Workgroup: Independent Stream
Internet-Draft:
draft-chuang-dkim-replay-problem-00
Published: 21 October 2022
Intended Status: Informational
Expires: 24 April 2023
A W. Chuang A. Robinson B. Gondwana
uGoogle, Inc. Google, Inc. Fastmail Pty Ltd
t
h
o
r
s
:

DKIM Replay Problem Statement and Scenarios

Abstract

DKIM [RFC6376] is an IETF standard for the cryptographic protocol to sign and authenticate email at the domain level and protect the integrity of messages during transit. In particular this enables DKIM to be able authenticate email through email forwarding. Section 8.6 of [RFC6376] defines a vulnerability called DKIM Replay as a spam message sent through a SMTP MTA DKIM signer, that then is sent to many more recipients, leveraging the reputation of the signer. This document defines the damage this causes to email delivery and interoperability, and the impacted mail flows. Part of the reason why this is such a difficult problem is that receivers have a hard time differentiating between legitimate forwarding flows and DKIM replay.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 24 April 2023.

Table of Contents

1. Introduction

The SMTP protocol was originally designed without authentication, and to allow indirect email flow.

In the years since, spam has become a major problem. In order to combat spam and help detect whether email has come from the source it claimed, various authentication mechanisms have been created. SPF, DKIM, ARC.

SPF lists the IP address ranges from which email may be sent from a domain. This is a very strong signal of correctness, however it does not allow for indirect email flows, such as email forwarding services or mailing lists.

DKIM adds a cryptographic signature over the contents and some of the headers of the email message, created with a private key only known to the sending service, and with the public key published in DNS to allow receiving systems to validate that the content hasn't been changed in transit. It is more robust in the face of indirect email flows, so long as they don't modify the content.

SRS rewriting allows for email forwarding in the face strict SPF or modifications that break DKIM (i.e. mailing lists which add footers with unsubscribe links or links to permanent archive versions of the message)

1.1. The problem

Since most domains (including bad actors) have SPF or DKIM records for their domain, receiving systems track reputation on SPF and DKIM (in addition to IP range) to help classify incoming emails.

In early 2022, a particular type of DKIM replay attack became common. Attackers would create or compromise an account at a site with a high reputation - which allowed them to generate an email with the content that they wanted to broadcast widely. They would then send an email with this content to an external account under their control. This single message would be delivered to the attacker's mailbox, at which point they would have an email with a valid DKIM signature, and with the high reputation of the sending site associated.

Due to the ability to BCC emails (which means that the ENVELOPE-TO does not match the To or Cc headers signed by the DKIM Signature), this message could then be replayed to arbitrary other mailboxes, sometimes thousands or millions of mailboxes would all have the same message delivered to them.

ARC was created as an experiment in 2018, to provide additional authentication at every hop in the email flow. It has the same "replay" issue however, as the DKIM signature is still present and an attacker can remove all ARC headers from the email before replaying it. There are no external indicators that ARC is expected to be applied, so a recipient can't tell the difference between a message that used to have ARC and one that never had ARC.

1.2. Glossary

This document borrows terminology from [RFC5321] and [RFC5598].

1.2.1. Administrative Management Domain (ADMD)

Defined in [RFC5598], this defines an actor who administrative policies are the same. An ADMD can be a Mailbox Provider, Enterprise Service Provider, and Internet Service Provider.

1.2.2. Originating Sender

This is a service that originates the email meaning is the first SMTP [RFC5321] MTA sender. It also is described as the first ADMD. Likely this was in response to a user submitting an email, but also potentially from automated systems. Regardless, it's up to the service operator to authenticate the user/system that is submitting mail to them, using whatever mechanisms make sense for them. The originating sender typically DKIM signs the message.

1.2.3. Outbound filtering service

This is a service that a sending service such as the originator can route some or all of their mail through, instead of sending directly to the recipient's configured server. This should typically be transparent to mail flows. It may provide spam or data loss protection services, and may modify the message.

1.2.4. Destination Receiver

This is defined as the service that receives the email and does not elect to forward it elsewhere. It is the last SMTP MTA receiver, and ADMD. For simplicity, this excludes things like autoforwarding and mailing lists, as these behave differently with respect to authentication signals.

1.2.5. Inbound filtering service

This is a service that is configured to accept inbound mail for a domain by being listed in the recipient's MX record. This service would also be configured with enough information to forward email along, either to a subsequent filtering service or to the destination service. These services are often used to do things like security scans and spam filtering, but may also modify messages before forwarding them.

1.2.6. Mailing list service

This is a service that accepts mail sent to an address, and sends copies of it to a configured list of other recipients. This service may change the message prior to sending to list members, for example changing the subject or adding a message footer. Mailing list message may be DKIM signed by the originating sender or sometimes when the message is modified by the mailing list.

1.2.7. Forwarder

This is a service that may be a standard destination service, but additionally performs forwarding of (some/all) mail for an address to another recipient. Split out explicitly due to challenges in authentication for flows involving this type of activity, and because flows involving the same service as both a forwarder and as a destination service can be thought of as two independent (though related) mail flows for the purposes of authentication signal propagation.

