independent
Internet-Draft
Intended status: Informational
Expires: April 18, 2016

. OAR-DEV Group
OAR-DEV Group
S. Jones
DMARC.org
J. Rae-Grant
Google
T. Adams
Paypal
K. Andersen, Ed.
LinkedIn

October 16, 2015

Recommended Usage of the Authenticated Received Chain (ARC)
draft-jones-arc-usage-00

Abstract

The Authentication Results Chain (ARC) provides a means to preserve email authentication results and verify the identity of email message handlers, each of which participates by inserting certain headers before passing the message on. But the specification does not indicate how intermediaries and receivers should interpret or utilize ARC. This document will provide guidance in these areas.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

[ARC] is intended to be used primarily by intermediaries, or message handlers - those parties who may forward or resend messages, with or without alterations, such that they will no longer pass the SPF, DKIM, and/or [RFC7489] authentication mechanisms. In such cases ARC may provide the final message recipient with useful information about the original sender.

2. How does ARC work?

Consider a mailing list as an example, where the message submitter's domain publishes a DMARC policy other than "p=none". The message is received, a prefix is added to the RFC5322.Subject header, some text is appended to the message body, and the message is sent to list members with the original RFC5322.From address intact. In this case SPF may pass because the mailing list operator uses their own domain in the RFC5321.MailFrom header, but this domain will not match the RFC5322.From address, thus the DMARC SPF result cannot be a "pass." Any DKIM signature from the message submitter's domain will be broken as the message body has been altered (and if included in the signature, the RFC5322.Subject header). Again, the DMARC DKIM result cannot be a "pass." And if the mailing list operator inserted an Authentication-Results: header it was most likely stripped and/or replaced by the next message receiver.

If the mailing list implemented ARC, it would record the contents of the Authentication-Results: header in the ARC-Authentication-Results: header. It would then create an an ARC-Message-Signature: header, which includes a cryptographic signature of the message itself, and then an ARC-Seal: header, which includes a cryptographic signature of a few key message headers - including the other ARC headers.

Any subsequent system participating in ARC that was not performing final delivery of the message within its ADMD boundaries would also generate and insert ARC headers whose signatures cover all ARC headers inserted into the message by previous message handlers. Thus the information from any previous ARC participants, including the ARC-Authentication-Results: header from the mailing list operator, would be signed at each ADMD that handled the message.

When the message reaches the final receiving system, the SPF and DKIM results will not satisfy the DMARC policy for the message author's domain. However if the receiving system implements ARC then it can check for and validate an ARC chain and verify that the contents of the ARC-Authentication-Results: header were conveyed intact from the mailing list operator. At that point the receiving system might choose to use those authentication results in the decision of whether or not to deliver the message, even though it failed to pass the usual authentication checks.

3. Guidance for Receivers/Validators

3.1. What is the significance of an intact ARC chain?

An intact ARC chain conveys authentication results like SPF and DKIM as observed by the first ARC participant. In cases where the message no longer produces passing results for DKIM, SPF, or DMARC but an intact ARC chain is present, the message receiver may choose to use the contents of the ARC-Authentication-Results: header in determining how to handle the message.

3.2. What exactly is an "intact" ARC chain?

Note that not all ADMDs will implement ARC, and receivers will see messages where one or more non-participating ADMDs handled a message before, after, or in between participating ADMDs.

An intact ARC chain is one where the ARC headers that are present can be validated, and in particular the ARC-Message-Signature: header from the last ARC participant can still be validated. This shows that, whether another ADMD handled the message after the last ARC participant or not, the portions of the message covered by that signature were not altered. If any non-participating ADMDs handled the message between ARC intermediaries but did not alter the message in a way that invalidated the most recent ARC-Message-Signature: present at that time, the chain would still be considered intact by the next ARC participant, and recorded as such in the ARC-Seal: header they insert.

Message receivers may make local policy decisions about whether to use the contents of the ARC-Authentication-Results: header in cases where a message no longer passes DKIM, DMARC, and/or SPF checks. Whether an ARC chain is intact can be used to inform that local policy decision.

