DRAFT

Client SMTP Validation

Best Practices

This document discusses the actual use of Client SMTP Validation, from the standpoint of both CSV publishers and receiving SMTP servers. It starts with a non-technical overview of CSV; then moves on to what managers of sending SMTP clients should do (and why), and finishes with what managers of receiving SMTP servers can do (and why).

Readers of this document are expected to have a minimal understanding of SMTP (Simple Mail Transport Protocol) and DNS (Domain Name Service), but the necessary details will be explained.

Overview

Client SMTP Validation concerns itself with the so-called “HELO” string. This string is the parameter contained in the first command issued for every SMTP session, and is intended to identify the sending SMTP client. Generally the “HELO” string is either the name of the particular computer which initiates the SMTP connection or the name of the domain responsible for managing it, although RFC 2821 allows for other values.

Client SMTP Validation does not impose any restrictions on what may be used for a “HELO” string, but instead allows for additional information to be found by a DNS query, and for that information to validate the intent by domain management to take some level of responsibility for what that sending SMTP client may send. This document will recommend particular values to be used in the “HELO” string, but CSV specifications do not require this.

Receiving SMTP servers which implement CSV will extract the “HELO” string and do a DNS lookup for the “client” service and “SMTP” protocol to find whether the management of that domain has published an intent for this particular client to be authorized to send email. CSV does not require any particular action based on that lookup.

CSV does not require any particular method of evaluating the reputation of the domain named in the “HELO” string, but senders should expect that receivers will do some reputation checking. This should present no problem if your (sending) domain has a good reputation. Should this not be the case, we explain the CSV accreditation system in the section on managing sending SMTP clients.

Managing sending SMTP clients

We recommend that managers of sending SMTP clients publish CSV records as soon as practical, in order to allow your good reputation to be recognized. By publishing the CSV “SRV” record, you allow receiving SMTP servers to authenticate your sending SMTP client by domain-name, instead of only by IP address.

Prior to CSV, essentially all reputation of sending SMTP clients was based only on the IP address of that client. This system has worked well for those clients whose IP address has never been blacklisted; but isn’t working well for those which wish to get removed from a blacklist. Many blacklists have no automatic process for removal of IP addresses which once had a problem but have been free of problems for months — or even years. What the managers of such clients want, of course, is to get removed from blacklists in a matter of minutes after the problem is resolved. This has proven impossible because:

1) Many blacklists don’t bother to tell you they’re listing your IP address;

2) For fear of lawsuits, many blacklists obscure their contact information;

3) Blacklist maintainers have no reason to believe your statement that a problem is fixed;

4) The folks “responsible” for the IP block usually disclaim all responsibility.

The result of all this is that well-behaved sending SMTP clients may be on dozens of blacklists before they learn there was a problem, and it’s likely they’ll never learn about all of them unless the senders of individual emails ask the intended recipients whether they received them. It can take hours of effort to find out *whether* there’s a way to get off each of these blacklists, and it is likely to be several days after that to find out whether it worked.

CSV sets out to support a “whitelisting” system to bypass any existing IP-based blacklists when:

1) the managers of the sending SMTP client explicitly advertise that they accept responsibility; and

2) a trusted reputation service agrees the sending domain will act responsibly.

We recommend publishing an A)ddress record matching the HELO string and returning the IP address which the sending SMTP client uses to originate SMTP sessions. (In all likelihood, you’re already doing this: if not, it might be easier to change the HELO string to match an A)ddress record you already are publishing.) If you use more than one SMTP client to send email, we recommend assigning separate subdomain names to each, and using those as the HELO string. If you use the same HELO string for very many clients, it may not be possible for receiving SMTP servers to authenticate that name to the IP address of a particular client.

In addition, you should publish a “SRV” record showing the “service” as “_CLIENT” and the “protocol” as “_SMTP”, returning the A)ddress record(s) for the HELO string. In the so-called “bind” format (the most prevalent DNS server software), this would be:

```
_CLIENT._SMTP.thishost.sendingclient.com    SRV   1 2 0   thishost.sendingclient.com
```

You should also evaluate potential CSV-compliant accreditation services, although there is no need to publish records for these unless and until your sending SMTP client is listed on some blacklist. However, it is best not to put this off too long, because you may not learn about being blacklisted until months after it happens.

CSV-compliant accreditation services agree to publish DNS records certifying that your domain has policies and practices which they believe show that email you send is unlikely to be abusive, and/or that you respond to reports of abuse promptly. These accreditation services must publish, in human-readable form, the conditions they accredit; and it is up to whatever reputation service that receiving SMTP servers may use to decide whether these conditions are sufficient to bypass the anti-spam checks that receiving SMTP server would normally use.

Once you select one or more CSV-compliant accreditation services, you publish “PTR” records for them. In “bind” format, these might be:

```
thishost.sendingclient.com                  PTR   _VOUCH._SMTP.accreditation.com
```

Since it is not under your control how the various reputation services might view a particular accreditation service, it is best to publish several of these. But we do not advise listing hundreds of them, because a reputation service might only accept as many of these as can fit in a single UDP reply. Half a dozen is probably a reasonable limit.

Managing receiving SMTP servers

We recommend that managers of receiving SMTP servers enable the CSV-checking option if already

DRAFT

available in their MTA software; or, if it is not yet available in their MTA software, set up an alternate server for testing purposes using MTA software which supports CSV-checking. Obviously most managers will not want to abandon their current MTA software until they are comfortable about managing the software they might switch to. However, there are numerous open-source options available for no initial cash outlay, and most of these can accomplish most MTA functions you might need.

This alternate server might receive email for a subdomain, e.g. <user@csv.mydomain.com> and pass it on directly to the regular user accounts if the SMTP session passes all CSV-checking and the reputation service shows a good reputation; but otherwise return an immediate error. Although this is certainly less convenient for senders than implementing CSV-checking on your main server, it does give motivated senders something they can do to immediately resolve the problem.

The CSV-checking option will inform you whether a particular SMTP session is coming from a sending SMTP client which is both authenticated and authorized by the domain listed in the EHLO string. Initially, most SMTP sessions will not pass both of these tests, and you will need to continue the anti-spam measures you currently use. But for sessions which do pass both of these tests, you can substitute a single reputation check for all the blacklist tests you do now. (It's not unusual to find yourself checking a dozen different blacklists while evaluating each individual email.)

The reward for the (initially extra) CSV-checking is a reduction in the costs associated with wrongly identifying legitimate emails as spam. It is almost impossible to avoid the use of blacklists which continue to list IP addresses as spam-producing long after the actual problem has been resolved; and the actual sending domain usually has no idea how to speed the process of removing their IP address from such an IP-based blacklist.

CSV-checking, on the other hand, documents a strong association with a domain-name instead of an IP address; and enables that domain to publish records pointing to accreditation services which will vouch for its policies and practices. Every reputation service you could choose will have access to these records, and can automatically upgrade its reports within minutes of the resolution of any problem which may occur. (Likewise, it could downgrade a report within minutes of the first sign of a problem.)

The net effect is that you can reliably trust any reputation service which checks for accreditation pointers to inform you whether you need to do all the blacklist checking you are currently doing; and you can avoid many of the problems which would arise from incorrectly evaluating legitimate email as spam.

(We further recommend that you pass on a pointer to explanations of CSV to anyone who reports the false identification of spam for a session which didn't pass CSV-testing. In most cases, CSV records could be published in a matter of minutes, and "permanently" resolve the issue in a reasonably short time.)