Forwarders are very similar to mailing lists, but have a few differences that may be interesting:

- Different headers (List-XYZ, X-Forwarded-By, etc.)
- Different features: mailing lists often change subject, add footers, etc. but simple forwarders such as auto-forwarders often resend the message as-is
- Envelope from address may not be handled the same. Lists usually send as themselves, forwarders not so much

1.2.8. ESP Bulk Sender

This is a forwarding service similar to mailing list service in that they send to a list of recipients. However they are commercial in nature and they often disclose and DKIM sign for the recipient in the headers. ESP Bulk Senders also typically do not modify the message body.

2. Mail Flow Scenarios

2.1. Direct delivery

This is the "standard" mail flow. In this case, the email is delivered directly to the MX servers for recipient.com, which stores the message for consumption by "someone@recipient.com" through their mail client. \

In this flow:

- SPF is likely easy and safe, assuming IP ranges allowed by the policy are stable and aren't shared (ex. Cloud-based mail providers)
- DKIM is effective, as there is nothing in the mail flow that can break the DKIM signature
- DMARC is often aligned and a strong DMARC policy is an option for some senders
- ARC isn't really necessary because there is only a single SMTP transaction, and all authentication signals are available to that transaction

ESP Bulk Sender

ESP bulk senders will originate messages given a message body and a list of recipients that then become the RFC822 To recipient. A ESP Bulk Sender may have DKIM signatures of the brand and the ESP Bulk Sender, or one or the other. The ESP Bulk Sender will administratively apply policies to protect its reputation.

- SPF is likely easy and safe, assuming IP ranges allowed by the policy are stable and aren't shared (ex. Cloud-based mail providers)
- DKIM signatures should survive ESP Bulk Sender flow
- It may be feasible to apply strict DMARC policy assuming the brand aligns the RFC822 From with SPF or DKIM domains. However in practice, many smaller brands do not pass DMARC though there is nothing technical that prevents them from doing so.

2.2. Delivery through an outbound filtering service

In this case, the originating service decides to send the email to an endpoint other than the one specified in the recipient's MX record.

In this flow:

* SPF is likely possible, but may not be advisable. The filtering service should be able to provide an SPF policy (ex. by publishing an SPF policy that covers their IPs and instructing users to reference it in their policy). More potential for IP sharing, which weakens the authentication.

* DKIM is likely possible, but may break if the filter performs any modifications of the message. Such modifications could be done if the filtering service has the ability to resign for DKIM after modifications on behalf of the sending domain.

* DMARC may not work, depending on whether SPF/DKIM work reliably.

* ARC could help in the absence of DKIM resigning capabilities, and may actually improve the state of things by removing the requirement that domain owners grant such services permission to sign on their behalf.

2.3. Delivery through an inbound filtering service

In this case, the email is sent directly from the originating service, but the service listed in the recipient's MX record is not the destination service. The inbound filtering service that is listed in the MX record will forward directly to the destination service. The envelope recipient and the header recipient remains the same.

In this flow:

* The inbound filtering service can do the same authentication as a normal receiving service. There is nothing between the originating service and the filter that would break SPF/DKIM.

* DMARC can be applied at the inbound filter service in the same way that a destination service would apply it during direct delivery. SPF is available and there is no previous intermediary to break DKIM.

* The destination service would be unable to perform SPF authentication, because the required information (connecting IP) is only available during the SMTP transaction between the originating service and the inbound filtering service.

* The destination service may be unable to perform DKIM authentication, in cases where the filter mutates the message in a way that invalidates the signature.

* ARC could be interesting in this flow. The inbound filtering service could produce an ARC set describing the authentication results it computed, and the destination service could verify authentication results through that.

2.3.1. Mailing List

Messages are sent to a mailing list which expands the list of recipients and forwards the message to the envelope recipients. Note

that the envelope recipient changes, but header recipient usually remains the same. The message may be modified such as when the subject is updated or a footer added to the message body.

In this flow:

- The mailing list can do the same authentication as a normal receiving service.
- DMARC can be applied at the mailing list in the same way that a destination service would apply it during direct delivery. SPF is available and there is no previous intermediary to break DKIM.
- The destination service would be unable to perform SPF authentication, because the required information (connecting IP) is only available during the SMTP transaction between the originating service and the mailing list.
- The destination service may be unable to perform DKIM authentication, in cases where the mailing list mutates the message e.g. adding a footer, in a way that invalidates the signature.
- ARC could be interesting in this flow. The mailing list could produce an ARC set describing the authentication results it computed, and the destination service could verify authentication results through that.

2.3.2. Forwarder

Messages sent to a forwarder will have the envelope recipient updated to a different recipient at a different MX, and routed to that different MX. A forwarder may add headers but typically doesn't break the DKIM signature.

