-equal/TERENA%20Postmaster.pdf

SpamAssassin: spamassassin.apache.org/

9. Lessons learned from operating Mail Dike, Magnus Strømdal, UNINETT

"The Dike" project started in 2004 as a central service operated by UNINETT. It provides users with self-service configuration. It currently serves over 200 domains and has had IPv6 capability added along with content inspection (2007), blacklisting (2008) and support for recipient lists (2009). The software being used in Dike is open source.

Magnus explained that the system is modular with Dike I comprising of blacklisting, greylisting and whitelisting functions whilst Dike II provides anti-virus and anti-spam services. The system runs on three machines (one being a hot spare) and processes around 2.3 million messages per day which is within its capabilities, however increasing the load above makes it struggle. The operators of Dike have found that logging is a bottleneck, but both machines send out daily logs to each customer postmaster.

The operators have also found that when a customer MTA has been offline for a while, the flood of mail from Dike when it returns can be a problem for the MTA systems.

The feedback from the end users is positive with 98% of all spam being blocked.

Presentation:

www.terena.org/activities/tf-msp/meetings/20091209-equal/Dike-equal-20091209.ppt

Link to Mail Dike information: <http://software.uninett.no/maildike/>

10. Open Source Mail Filtering at the ETH Zurich, David Bruce McLaughlin, ETH

The ETH-Z mail filter does not accept message and then silently delete them. It either accepts them or explicitly rejects them. It also provides users with a tag-only option along with personal blacklisting & whitelisting of sender addresses and domains. Malware & phishing messages however are rejected without regard to user preferences.

The open source components that are used to build the system include:

- EXIM - mail gateway software
- ClamAV - virus/phishing detection network + 3rd-party signatures
- SpamAssassin - spam content-filtering software + SARE & local rule-sets
- DCC - mass-mailing detection network (called from SpamAssassin)
- Razor - spam detection network (called from SpamAssassin)

David explained the pipeline like operation of the system with some graphics that can be found in the presentation.

www.terena.org/activities/tf-msp/meetings/20091209-equal/ethz-mailfiltering.pdf

11. E-Mail Quality is a matter of good System Hygiene, Eliot Lear, CISCO

Eliot started his presentation by asserting that reputation is both important and dynamic. The identity of the sender should be known to provide some assurances as to who sent what. A recent estimation is that 85% of all mail worldwide is spam and of that, 95% of this spam is coming from BOTS. Only around 14% of worldwide e-mail is thought to be legitimate e-mail.

In addition to the normal phishing attacks, the phenomena of spear-phishing which targets and individual or organisations is thought to be growing rapidly. Because spear-phishing is so tightly focussed it does not show up in the statistics. Whilst traditional spam has almost been conquered by the approaches presented during this workshop, new threats such as spear-phishing are becoming more problematic.

Users of social networking sites such as: Facebook (350 million users); LinkedIn (53 million users) provide those intent on spamming and the like with huge repositories of private information. This information is what hackers and spear-phishers need to operate. Eliot went on to demonstrate some examples of problems end-users can encounter differentiating legitimate mails and from sent by bogus sites.

What is needed is a strong and consistent policy around the community. The TERENA community is a unique position to be able to agree such a uniform policy for its constituency. Such a policy might include the use of something like DomainKeys Identified Mail (DKIM) and Author Domain Signing Practices (ADSP). With ADSP domain owners can publish that all mails from their domain have a corresponding DKIM signature as defined in RFC 4871.

The main objective of these activities is to devalue the BOT.

He also mentioned possible approaches to awarding privilege and assurance, but some of these rely on relatively expensive hardware for users to carry with them.

As the availability of trusted federated identity providers becomes more pervasive it should become the required method of authentication. It might even provide a range levels of assurance for different purposes.

To conclude Eliot mentioned the work of the Mail Anti-Abuse Working Group (MAAWG). The working group is a global organization focusing on preserving electronic messaging from online exploits and abuse. Their goal is to enhance user trust and confidence, while ensuring the deliverability of legitimate message. MAAWG publishes a range of documentation on Best-Practice, Surveys and Recommendations.

MAAWG meets three times each year, usually once in Europe. The next European meeting will be 8-10 June 2010 in Barcelona, Spain. It was suggested that TERENA approach MAAWG to discuss the possibility of membership on behalf of the community.

Presentation:

<http://www.terena.org/activities/tf-msp/meetings/20091209-equal/cisco-eliot-lear.ppt>

Link to MAAWG: <http://www.maawg.org/home>

12. Discussions and the Way Forward.

The work of Internet2, the Internet Society, and the In Common Federation are collaborating on an activity to promote deployment of DKIM technology describe in the IETF RFCs at higher-education and research sites was mentioned. DKIM promises a variety of improvements in email reliability, trustworthiness, and usability when deployed in a way that takes advantage of existing trust communities, such as research and education federations. There is also a proposal for a DDX Pilot Project using DKIM to create an e-mail trust channel.

