Network Working Group
Request for Comments: 5016
Category: Informational

M. Thomas
Cisco Systems

October 2007

Requirements for a DomainKeys Identified Mail (DKIM) Signing Practices Protocol

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

DomainKeys Identified Mail (DKIM) provides a cryptographic mechanism for domains to assert responsibility for the messages they handle. A related mechanism will allow an administrator to publish various statements about their DKIM signing practices. This document defines requirements for this mechanism, distinguishing between those that must be satisfied (MUST), and those that are highly desirable (SHOULD).

Table of Contents

1. Introduction 2
2. Definitions and Requirements Language 3
3. SSP Problem Scenarios 4
   3.1. Problem Scenario 1: Is All Mail Signed with DKIM? 4
   3.2. Problem Scenario 2: Illegitimate Domain Name Use 5
4. SSP Deployment Considerations 6
   4.1. Deployment Consideration 1: Outsourced Signing 6
   4.2. Deployment Consideration 2: Subdomain Coverage 6
   4.3. Deployment Consideration 3: Resent Original Mail 7
   4.4. Deployment Consideration 4: Incremental Deployment of Signing 7
   4.5. Deployment Consideration 5: Performance and Caching 8
   4.6. Deployment Consideration 6: Human Legibility of Practices ..8
   4.7. Deployment Consideration 7: Extensibility 8
   4.8. Deployment Consideration 8: Security 8
5. Requirements 9
   5.1. Discovery Requirements 9
   5.2. SSP Transport Requirements 10
   5.3. Practice and Expectation Requirements 10
   5.4. Extensibility and Forward Compatibility Requirements 13
6. Requirements for SSP Security 13
7. Security Considerations 13
8. Acknowledgments 13
9. References 14
   9.1. Normative References 14

1. Introduction

DomainKeys Identified Mail [RFC4871] defines a message level signing and verification mechanism for email. While a DKIM signed message speaks for itself, there is ambiguity if a message doesn’t have a valid first party signature (i.e., on behalf of the [RFC2822].From address): is this to be expected or not? For email, this is an especially difficult problem since there is no expectation of a priori knowledge of a sending domain’s practices. This ambiguity can be used to mount a bid down attack that is inherent with systems like email that allow optional authentication: if a receiver doesn’t know otherwise, it should not assume that the lack of a valid signature is exceptional without other information. Thus, an attacker can take advantage of the ambiguity and simply not sign messages. If a protocol could be developed for a domain to publish its DKIM signing practices, a message verifier could take that into account when it receives an unsigned piece of email.

This document defines the requirements for a mechanism that permits the publication of Sender Signing Practices (SSP). The document is organized into two main sections: first, a Problem and Deployment Scenario section that describes the problems that SSP is intended to address as well as the deployment issues surrounding the base problems, and the second section is the Requirements that arise because of those scenarios.

2. Definitions and Requirements Language

- Domain Holder: the entity that controls the contents of the DNS subtree starting at the domain, either directly or by delegation via NS records it controls.

- First Party Address: for DKIM, a first party address is defined to be the [RFC2822].From address in the message header; a first party address is also known as an Author address.

- First Party Signature: a first party signature is a valid signature where the signing identity (the d= tag or the more specific identity i= tag) matches the first party address. "Matches" in this context is defined in [RFC4871].

- Third Party Signature: a third party signature is a valid signature that does not qualify as a first party signature. Note that a DKIM third party signature is not required to correspond to a header field address such as the contents of Sender or List-Id, etc.

- Practice: a statement according to the [RFC2822].From domain holder of externally verifiable behavior in the email messages it sends.

- Expectation: an expectation combines with a practice to convey what the domain holder considers the likely survivability of the practice for a receiver, in particular receivers that may be more than one SMTP hop away.

- DKIM Signing Complete: a practice where the domain holder asserts that all legitimate mail will be sent with a valid first party signature.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

3. SSP Problem Scenarios

The email world is a diverse place with many deployment considerations. This section outlines expected usage scenarios where DKIM signing/verifying will take place, and how a new protocol might help to clarify the relevance of DKIM-signed mail.

3.1. Problem Scenario 1: Is All Mail Signed with DKIM?

After auditing their outgoing mail and deploying DKIM signing for all of their legitimate outgoing mail, a domain could be said to be DKIM signing complete. That is, the domain has, to the best of its ability, ensured that all legitimate mail purporting to have come from that domain contains a valid DKIM signature.