So for example one message receiver may decide that, for messages with an intact ARC chain where a DMARC evaluation does not pass, but the ARC-Authentication-Results: header indicates a DKIM pass was reported that matches the domain in the RFC5322.From header, it will override a DMARC "p=reject" policy. Another message receiver may decide to do so for intact ARC chains where the ARC-Authentication-Results: header indicates an SPF pass. A third message receiver may use very different criteria, according to their requirements, while a fourth may choose not to take ARC information into account at all.

3.3. What is the significance of an invalid ("broken") ARC chain?

An ARC chain is not considered to be valid if the signatures in the ARC-Seal: headers cannot be verified. For example the remote server delivering the message to the local ADMD is not reflected in any ARC headers, perhaps because they have not implemented ARC, but they modified the message such that ARC and DKIM signatures already in the message were invalidated.

In such cases the ARC-Authentication-Results: header should not have any influence on the disposition of the message. For example, a message that fails under DMARC and has an invalid ARC chain would be subject to that DMARC policy, which may cause it to be quarantined or rejected.

3.4. What does the absence of an ARC chain in a message mean?

The absence of an ARC chain means nothing. ARC is intended to allow a participating message handler to preserve certain authentication results when a message is being forwarded and/or modified such that the final recipient can evaluate the source. If they are absent, there is nothing extra that ARC requires the final recipient to do.

3.5. What reasonable conclusions can you draw based upon seeing lots of mail with ARC chains?

With sufficient history, ARC can be used to augment DMARC authentication policy (i.e. a message could fail DMARC, but pass ARC and therefore could be considered as validly authenticated as reported by the first ARC participant).

If the validator does content analysis and reputation tracking, the ARC participants in a message can be credited or discredited for good or bad content. By analyzing different ARC chains involved in "bad" messages, a validator might identify malicious participating intermediaries.

With a valid chain and good reputations for all ARC participants, receivers may choose to apply a "local policy override" to the DMARC policy assertion for the domain authentication evaluation, depending on the ARC-Authentication-Results: header contents. Normal content analysis should never be skipped.

3.6. What if none of the intermediaries have been seen previously?

This has no impact on the operation of ARC, as ARC is not a reputation system. ARC conveys the results of other authentication mechanisms such that the participating message handlers can be positively identified. Final message recipients may or may not choose to examine these results when messages fail other authentication checks. They are more likely to override, say, a failing DMARC result in the presence of an intact ARC chain where the participating ARC message handlers have been observed to not convey "bad" content in the past, and the initial ARC participant indicates the message they received had passed authentication checks.

3.7. What about ARC chains where some intermediaries are known and others are not?

Validators may choose to build reputation models for ARC message handlers they have observed. Generally speaking it is more feasible to accrue positive reputation to intermediaries when they consistently send messages that are evaluated positively in terms of content and ARC chains. When messages are received with ARC chains that are not intact, it is very difficult identify which intermediaries may have manipulated the message or injected bad content.

3.8. What should message handlers do when they detect malicious content in messages where ARC is present?

Message handlers should do what they normally do when they detect malicious content in a message - hopefully that means quarantining or discarding the message. ARC information should never make malicious content acceptable.

In such cases it is difficult to determine where the malicious content may have been injected. What ARC can do in such cases is verify that a given intermediary or message handler did in fact

handle the message as indicated in the headers. In such cases a message recipient who maintains a reputation system about email senders may wish to incorporate this information as an additional factor in the score for the intermediaries and sender in question. However reputation systems are very complex, and usually unique to those organizations operating them, and therefore beyond the scope of this document.

3.9. What feedback does a sender or domain owner get about ARC when it is applied to their messages?

ARC itself does not include any mechanism for feedback or reporting. It does however recommend that message receiving systems that use ARC to augment their delivery decisions, who use DMARC and decide to deliver a message because of ARC information, should include a notation to that effect in their normal DMARC reports. These notations would be easily identifiable by report processors, so that senders and domain owners can see where ARC is being used to augment the deliverability of their messages.

3.10. What prevents a malicious actor from removing the ARC headers, altering the content, and creating a new ARC chain?

ARC does not prevent a malicious actor from doing this. Nor does it prevent a malicious actor from removing all but the first ADMD's ARC headers and altering the message, eliminating intervening participants from the ARC chain. Or similar variations.