Another issue raised was that of the introduction of carrier-grade NAT routers in an attempt by ISPs to mitigate the effects of the shortage of IPv4 address space. The effect of the further proliferation of NAT will be to prevent the disclosure of the precise identification of the end-node. Knowing the precise address of end-nodes is extremely helpful for security purposes. If or when the world migrates to IPv6 this should provide a one-to-one relationship between end-node and address and lead to increased security.

There seemed genuine surprise about the amount of work that is currently underway in our community and a strong desire to capitalise on all the knowledge and expertise by sharing and collaborating. Delegates thought that would best be achieved under the auspices of TERENA. The work cuts across several TERENA taskforces such as TF-EMC2, TF-CSIRT, TF-MOBILITY and TF-MSP. Consequently it was agreed that further work and discussions should initially take place in an ad-hoc EQUAL group with invitations for members of the existing taskforces to participate. In order to make this easier it was agreed that the ad-hoc group should try and co-locate with a meeting of the one/some of the interested taskforces.

The consensus is that ad-hoc group should consider undertaking the following activities:

- Continued exchange of information between NRENs
- Considering a community-wide whitelist / blacklist exchange or system
- Establishing and recommending: Common Policies; Standards and Best Practice
- Possible joint-NREN procurement of commercial offerings
- Exploring what can be done in respect of the proliferation of private information online
- Seeking community membership of MAAWG
- Maintaining the EQUAL email distribution list and hosting online material about the activity either on the TERENA website or the secure wiki hosted by TERENA.

The timetable for the activities is:

- Meet early in 2010 for ½-1 day, in order to establish a concrete plan for further activities. The meeting should possibly be co-located with one of the following meetings:
 - TF-CSIRT, 25-26 January 2010, Hamburg, Germany
 - TF-MSP, 4 February 2010, Rome, Italy
 - TF-EMC2 and TF-MOBILITY, 16-18 February 2010, Vienna, Austria
 - MAAWG, 8-10 June 2010, Barcelona, Spain
- A BoF should be arranged during TNC2010 (31 May - 3 June 2010, Vilnius, Lithuania) in order to give the work greater exposure to the community and recruit more resources to the activity.

Additional Links:

TERENA Task Forces: <http://www.terena.org/>

DDX BOF (DKIM Deployment) hosted by Internet2:

<https://spaces.internet2.edu/display/ddx/DKIM+Deployment>

Vouch By Reference (VBR):

<http://www.domain-assurance.org/>

Registrations for the Meeting

In Person

First name	Family name	Organisation	Country
Aris	Adamantiadis	BELNET	Belgium
Serge	Aumont	CRU	France
Filiz	Bektas	Switch	Switzerland
Wolfgang	Breyha	University of Vienna/Aconet	Austria
Roberto	Cecchini	GARR	Italy
Paul	Dekkers	SURFnet	Netherlands
Dirk	Dupont	Belnet	Belgium
Jerome	Durand	RENATER	France
John	Dyer	TERENA	- N/A -
Licia	Florio	TERENA	- N/A -
Aurelija	Gefeniene	Vilnius University	Lithuania
Francesco	Gennai	ISTI - CNR	Italy
Martin	Kämpf	SWITCH	Switzerland
Alexandros	Kosiaris	GRNET S.A.	Greece
Andrea	Kropacova	CESNET	Czech Republic
Eliot	Lear	Cisco Systems Ecole des Mines de Paris	Switzerland France
Jose-Marcio	Martins da Cruz	ETH Zurich	Switzerland
David	McLaughlin	RedIRIS	Spain
Francisco	Montserrat	RedIRIS	Spain
Jaime	Perez	RedIRIS	Spain
Fredrik	Pettai	NORDUnet A/S	Sweden
Jernej	Porenta	ARNES	Slovenia
Lino	Santos	FCCN	Portugal
Jesus	Sanz de las Heras	RedIRIS	Spain
Brook	Schofield	TERENA	Netherlands
Alba	Shahin	ISTI - CNR	Italy
Magnus	StrÅmdal	UNINETT	Norway
Maurice	van den Akker	SURFnet	Netherlands
JP	Velders	SURFcert / UvA	Netherlands
Dick	Visser	TERENA	Netherlands
Klaas	Wierenga	Cisco	Netherlands

Remote participation

Claudio	Allocchio	GARR	Italy
Brian	Boyle	HEAnet	Ireland
Virginia	Calabritto	CASPUR	Italy
Massimiliano	Filacchioni	CASPUR	Italy
Carles	Fragoso Mariscal	CESICAT	Spain