A receiver in the general case doesn’t know what the practices are for a given domain. Thus, the receiver is at a disadvantage in not knowing whether or not it should expect all mail to be signed from a given domain. This knowledge gap leads to a trivially exploitable bid-down attack where the attacker merely sends unsigned mail; since the receiver doesn’t know the practices of the signing domain, it cannot treat the message any more harshly for lack of a valid signature.

An information service that allows a receiver to query for the practices and expectations of the first party domain when no valid first party signature is found could be useful in closing this gap. A receiver could use this information to treat such questionable mail with varying degrees of prejudice.

Note that, for the foreseeable future, unrestricted use patterns of mail (e.g., where users may be members of mailing lists, etc.) will likely suffer occasional, non-malicious signature failure in transit. While probably not a large percentage of total traffic, this kind of breakage may be a significant concern for those usage patterns. This scenario defines where the sender cannot set any expectation as to whether an individual message will arrive intact.

Even without that expectation, a receiver may be able to take advantage of the knowledge that the domain’s practice is to sign all mail and bias its filters against unsigned or damaged in transit mail. This information should not be expected to be used in a binary yes/no fashion, but instead as a data point among others in a filtering system.

The following exchange illustrates problem scenario 1.

1. Mail with a [RFC2822].From domain Alice is sent to domain Bob with a missing or broken DKIM first party signature from Alice.

2. Domain Bob would like to know whether that is an expected state of affairs.

3. Domain Alice provides information that it signs all outgoing mail, but places no expectation on whether it will arrive with an intact first party signature.

4. Domain Bob could use this information to bias its filters to examine the message with some suspicion.

3.2. Problem Scenario 2: Illegitimate Domain Name Use

A class of mail typified by transactional mail from high-value domains is currently the target of phishing attacks. In particular, many phishing scams forge the [RFC2822].From address in addition to spoofing much of the content to trick unsuspecting users into revealing sensitive information. Domain holders sending this mail would like the ability to give an enhanced guarantee that mail sent with their domain name should always arrive with the proof that the domain holder consented to its transmission. That is, the message should contain a valid first party signature as defined above.

From a receiver’s standpoint, knowing that a domain not only signs all of its mail, but places a very high value on the receipt of a valid first party signature from that domain is helpful. Hence, a receiver knows that the sending domain signs all its mail, and that the sending domain considers mail from this domain particularly sensitive in some sense, and is asking the receiver to be more suspicious than it otherwise might be of a broken or missing first-party signature. A receiver with the knowledge of the sender’s expectations in hand might choose to process messages not conforming to the published practices in a special manner. Note that the ability to state an enhanced guarantee of a valid signature means that senders should expect mail that traverses modifying intermediaries (e.g., mailing lists, etc.) will likely be quarantined or deleted; thus, this scenario is more narrow than problem scenario 1.

> Informative Note: a receiving filter may choose to treat scenario 2 much more harshly than scenario 1; where scenario 1 looks odd, scenario 2 looks like something is very wrong.

1. Mail with a [RFC2822].From domain Alice is sent to domain Bob with a missing or broken first party DKIM signature from domain Alice.

2. Domain Bob would like to know whether that is an expected state of affairs.

3. Domain Alice provides information that it signs all outgoing mail, and furthermore places an expectation that it should arrive with an intact first party signature, and that the receiver should be much more wary if it does not.

4. Domain Bob could use this information to bias its filters such that it examines the message with great suspicion.

4. SSP Deployment Considerations

Given the problems enumerated above for which we’d like SSP to provide information to recipients, there are a number of scenarios that are not related to the problems that are to be solved, per se, but the actual mechanics of implementing/deploying the information service that SSP would provide.

4.1. Deployment Consideration 1: Outsourced Signing

Many domains do not run their own mail infrastructure, or may outsource parts of it to third parties. It is desirable for a domain holder to have the ability to delegate to other entities the ability to sign for the domain holder. One obvious use scenario is a domain holder from a small domain that needs to have the ability for their outgoing ISP to sign all of their mail on behalf of the domain holder. Other use scenarios include outsourced bulk mail for marketing campaigns, as well as outsourcing various business functions, such as insurance benefits, etc.

