Network Working Group
Internet-Draft
Obsoletes: 8617 (if approved)
Intended status: Informational
Expires: 7 June 2026

T. Adams
Proofpoint
J. Levine
Taughannock Networks

4 December 2025

Concluding the ARC Experiment
draft-adams-arc-experiment-conclusion-01

Abstract

This document calls for a conclusion to the experiment defined by “The Authenticated Received Chain (ARC) Protocol,” (RFC8617) and recommends that ARC no longer be deployed or relied upon between disparate senders and receivers. The document summarizes what ARC set out to do, reports on operational experience, and explains how the experience gained during the experiment is being incorporated into the proposed DKIM2 work as the successor to DomainKeys Identified Mail (DKIM). To avoid any future confusion, it is therefore requested that ARC (RFC8617) be marked “Obsolete” by the publication of this Internet-Draft.

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 7 June 2026.

Copyright Notice

Table of Contents

1. Introduction

Following the deployment of DMARC [RFC7489] that aligned author domains with SPF [RFC7208] / DKIM [RFC6376] and provided a method to request receiver handling for authentication failures, while DKIM continued to provide message-level signatures, it became clear that there was a failure case that needed to be addressed. Real-world forwarding and modifications performed by mailing list managers frequently broke the authentication protocols that underpin DMARC, motivating the ARC experiment as a potential mitigation.

As a response, ARC [RFC8617] was introduced as an experiment to determine whether a cryptographically verifiable "chain of custody" for email, as assembled by intermediaries rewriting messages, could preserve the original sender's authentication results across forwarding, mailing lists, and other intermediaries. ARC's premise was that each handler could record its view of upstream authentication and then sign that record, enabling downstream evaluators to see what happened along the path.

This document reports the experiment's results and explains why, with the emergence of the proposed successor to DKIM, currently known as DKIM2, the community should retire ARC and incorporate the useful pieces directly into the successor to DKIM.

2. Background

2.1. Problem Space: DMARC Breakage at Intermediaries

DMARC relies on successful SPF and/or DKIM authentication along with alignment with the Author Domain. When intermediaries modify a message (for example, subject or body changes, footer insertion, MIME adjustments), DKIM signatures from the originator can fail to verify; when an intermediary relays mail through different IPs than are defined within the originator's SPF record, SPF authentication can fail. As a result, messages that were legitimate at origination can appear unauthenticated downstream, even if the intermediary handling is benign. ARC was proposed to let trustworthy intermediaries attest to what they saw before the breakage occurred and add a new signature to the message, essentially creating a signature chain.

Forwarding remains one of the most pervasive sources of broken authentication results. When a recipient's mail is automatically forwarded (for example, via a mailing list, auto-forward rule, or redirect), the forwarding infrastructure appears as the sending IP, not the IP of the original sending domain, so SPF authentication fails by design. DKIM may survive only if the signature remains intact through forwarding, but many forwarding systems change headers or bodies (footers, mailing list tags, encodings), thus invalidating DKIM and causing DMARC to fail.

Because the forwarding party is typically not in the author's domain control and cannot easily be enumerated in the author's SPF record, it becomes operationally infeasible for senders to cover every possible forwarder. As such, broken authentication at forwarders represents a structural gap in DMARC deployment.

The forwarder's participation and transformations therefore form the very scenario that the ARC experiment targeted, namely intermediaries rewriting messages and breaking original authentication signals, and the hope that those intermediaries could attest to the original author's state via a chain of custody.

2.2. ARC Overview

To address these failure modes, ARC defines header fields (ARC-Seal, ARC-Message-Signature, ARC-Authentication-Results) that allow each intermediary to:

* capture its input authentication assessment (typically DMARC-related),

* indicate that some transformation happened but not what changed, and

* sign its contribution, forming a verifiable chain that subsequent receivers can validate.

ARC does not assert message authorship; it asserts a sequence of handling and observations only by those participating in ARC signing. Verification yields two outputs: (1) whether the chain is cryptographically valid, and (2) what those upstream assessments (if any) were. It makes no value assertion of the email nor if there were any intermediary handlers not participating in ARC handling or signing.

