

Email Authentication Summit- Comments Matter Number P044411

Federal Trade Commission
Office of the Secretary, Room 159-H, Annex V
600 Pennsylvania Avenue NW.
Washington, DC 20580
authenticationsummit@ftc.gov

Key Points:

- An authenticated HELO domain is the only sufficiently strong name identifier within a mail stream to safely identify the message source and allow a reputation assessment of those accountable.
- The HELO domain is analogous to the letter postmark, but is currently unreliable due to a protocol defect resulting from limitations in DNS information. Without changes to RFC 2821, the HELO domain must be allowed to fail authentication.
- Adding a DNS Service record to assure HELO domain authentication and additionally validating authorization can be implemented without any negative impact upon existing applications.
- Making the authenticated HELO domain name visible to the user would be an effective and safe deterrent against phishing and spoofing, as this identity would be the most difficult to spoof as a means to provide false assurances.
- The authenticated HELO domain identifies the server holding the SMTP log needed for criminal enforcement.
- Authenticated HELO domains enable a simple and safe means to associate authorized mail transfer agents with the mailbox domain through the publishing of a simple name list. The alternative text scripts, as used with SPF or Sender-ID, to obtain addresses for a large array of hosts is inherently perilous, and must be discouraged.

Reduction in Spam

A strategy predominately used to reduce spam and preserve network resources is called the Real-time Black-hole List (RBL). This list is based upon the IP addresses of the connecting client where, when a record is returned after the address is queried, the client is refused with a reply that often names the list responsible for the rejection. The address is used as the identifier because there are currently no validated names associated with a mail stream. The lack of a validated name makes it difficult to know when an address has been reassigned and the reverse DNS name directory is too poorly maintained to be of use for this purpose as well.

An address is currently the only strong identifier that is authenticated by way of interaction within the transport protocol within a mail stream. With the potential number of addresses being fewer than the potential number of names, tracking addresses that are

sending abusive mail is effective at making access difficult for abusers. The success of the RBL, as a tracking mechanism, has led to tactics such as Trojan programs as a means to commandeer new addresses. Introduction of IPv6 addressing also changes the number of addresses compared to names. To move away from the address based methods for tracking abuse, a strongly authenticated name is required.

The strength of this authenticated name must be as strong as that of the IP address due to litigation risks. The strength of the abuser's identity is paramount for defending a reputation assertion. The selected name must identify the domain administrator which controls the specific mail transfer agent and is accountable for the requisite security. For the purpose of abating spam, there is only a single domain name within the mail stream which meets this requirement. This would be the HELO domain name.

All other domain names within the mail stream are based upon an unverifiable assumption of the mail stream integrity, regardless of any association within domain name server records. Any such association which serves to authorize the sending of mail on behalf of a specific mailbox domain, does not imply the mailbox domain administrator accepts accountability for the performance of this authorized mail transfer agent when controlled by a different administrator. The administrator of the mailbox domain not in control of the mail transfer agent can not be expected to make assurances or take corrective measures. The consumer, as well as the reputation service, would be threatened by an overly broad application of accountability where litigation becomes the only recourse for an unfair reputation assessment. Accountability must be constrained to the domain administrator, identified by the HELO domain, in control of the mail transfer agent. This administrator is expected to monitor SMTP error logs, track abuse@ complaints, maintain security, and disable problematic accounts as a means to control access.

The SMTP protocol has all the necessary components to assess this accountability. SMTP presents a client name in the HELO exchange at the beginning of an SMTP session that is recorded within the message RECEIVED headers. Unfortunately, failure to authenticate the HELO domain name must not be used to refuse mail from the client, as required by the SMTP standard RFC 2821. Any name may be presented in the HELO exchange, where authentication, although vital, is often missing. The steps to rehabilitate the HELO name to ensure it will authenticate also enables an assertion of authorization to protect machines that may be compromised. When this domain name is both authenticated and the authorization for sending mail is validated, then this domain can be safely accredited "by name" for traffic emitted.

Having authenticated and verified the authorization of the SMTP client by name, should there then be criminal activity detected, the server logs can be readily uncovered. This would be a substantial improvement over just having an IP address, as now an

authenticated domain has made expressed authorization for the mail sent. The registration of the domain also offers the name of the applicant and a possible money trail from the purchase of the domain.

Prior to the break up of the IETF MARID working group, there was a suite of standards proposed under the title of Client SMTP Validation, CSV by Dave Crocker, Douglas Otis, and John Leslie. A new working group will carry forward these standards, as the chairs of the now defunct MARID working group held back consideration of these standards. The salient aspect of this suite of proposed standards was aimed directly at repairing the problem which prevents the HELO domain from being authenticated. This suite further requires the domain to expressly authorize the client for sending mail.