4.2. Deployment Consideration 2: Subdomain Coverage

An SSP client will perform lookups on incoming mail streams to provide the information as proposed in the problem scenarios. The domain part of the first address of the [RFC2822].From will form the basis to fetch the published information. A trivial attack to circumvent finding the published information can be mounted by simply using a subdomain of the parent domain that doesn’t have published information. This attack is called the subdomain attack: that is, a domain wants to not only publish a policy for a given DNS label it controls, but it would also like to protect all subdomains of that label as well. If this characteristic is not met, an attacker would need only create a possibly fictitious subdomain that was not covered by the SSP’s information service. Thus, it would be advantageous for SSP to not only cover a given domain, but all subdomains of that domain as well.

4.3. Deployment Consideration 3: Resent Original Mail

Resent mail is a common occurrence in many scenarios in the email world of today. For example, domain Alice sends a DKIM-signed message with a published practice of signing all messages to domain Bob’s mailing list. Bob, being a good net citizen, wants to be able to take his part of the responsibility of the message in question, so he DKIM signs the message, perhaps corresponding to the Sender address.

Note that this scenario is completely orthogonal to whether Alice’s signature survived Bob’s mailing list: Bob merely wants to assert his part in the chain of accountability for the benefit of the ultimate receivers. It would be useful for this practice to be encouraged as it gives a more accurate view of who handled the message. It also has the side benefit that remailers that break DKIM first party signatures can be potentially assessed by the receiver based on the receiver’s opinion of the signing domains that actually survived.

4.4. Deployment Consideration 4: Incremental Deployment of Signing

As a practical matter, it may be difficult for a domain to roll out DKIM signing such that they can publish the DKIM Signing Complete practice given the complexities of the user population, the outsourced vendors sending on its behalf, etc. This leaves open an exploit that high-value mail, such as in Problem Scenario 2, must be classified to the least common denominator of the published practices. It would be desirable to allow a domain holder to publish a list of exceptions that would have a more restrictive practices statement. NB: this consideration has been deemed met by the mechanisms provided by the base DKIM signing mechanism; it is merely documented here as having been an issue.

For example, bigbank.example.com might be ready to say that statements@bigbank.example.com is always signed, but the rest of the domain, say, is not. Another situation is that the practices of some address local parts in a given domain are not the same as practices of other local parts. Using the same example of statements@bigbank.example.com being a transactional kind of email that would like to publish very strong practices, mixed in with the rest of the user population local parts, which may go through mailing lists, etc., for which a less strong statement is appropriate.

It should be said that DKIM, through the use of subdomains, can already support this kind of differentiation. That is, in order to publish a strong practice, one only has to segregate those cases into different subdomains. For example: accounts.bigbank.example.com would publish constrained practices, while corporateusers.bigbank.example.com might publish more permissive practices.

4.5. Deployment Consideration 5: Performance and Caching

Email service provides an any-any mesh of potential connections: all that is required is the publication of an MX record and an SMTP listener to receive mail. Thus, the use of SSP is likely to fall into two main scenarios, the first of which are large, well-known domains that are in constant contact with one another. In this case, caching of records is essential for performance, including the caching of the non-existence of records (i.e., negative caching).

The second main scenario is when a domain exchanges mail with a much smaller volume domain. This scenario can be both perfectly normal as with the case of vanity domains, and unfortunately, a vector for those sending mail for anti-social reasons. In this case, we’d like the message exchange burden to SSP querier to be low, since many of the lookups will not provide a useful answer. Likewise, it would be advantageous to have upstream caching here as well so that, say, a mailing list exploder on a small domain does not result in an explosion of queries back at the root and authoritative server for the small domain.

4.6. Deployment Consideration 6: Human Legibility of Practices

While SSP records are likely to be primarily consumed by an automaton, for the foreseeable future they are also likely to be inspected by hand. It would be nice to have the practices stated in a fashion that is also intuitive to the human inspectors.

4.7. Deployment Consideration 7: Extensibility

While this document pertains only to requirements surrounding DKIM signing practices, it would be beneficial for the protocol to be able to extend to other protocols.

4.8. Deployment Consideration 8: Security

SSP must be able to withstand life in a hostile, open-Internet environment. These include DoS attacks, and especially DoS attacks that leverage themselves through amplification inherent in the protocol. In addition, while a useful protocol may be built without strong source authentication provided by the information service, a path to strong source authentication should be provided by the protocol, or underlying protocols.