A valid ARC chain does not provide any automatic benefit. With an intact ARC chain, the final message recipient may choose to use the contents of the ARC-Authentication-Results: header in determining how to handle the message. The decision to use the ARC-Authentication-Results: header is dependent on evaluation of those ARC intermediaries.

In the first case, the bad actor has succeeded in manipulating the message but they have attached a verifiable signature identifying themselves. While not an ideal situation, it is something they are already able to do without ARC involved, but now a strong link to the domain responsible for the manipulation is present.

Additionally in the second case it is possible some negative reputational impact might accrue to the first ARC participant left in place until more messages reveal the pattern of activity by the bad actor. But again, a bad actor can similarly manipulate a sequence of RFC5322.Received headers today without ARC, and with ARC that bad actor has verifiably identified themselves.

4. Guidance for Intermediaries

4.1. What is an Intermediary under ARC?

In the context of ARC, an Intermediary is typically an Administrative Management Domain [RFC5598] that is receiving a message, potentially manipulating or altering it, and then passing it on to another ADMD for delivery. Common examples of Intermediaries are mailing lists, alumni or professional email address providers that forward messages such as universities or professional organizations, et cetera.

4.2. What are the minimum requirements for an ARC Intermediary?

A participating ARC intermediary must validate the ARC chain on a message it receives, if one is present. It then attaches its own ARC seal and signature, including an indication if the chain failed to validate upon receipt.

4.2.1. More specifically a participating ARC intermediary must do the following:

1. Validate that the ARC chain, if one is already present in the message, is intact and well-formed.

2. Validate that the most recent sender matches the last entry in the ARC chain (if present).

3. Validate that the most recent sender's DKIM signature is attached, and matches the reference to it in the ARC chain (if present).

4. Generate a new ARC Signature and add it to the message according to the ARC specification.

5. Generate a new ARC Seal and add it to the message according to the ARC specification.

4.3. Should every MTA be an ARC participant?

Generally speaking, ARC is designed to operate at the ADMD level. When a message is first received by an ADMD, the traditional authentication results should be captured and preserved - this could be the common case of creating an Authentication-Results: header. But when it is determined that the message is being sent on outside of that ADMD, that is when the ADMD should add itself to the ARC chain - before sending the message outside of the ADMD.

Some organizations may operate multiple ADMDs, with more or less independence between them. While they should make a determination based on their specific circumstances, it may be useful and appropriate to have one or both ADMDs be ARC participants.

4.4. What should an intermediary do in the case of an invalid or "broken" ARC chain?

In general terms, a participating ARC intermediary will note that an ARC chain was present and invalid, or broken, when it attaches its own ARC seal and signature. However the fact that the ARC chain was invalid should have no impact on whether and how the message is delivered.

4.5. What should I do in the case where there is no ARC chain present in a message?

A participating ARC intermediary receiving a message with no ARC chain, and which will be delivered outside its ADMD, should start an ARC chain according to the ARC specification. This will include capturing the normal email authentication results for the intermediary (SPF, DKIM, DMARC, etc), which will be conveyed as part of the ARC chain.

4.6. How could ARC affect my reputation as an intermediary?

Message receivers often operate reputation systems, which build a behavioral profile of various message handlers and intermediaries. The presence or absence of ARC is yet another data point that may be used as an input to such reputation systems. Messages deemed to have good content may provide a positive signal for the intermediaries that handled it, while messages with bad content may provide a negative signal for the those intermediaries. Intact and valid ARC elements may amplify or attenuate such signals, depending on the circumstances.

Reputation systems are complex and usually specific to a given message receiver, and a meaningful discussion of such a broad topic is beyond the scope of this document.

4.7. What can I do to influence my reputation as an intermediary?

Today it is extremely simple for a malicious actor to construct a message that includes your identity as an intermediary, even though you never handled the message. It is possible that an intermediary implementing ARC on all traffic it handles might receive some reputational benefit by making it easier to detect when their involvement in conveying bad traffic has been "forged."

As mentioned previously reputation systems are very complex and usually specific to a given message receiver, and a meaningful discussion of such a broad topic is beyond the scope of this document.

5. Guidance for Originators

5.1. How can ARC impact my email?

Prior to ARC, certain DMARC policies on a domain would cause messages using those domains in the RFC5322.From field, and which pass through certain kinds of intermediaries (mailing lists, forwarding services), to fail authentication checks at the message receiver. As a result these messages might not be delivered to the intended recipient.