2.3. Scope and Non-Goals of ARC

The experiment explicitly limited ARC's role to signaling: it could reveal that certain intermediaries participated and re-signed messages, but a validated ARC chain was not intended to convey trust in any signer on its own. Trust decisions were left to receivers' local policy. As such, without a robust reputation system, ARC in-and-of itself cannot convey trust in an email that fails DMARC.

Another limitation of the design was that the ARC signature only indicated the intermediaries handling the message, but was silent about any changes the intermediaries made to the message. As such, a fully validated ARC chain might include a modified message without the final evaluator knowing what changes were made.

3. Analysis of the ARC Experiment

3.1. Operational Experience

This section summarizes widely reported deployment observations from operators and implementers during the ARC experiment.

* Data-center and intra-domain utility: Early effective use of ARC occurred inside single administrative domains or tightly controlled data centers, where messages traversed multiple internal hops. Operators applied ARC to ensure messages were not modified unexpectedly between their own servers. In these cases, operators already had implicit trust and operational control.

* Internet-scale dependency on reputation: For broad interoperability, ARC required evaluators to run a reputation system for ARC signers. Verifying the cryptography was necessary but insufficient; evaluators needed to decide whether to trust each signer in the chain before using their assertions to override DMARC outcomes. This created a parallel trust infrastructure, separate from (but interacting with) existing domain or IP reputation.

* Limited reputation deployment:Even early deployers that validated ARC chains did not deploy robust, dynamic reputation. Instead, they implemented “allow lists” of intermediaries whose ARC assertions were always (or mostly) accepted. This provided utility in specific bilateral or consortium relationships but did not scale to the open Internet.

* Complex evaluator policy: Receivers faced policy questions: how many hops of ARC to honor, how to treat partial or broken chains, how to reconcile conflicting assessments across chain links, and under what conditions ARC could influence DMARC enforcement. The resulting diversity limited predictable interoperation across receivers.

* Forwarding-driven breakage still dominant: Because email forwarding automatically changes the apparent sender infrastructure (for example, the forwarding system’s IP rather than the original domain), many well-authenticated messages fail DMARC at the final recipient purely due to forwarders. Forwarding often results in SPF failure by design and DKIM failure due to header or body modifications. This reinforces that any intermediary authentication or chaining mechanism (such as ARC) must address the uncontrolled nature of forwarding, which spans countless unknown and dynamic systems, rather than only known mailing lists or enterprise relays.

* Abuse of trusted chains: Attackers have been leveraging situations where recipients trusted ARC chains by modifying email after an initial attestation and then passing it to another entity which trusted that attestation. Assessing whether the email came directly from the entity at the head of the ARC chain is complex and assessing whether or not any changes to the email were made by a trusted signer or a malicious third party has proved to be impossible.

* Limited deployment and use compared to DKIM: Implementation of ARC has been limited at Internet scale with only about 10,000 domains having been observed to send email with ARC-signed headers. Even fewer have done so in full compliance with the specification since

it was published as an RFC in 2019. For comparison, 9.5 million DKIM records were found via passive DNS inspection over the same six years with 460 million DKIM signatures observed from real-world email headers.

* Ecosystem shift to successor work: As the community prioritized addressing DKIM replay and strengthening end-to-end authenticity, the DKIM working group has initiated work on what is being called DKIM2. That effort explicitly considers incorporating ARC-like "handling assertions" where they add value, while avoiding a separate global trust fabric for intermediaries. As focus has shifted from ARC to DKIM2, incorporating the learning from the ARC experiment, there is no longer any meaningful effort to continue developing and deploying ARC.

3.2. ARC's Core Lesson: Signatures Are Not Trust

ARC successfully demonstrated that intermediaries can publish a cryptographically verifiable history of handling. However, verifiable history without reputation does not enable safe override of DMARC or other enforcement policies. Any Internet-wide solution must pair verifiable signals with a scalable, abuse-resistant trust model; ad hoc allow lists are not sufficient.

3.3. No Indication of Modifications

When the content of an email was modified by an intermediary, breaking the DKIM signature, ARC was able to identify the intermediary that performed the modification via a signature, through ARC doesn't define a mechanism to identify what was modified in the message or why it was modified. This left the interpretation of whether or not the email should be accepted up to the evaluator's ability to determine the reputation of the intermediary.