5. Requirements

This section defines the requirements for SSP. As with most requirements documents, these requirements define the MINIMUM requirements that a candidate protocol must provide. It should also be noted that SSP must fulfill all of the requirements.

5.1. Discovery Requirements

Receivers need a means of obtaining information about a sender’s DKIM practices. This requires a means of discovering where the information is and what it contains.

1. The author is the first-party sender of a message, as specified in the [RFC2822].From field. SSP’s information is associated with the author’s domain name, and is published subordinate to that domain name.

2. In order to limit the cost of its use, any query service supplying SSP’s information MUST provide a definitive response within a small, deterministic number of message exchanges under normal operational conditions.

    Informative Note: this, for all intents and purposes is a prohibition on anything that might produce loops or result in extended delays and overhead; also though "deterministic" doesn’t specify how many exchanges, the expectation is "few".

    Refs: Deployment Considerations, Sections 4.2 and 4.5.

3. SSP’s publishing mechanism MUST be defined such that it does not lead to multiple resource records of the same type for different protocols residing at the same location.

    Informative note: an example is multiple resource record of the same type within a common DNS leaf. Hence, uniquely defined leaf names or uniquely defined resource record types will ensure unambiguous referencing.

    Refs: Deployment Consideration, Section 4.2.

4. SSP retrieval SHOULD provide coverage for not only a given domain but all of its subdomains as well. It is recognized that there is some reasonable doubt about the feasibility of a widely accepted solution to this requirement. If the working group does not achieve rough consensus on a solution, it MUST document the relevant security considerations in the protocol specification.

   Refs: Deployment Considerations, Sections 4.2 and 4.5.

5.2. SSP Transport Requirements

The publication and query mechanism will operate as an internet-based message exchange. There are multiple requirements for this lower-layer service:

1. The exchange SHOULD have existing widespread deployment of the transport layer, especially if riding on top of a true transport layer (e.g., TCP, UDP).

   Refs: Deployment Considerations, Sections 4.5 and 4.7.

2. The query/response in terms of latency time and the number of messages involved MUST be low (less than three message exchanges not counting retransmissions or other exceptional conditions).

   Refs: Deployment Consideration, Section 4.5.

3. If the infrastructure doesn’t provide caching (a la DNS), the records retrieved MUST provide initiators the ability to maintain their own cache. The existing caching infrastructure is, however, highly desirable.

   Refs: Deployment Consideration, Section 4.5.

4. Multiple geographically and topologically diverse servers MUST be supported for high availability.

   Refs: Deployment Considerations, Sections 4.5 and 4.7.

5.3. Practice and Expectation Requirements

As stated in the definitions section, a practice is a statement according to the [RFC2822].From domain holder of externally verifiable behavior in the email messages it sends. As an example, a practice might be defined such that all email messages will contain a DKIM signature corresponding to the [RFC2822].From address. Since there is a possibility of alteration between what a sender sends and a receiver examines, an expectation combines with a practice to convey what the [RFC2822].From domain considers the likely outcome of the survivability of the practice at a receiver. For example, a practice that a valid DKIM for the [RFC2822].From address is present when it is sent from the domain, and an expectation that it will remain present and valid for all receivers whether topologically adjacent or not.

1. SSP MUST be able to make practices and expectation assertions about the domain part of a [RFC2822].From address in the context of DKIM. SSP will not make assertions about other addresses for DKIM at this time.

   Refs: Problem Scenarios 1 and 2, Sections 3.1 and 3.2.

2. SSP MUST provide a concise linkage between the [RFC2822].From and the identity in the DKIM i= tag, or its default if it is missing in the signature. That is, SSP MUST precisely define the semantics of what qualifies as a first party signature.

   Refs: Problem Scenarios 1 and 2, Sections 3.1 and 3.2.

3. SSP MUST be able to publish a practice that the domain’s signing behavior is "DKIM Signing Complete". That is, all messages were transmitted with a valid first party signature.

   Refs: Problem Scenario 1, Section 3.1.

4. SSP MUST be able to publish an expectation that a verifiable first party DKIM signature should be expected on receipt of a message.