ARC seeks to provide these so-called "indirect mailflows" with a means to preserve email authentication results as seen by participating intermediaries. Message receivers may accept ARC results to supplement the information that DMARC provides, potentially deciding to deliver the message even though a DMARC check did not pass.

The net result for domain owners and senders is that ARC may allow messages routed through participating ARC intermediaries to be delivered, even though those messages would not have been delivered in the absence of ARC.

5.2. How can ARC impact my reputation as a message sender?

Message receivers often operate reputation systems, which build a behavioral profile of various message senders (and perhaps intermediaries). The presence or absence of ARC is yet another data point that may be used as an input to such reputation systems. Messages deemed to have good content may provide a positive signal for the sending domain and the intermediaries that handled it, while messages with bad content may provide a negative signal for the sending domain and the intermediaries that handled it. Intact and valid ARC elements may amplify or attenuate such signals, depending on the circumstances.

Reputation systems are complex and usually specific to a given message receiver, and a meaningful discussion of such a broad topic is beyond the scope of this document.

5.3. Can I tell intermediaries not to use ARC?

At present there is no way for a message sender to request that intermediaries not employ ARC.

6. References

6.1. Normative References

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

[RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, DOI 10.17487/RFC5322, October 2008, <http://www.rfc-editor.org/info/rfc5322>.

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

[RFC7601] Kucherawy, M., "Message Header Field for Indicating Message Authentication Status", RFC 7601, DOI 10.17487/RFC7601, August 2015, <http://www.rfc-editor.org/info/rfc7601>.

6.2. Informative References

[ARC] Andersen, K., Rae-Grant, J., Long, B., Adams, T., and S. Jones, "Authenticated Received Chain (ARC)", October 2015, <https://tools.ietf.org/html/draft-andersen-arc-00>.

[DMARC-INTEROP] Martin, F., Lear, E., Draegen, T., Zwicky, E., and K. Andersen, "Interoperability Issues Between DMARC and Indirect Email Flows", August 2015, <https://tools.ietf.org/html/draft-ietf-dmarc-interoperability-06>.

[ENHANCED-STATUS] "IANA SMTP Enhanced Status Codes", n.d., <http://www.iana.org/assignments/smtp-enhanced-status-codes/smtp-enhanced-status-codes.xhtml>.

[OAR] Chew, M. and M. Kucheraway, "Original-Authentication-Results Header Field", February 2012, <https://tools.ietf.org/html/draft-kucherawy-original-authres-00>.

[RFC7489] Kucherawy, M., Ed. and E. Zwicky, Ed., "Domain-based Message Authentication, Reporting, and Conformance (DMARC)", RFC 7489, DOI 10.17487/RFC7489, March 2015, <http://www.rfc-editor.org/info/rfc7489>.

6.3. URIs

[1] mailto:arc-discuss@dmarc.org

Appendix A. GLOSSARY

ADMD Administrative Management Domain as used in [RFC5598] and similar references refers to a single entity operating one or more computers within one or more domain names under said entity's control. One example might be a small company with a single server, handling email for that company's domain. Another example might be a large university, operating many servers that fulfill different roles, all handling email for several different domains representing parts of the university.

ARC ARC is an acronym: Authentication Results Chain - see also [ARC]

ARC-Seal An [RFC5322] message header formed in compliance with the ARC specification. It includes certain content from all prior ARC participants, if there are any.

ARC-Message-Signature An [RFC5322] message header formed in compliance with the [ARC] specification. It includes certain content about the message as it was received and manipulated by the intermediary who inserted it.

Authentication Results Chain (ARC) A system that allows a Message Receiver to identify Intermediaries or Message Handlers who have conveyed a particular message. For more information see the Abstract of this document, or refer to [ARC].

Domain Naming System Block List (DNSBL) This is a system widely used in email filtering services whereby information about the past activity of a set of hosts or domains indicates that messages should not be accepted from them, or at least should be subject to greater scrutiny before being accepted. Common examples would be SpamCop, Spamhaus.org, SORBS, etc.

Email Service Provider (ESP) An Email Service Provider is typically a vendor or partner firm that sends mail on behalf of another company. They may use email addresses in Internet domains belonging to the client or partner firm in various [RFC5321]

fields or [RFC5322] message headers of the messages they send on their behalf.