3.4. Reputation at Each Hop Is Operationally Heavy

Operating, sharing, and refreshing reputation for potentially thousands of intermediaries is expensive and complex. Without a common reputation framework, ARC yielded inconsistent receiver behavior and created incentive for attackers to infiltrate or mimic "trusted" intermediaries.

The forwarding problem illustrates this operational burden: the number of potential forwarders is vast and dynamic, making it unrealistic to maintain allow-lists or reputation records for all of them.

Attempts to create internet-scale reputation systems for ARC have not been successful during the ten years of the experiment, and it as there is no known plan for one in development, it is unlikely there will be one in the future.

3.5. Favoring the DKIM2 Approach

The DKIM2 motivation identifies replay as a critical gap and proposes signing the source and destination for each message, along with mechanisms better aligned with modern routing patterns. Incorporating ARC's useful elements (for example, signed assertions about handling) into DKIM2 avoids a parallel chain or signature stack and reduces reliance on separate hop-by-hop reputation.

3.6. Conclusions of the ARC Experiment

Based on community experience and the direction of the DKIM2 work:

a. The ARC experiment is over. Implementers and operators should not rely on ARC going forward, and should cease further Internet-wide deployments. Existing ARC deployers should plan to decommission them or confine their usage to controlled, intra-domain contexts where bilateral policy suffices.

b. Experience from the ARC experiment is informing the development of DKIM2. The DKIM working group is actively developing DKIM2; relevant ARC insights, such as durable capture of upstream authentication state and intermediary handling, shall inform DKIM2 design where appropriate.

c. RFC 8617 should be marked Obsolete. This document requests that the RFC Editor and IESG mark RFC 8617 as "Obsolete" upon publication of this draft (or its successor) to conclude the experiment and prevent new deployments.

4. Guidance to Implementers and Operators

* Receivers that still parse ARC headers may continue to verify them for forensic or intra-domain purposes, but should not make delivery decisions based on ARC chain validity without robust reputational trust signals and associated policies.

* New ARC deployments are discouraged since they are unlikely to provide useful information for mail processing.

* Anyone interested in ARC should follow the development of DKIM2 as it matures through the IETF process.

5. Acknowledgements

The authors of RFC8617 and the many operators who deployed and evaluated ARC provided the data and experience that made these conclusions possible. The DKIM working group's current efforts, including the DKIM2 motivation and related drafts, informed the direction recommended here. Thanks also to those who helped review and edit this draft including (but not limited to) Todd Herr, Richard Clayton, Alex Brotman, Marc Bradshaw, and Emanuel Schorsch.

6. IANA Considerations

This document has no IANA actions.

7. Security Considerations

ARC's separation of "verification" from "trust" created risks when evaluators accepted chains from low-reputation or compromised intermediaries. Attackers could attempt to route through permissive handlers to gain favorable treatment. Ending the experiment and migrating learnings into DKIM2, along with explicit controls to mitigate replay and stronger binding of message context, should reduce these risks. Operators must treat residual ARC processing as diagnostic only, unless backed by robust, auditable trust frameworks.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[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>.

[RFC7208] Kitterman, S., "Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1", RFC 7208, DOI 10.17487/RFC7208, April 2014, <https://www.rfc-editor.org/info/rfc7208>.

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

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8617] Andersen, K., Long, B., Ed., Blank, S., Ed., and M. Kucherawy, Ed., "The Authenticated Received Chain (ARC) Protocol", RFC 8617, DOI 10.17487/RFC8617, July 2019, <https://www.rfc-editor.org/info/rfc8617>.

8.2. Informative References

[ARC-SPEC] "ARC Specification site, background and history", 2019, <https://arc-spec.org/>.

[I-D.ietf-dkim-dkim2-motivation] Gondwana, B., Clayton, R., and W. Chuang, "DKIM2 - signing the source and destination of every email", Work in Progress, Internet-Draft, draft-ietf-dkim-dkim2-motivation-02, 2 November 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-dkim-dkim2-motivation-02>.

Authors' Addresses

J. Trent Adams
Proofpoint
105 Edgeview Drive, Suite 440
Broomfield, CO 80021
United States of America
Email: tadams@proofpoint.com

John Levine
Taughannock Networks
PO Box 727
Trumansburg, NY 14886
United States of America
Email: standards@taugh.com