   Refs: Problem Scenario 2, Section 3.2.

5. Practices and expectations MUST be presented in SSP syntax using as intuitive a descriptor as possible. For example, p=? would be better represented as p=unknown.

   Refs: Deployment Consideration, Section 4.6.

6. Because DKIM uses DNS to store selectors, there is always the ability for a domain holder to delegate all or parts of the _domainkey subdomain to an affiliated party of the domain holder’s choosing. That is, the domain holder may set an NS record for _domainkey.example.com to delegate to an email provider who manages the entire namespace. There is also the ability for the domain holder to partition its namespace into subdomains to further constrain third parties. For example, a domain holder could delegate only _domainkey.benefits.example.com

to a third party to constrain the third party to only be able to produce valid signatures in the benefits.example.com subdomain. Last, a domain holder can even use CNAME’s to delegate individual leaf nodes. Given the above considerations, SSP need not invent a different means of allowing affiliated parties to sign on a domain’s behalf at this time.

Refs: Deployment Consideration, Section 4.4.

7. Practices and expectations MUST be presented as an information service from the signing domain to be consumed as an added factor to the receiver’s local policy. In particular, a practice or expectation MUST NOT mandate any disposition stance on the receiver.

   Refs: Problem Scenarios 1 and 2, Sections 3.1 and 3.2.

8. There is no requirement that SSP publish practices of any/all third parties that MUST NOT sign on the domain holder’s behalf. This should be considered out of scope.

   INFORMATIVE NOTE: this is essentially saying that the protocol doesn’t have to concern itself with being a blacklist repository.

   Refs: Problem Scenarios 1 and 2, Sections 3.1 and 3.2.

9. SSP MUST NOT be required to be invoked if a valid first party signature is found.

   Refs: Deployment Consideration, Section 4.2.

10. SSP MUST NOT provide a mechanism that impugns the existence of non-first party signatures in a message. A corollary of this requirement is that the protocol MUST NOT link practices of first party signers with the practices of third party signers.

    INFORMATIVE NOTE: the main thrust of this requirement is that practices should only be published for that which the publisher has control, and should not meddle in what is ultimately the local policy of the receiver.

    Refs: Deployment Consideration, Section 4.3.

5.4. Extensibility and Forward Compatibility Requirements

1. SSP MUST NOT extend to any protocol other than DKIM for email at this time. SSP SHOULD be extensible for protocols other than DKIM.

   Refs: Deployment Consideration, Section 4.7.

2. SSP MUST be able to add new practices and expectations within the existing discovery/transport/practices in a backward compatible fashion.

   Refs: Deployment Consideration, Section 4.7.

6. Requirements for SSP Security

1. SSP for a high-value domain is potentially a high-value DoS target, especially since the unavailability of SSP’s record could make unsigned messages less suspicious.

2. SSP MUST NOT make highly leveraged amplification or make-work attacks possible. In particular, the work and message exchanges involved MUST be order of a constant.

   Refs: Deployment Consideration, Section 4.8.

3. SSP MUST have the ability for a domain holder to provide SSP’s data such that a receiver can determine that it is authentically from the domain holder with a large degree of certainty. SSP may provide means that provide less certainty in trade off for ease of deployment.

   Refs: Deployment Consideration, Section 4.8.

7. Security Considerations

This document defines requirements for a new protocol and the security related requirements are defined above. Since it is expected that the new protocol will use the DNS as a basis for the published SSP information, most if not all of the threats described in [RFC4686] will be applicable.

8. Acknowledgments

Dave Crocker and Jim Fenton provided substantial review of this document. Thanks also to Vijay Gurbani and David Harrington for their helpful last call reviews.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2822] Resnick, P., Ed., "Internet Message Format", RFC 2822, April 2001.

[RFC4686] Fenton, J., "Analysis of Threats Motivating DomainKeys Identified Mail (DKIM)", RFC 4686, September 2006.

[RFC4871] Allman, E., Callas, J., Delany, M., Libbey, M., Fenton, J., and M. Thomas, "DomainKeys Identified Mail (DKIM) Signatures", RFC 4871, May 2007.

Author's Address

Michael Thomas
Cisco Systems
606 Sanchez St
San Francisco, California 94114
USA

Phone: +1-408-525-5386
Fax: +1-408-525-5386
EMail: mat@cisco.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.