Intermediary In the context of [ARC], an Intermediary is typically an Administrative Management Domain (per [RFC5598]) that is receiving a message, potentially manipulating or altering it, and then passing it on to another ADMD for delivery. Also see [DMARC-INTEROP] for more information and discussion. Common examples of Intermediaries are mailing lists, alumni or professional email address providers like universities or professional organizations, et cetera.

Mail/Message Transfer Agent (MTA) This refers to software that sends and receives email messsages across a network with other MTAs. Often run on dedicated servers, common examples are Exim, Microsoft Exchange, Postfix, and Sendmail.

Mailflow A group of messages that share features in common. Typical examples would be all messages sent by a given Message Sender to a Message Receiver, related to a particular announcement, a given mailing list, et cetera.

Malicious Actor A Malicious Actor is a party, often an Intermediary, that will take actions that seek to exploit or defraud the ultimate recipient of the message, or subvert the network controls and infrastructure of the Message Receiver. Typical examples would be a spammer who forges content or attributes of a message in order to evade anti-spam measures, or an entity that adds an attachment containing a virus to a message.

Message Handler A Message Handler is another name for an Intermediary.

Message Receiver In the transmission of an email message from one ADMD to another, this is the organization receiving the message on behalf of the intended recipient or end user. The Message Receiver may do this because the intended recipient is an employee or member of the organization, or because the end user utilizes email services provided by the Message Receiver (Comcast, GMail, Yahoo, QQ, et cetera).

Message Sender In the transmission of an email message from one ADMD to another, this is the organization sending the message on behalf of the Originator or end user.

Originator This refers to the author of a given email message. In different contexts it may refer to the end-user writing the message, or the ADMD providing email services to that end-user.

Reputation In the larger context of email hygiene - blocking spam and malicious messages - reputation generally refers to a wide variety of techniques and mechanisms whereby a message receiver uses the past actions of a sending host or domain to influence the handling of messages received from them in the future. One of the classic examples would be a Spamhaus-style DNSBL, where individual IP addresses will be blocked from sending messages because they've been identified as being bad actors. Very large message receivers may build and maintain their own reputation systems of this kind, whereas other organizations might choose to use commercial products or free services.

Reputation Service Provider A Reputation Service Provider would be a source of reputation information about a message sender. In this context, the DNSBL services offered by Spamhaus would allow them to be referred to as an RPS. Many spam and virus filtering vendors incorporate similar functionality into their services.

Request For Comment (RFC) RFCs are memoranda that "contain technical and organizational notes about the Internet." Created and managed by the Internet Engineering Task Force (IETF), they are de facto standards for various methods of communicating or collaborating over the Internet.

RFC5321 - Simple Mail Transfer Protocol This document describes the protocol used to transfer email messages between Message Transfer Agents (MTA) over a network. Link: [RFC5321]

RFC5322 - Internet Message Format This document describes the format of Internet email messages, including both the headers within the message and various types of content within the message body. Link: [RFC5322]

Validator A Message Receiver that attempts to validate the ARC chain in a message.

Appendix B. References

Appendix C. Acknowledgements

This draft is the work of OAR-Dev Group.

The authors thanks the entire OAR-Dev group for the ongoing help, innumerable diagrams and discussions from all the participants, especially: Alex Brotman, Brandon Long, Dave Crocker, Elizabeth Zwicky, Franck Martin, Greg Colburn, J. Trent Adams, John Rae-Grant, Mike Hammer, Mike Jones, Steve Jones, Terry Zink, Tim Draegen.

Appendix D. Comments and Feedback

Please address all comments, discussions, and questions to arc-discuss@dmarc.org [1][mailto:arc-discuss@dmarc.org].

Authors' Addresses

OAR-DEV Group

Email: arc-discuss@dmarc.org

Steven Jones
DMARC.org

Email: smj@crash.com

John Rae-Grant
Google

Email: johnrg@google.com

J. Trent Adams
Paypal

Email: trent.adams@paypal.com

Kurt Andersen (editor)
LinkedIn
2029 Stierlin Ct.
Mountain View, California 94043
USA

Email: kurta@linkedin.com