Modification of standards

There is one aspect of SMTP which needs to be addressed. The authentication of the HELO domain name. This can be done easily by introducing a DNS record specifically tasked to ensure a complete list of addresses will be available when queried. This can be done within a single binary query of a DNS Service record. This solution will be addressed by a new work group formed within the IETF to carry forward the Client SMTP Validation (CSV) specification suite that was started in the now defunct MARID work group. See:

<http://www.csvmail.org/>

<http://www.ietf.org/internet-drafts/draft-ietf-marid-csv-intro-01.txt>

<http://www.ietf.org/internet-drafts/draft-ietf-marid-csv-csa-01.txt>

If there is a desire to constrain the mail transfer agents authorized to send mail on behalf of a mailbox domain, there is also an IETF draft named “Mail Policy Record (MPR)” that illustrates how a simple list of names can define the authorized sources of mail. The use of this name list does not invite exploitation, as it will not normally act as a gate-keeper. This list can alert the user to a message carried outside the nominal mail channel. Financial institutions likely the subject of phishing would be well advised to use consistent names for their HELO domains and to publish their nominal sources by way of this name list. See:

<http://www.ietf.org/internet-drafts/draft-otis-marid-mpr-00.txt>

Should there be a need for an immediate solution that rejects obviously spoofed mail, then the use of the Mail Policy Record draft name list used in combination with Client SMTP Validation, offers a safe alternative to either Sender-ID or SPF. Use of the name list to indicate nominal sources for a specific mailbox domain avoids the perilous use of domain name server text scripts that attempt to compile all the addresses for a vast array of hosts. The text scripts used to implement both the SPF and Sender-ID schemes

promoted by POBOX.COM and Microsoft represent a hazard from several perspectives. The time needed to process these SPF and Sender-ID address lists can be exceeding long. If these drafts adhered to timeouts for domain name server transactions, the time required may exceed hours. By not adhering to these timeouts, the User Datagram Protocol (UDP) exponential back-off is violated and thus does not provide the requisite congestion avoidance. Even with the early timeout of 200 seconds as specified, malicious scripts could easily effect a denial of server attack aimed at disabling the checks.

Because both of these script schemes require a specialized parsing program, the source port used for domain name server queries may be constrained by the application running the parser. The hundreds of potential records and different domain name servers referenced allows a “birthday” attack staged using just a single DSL network to poison the records held in the domain name server’s cache. This risk is acute when the domain name server does not aggregate pending queries as with Bind 8, a popular version of the domain name server.

DNS cache security overview by Joe Stewart:
<http://www.securityfocus.com/guest/17905>

Compatible with Legacy

CSV offers true consumer protections when the authenticated HELO domain name is visible and when those responsible for security receive the reputation assessment. Until the consumer is using a mail client able to present the authenticated HELO domain name, there would be no expectations of increased assurances. Unlike SPF or Sender-ID, the use of CSV will not interfere with the normal use of mail and allows rapid deployment, as the impact to legacy systems would be negligible. Unlike SPF or Sender-ID, the user is not expected to forgo use of their favorite mailbox address. Unlike SPF or Sender-ID, it is the provider that receives the reputation allowing their customers the freedom to find other providers in the event the provider is blocked for allowing abusive mail. Unlike SPF or Sender-ID, older mail clients can not be used to provide false assurances.

Handling of unauthenticated messages

Initially, the phase-in of CSV moves from the reliance upon the client’s address to the authenticated and authorized HELO domain. Unauthenticated messages that fail to provide the newer HELO domain authentication records may receive a “slow path” approach as a means to limit the extent of potential damages. By making only the authenticated HELO domain names visible to the user, this too would be an incentive to deploy CSV as a means of providing this increased assurance. This added assurance may also be used to reduce the chance of being “filtered.”

Erroneous results

Unlike Sender-ID or SPF, there is little opportunity for erroneous results with CSV. With either Sender-ID or SPF, a shared mail transport agent opens up the possibility of spoofing the mailbox domain. Both of these schemes depend upon the unverifiable assumption that the sending mail transfer agent has performed the needed checks and that the lists are “closed.” With the information published for Sender-ID or SPF, a malicious attack may only require an assertion of the mailbox domain. CSV records are unique for each mail transport agent and make no assumptions of the mail stream integrity. Unlike Sender-ID or SPF, there is little risk of the domain name provided by the HELO transaction being spoofed. This provides consumers an identity which can be safely relied upon whereas SPF and Sender-ID do not. Institutions would be well advised to be consistent with their naming conventions within the HELO transaction to ensure reliable recognition of the domain name.

Compatible Enhancements

The use of the Mail Policy Record proposal, which allows the listing of the nominal mail channel for a mailbox domain, can be added without incurring any significant overhead. The use of the name list can be used to alert users of a possible spoofing without exposing the domain name system to the risk of a denial of service attack or being poisoned as is possible with Sender-ID and SPF. The alternative Mail Policy Record proposal permits a simple name list be obtained within a single domain name server query to ensure proper source recognition.

Proprietary Nature

Unlike Sender-ID, there is no proprietary algorithm for sorting the field being checked with CSV. By always using the same HELO parameter there is less risk of differences in the application of a complex algorithm that could allow a sneak path for spoofing.

Proprietary Services

CSV does not depend upon any proprietary services. Right Hand Side Black-hole Lists (RHSBL) have been in general use and are based upon structures established by MAPS, which provided the original RBL services.

Effect on forwarding

CSV will not interfere with forwarding. SPF and Sender-ID do interfere.

Effect on mobile use, roving users, and Mailing lists

CSV will not interfere with the operation of any of these mail applications or uses. SPF will seriously impact mobile and roving users. Both Sender-ID and SPF may potentially