- Forwarder can do the same authentication as a normal receiving service.
- DMARC can be applied at the forwarder in the same way that a destination service would apply it during direct delivery. SPF is available and there is no previous intermediary to break DKIM.
- The destination service would be unable to perform SPF authentication, because the required information (connecting IP) is only available during the SMTP transaction between the originating service and the forwarder, and it's not scalable (or secure) to have forwarders controlled by the recipient listed in a sender's SPF policy.
- DKIM signature should survive forwarding, unless the forwarder also performs modifications to the message's content or any signed headers.
- DMARC at the destination only works if the message was DKIM-signed, because SPF won't work.

2.3.3. Forwarder (rewrites envelope from)

Similar to the normal forwarding case, but also updates the envelope from (MAIL FROM) address to be something related to the forwarder, instead of simply reusing the envelope from address from the initial SMTP conversation.

- Forwarder can do the same authentication as a normal receiving service.
- DMARC can be applied at the forwarder in the same way that a destination service would apply it during direct delivery. SPF is available and there is no previous intermediary to break DKIM.
- The destination service can perform SPF checks, but the check would have no relationship to the originating domain because the envelope address was rewritten to something owned by the forwarder.
- DKIM signature should survive forwarding, unless the forwarder also performs modifications to the message's content or any signed headers.
- DMARC at the destination only works if the message was DKIM-signed. While SPF may pass, it won't be for an aligned identifier, so the result will not be considered for DMARC evaluation.

3. DKIM Replay

3.1. Scenario

A spammer will find a DKIM signer with a high reputation at a victim mailbox provider who the spammer intends to send spam messages to. The spammer sends a message with spam content through the DKIM signer to some spammer controlled account where they can obtain the signed message. This captured message is sometimes updated with additional headers such as To and Subject that do not damage the existing DKIM signature (such as leaving off those headers in the original message). That message is then sent at scale aka amplified to the victim domain. Because the signed message has a high reputation, the message with spam content gets through to the inbox. This is an example of a spam classification false negative.

The victim domain spam system starts to observe a large amount of spam, and reacts by dropping the reputation of the DKIM signer. Benign mail from the signer's domain starts to go to the spam folder. This is an example of a spam classification false positive.

3.2. Indistinguishable from Forwarding Flows

DKIM replay attacks always have a valid DKIM signature as this is needed to take advantage of the good reputation of the DKIM signer. After that the spammer utilizes whatever other authentication needed to get past the spam filter. Typically such a message has a DMARC aligned RFC822 FROM. They often have a valid SPF though not aligned with the RFC822 FROM. This DKIM pass and aligned and potential SPF pass but not aligned looks like an indirect forwarding flow message.

Legitimate indirect flows often look like this. For unmodified messages, the often will have a valid DKIM signature that is aligned but will fail SPF or will be unaligned. Example benign indirect flows are outbound and inbound gateway, mailing lists, forwarders, and ESP bulk senders.

4. Proposed Solution Space

Here are some potential solutions to the problem, and their pros and cons

* Include the ENVELOPE-TO in the signature somehow.
  - This avoids replay to destination addresses not anticipated by the DKIM signer.
  - Many indirect email flows are impacted just as badly as they are by SPF, since forwarding involves rewriting the ENVELOPE-TO.
  - This will detect DKIM replays, but not distinguish them from all other forwarding.
* Cache known DKIM signatures.
  - Since the exact same signature is being replayed over and over, this allows a receiving site with many mailboxes to detect whether a message is part of a DKIM replay set, and suppress it.
  - Mailing list traffic, exploder aliases, or regular BCCs will also show up as duplicates, so this is very much a heuristic guess of whether the amount of duplication is expected or not. This may lead to spam filtering false positives.
* Strip DKIM signatures on mailbox delivery.
  - Messages delivered to a mailbox are not able to be replayed any more.
  - May require two signatures, one that's kept and one that's removed, with both required over the wire.
  - Doesn't help against an evil mailbox server, as the attacker would just avoid stripping the headers. \
* Add a per-hop signature, specifying the destination domain for the next hop
  - Messages with this kind of signature cannot be replayed down a different pathway, since the destination won't match.
  - Requires every site along the path to support this spec, so it will need to detect whether the next stop is making a commitment to follow the spec.

-If email goes outside of sites with this spec (without disclosure), traversing a naive forwarder remains indistinguishable from replay.

5. Other

6. Normative References

[RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, DOI 10.17487/RFC5321, October 2008, <https://www.rfc-editor.org/info/rfc5321>.

[RFC5598] Crocker, D., "Internet Mail Architecture", RFC 5598, DOI 10.17487/RFC5598, July 2009, <https://www.rfc-editor.org/info/rfc5598>.

[RFC6376] Crocker, D., Ed., Hansen, T., Ed., and M. Kucherawy, Ed., "DomainKeys Identified Mail (DKIM) Signatures", STD 76, RFC 6376, DOI 10.17487/RFC6376, September 2011, <https://www.rfc-editor.org/info/rfc6376>.

Appendix A. Acknowledgments

Thanks goes to Emanuel Schorsch and Evan Burke for their advice.

Authors' Addresses

Weihaw Chuang
Google, Inc.

Email: weihaw@google.com

Allen Robinson
Google, Inc.

Email: arobins@google.com

Bron Gondwana
Fastmail Pty Ltd

Email: brong@fastmailteam.com