TAG + digicert®

Leveraging DMARC in the Context of the Modern AI and MCP-Enabled Enterprise

Analyst Report | 2026

Leveraging DMARC in the Context of the Modern AI and MCP-Enabled Enterprise

Dr. Edward Amoroso
Founder and CEO, TAG | Research Professor, NYU
eamoroso@tag-cyber.com

Alexander Garcia-Tobar
Founder and Former CEO, ValiMail | Strategic Advisor, DigiCert
alex@valimail.com

Version 1.0
June 12, 206

Abstract

This research report explains how the use of Domain-based Message Authentication, Reporting, and Conformance (DMARC) in enterprise has evolved over the past decades, how it is being used in enterprise today for secure and authenticated email, and how it provides a useful security framework to improve authentication for emerging AI-based workload and agent communications using Model Context Protocol (MCP).

Introduction: What is DMARC?

Domain-based Message Authentication, Reporting, and Conformance (DMARC) emerged in the early 2010s to address email spoofing and phishing attacks that were exploiting weaknesses in email authentication methods. Hackers had quickly identified means for spoofing emails for the purpose of phishing, fraud, and other forms of business email compromise. Improved authentication was an obvious means for reducing the risk of such attacks.

Prior to DMARC, organizations had used SPF (Sender Policy Framework) and DKIM (DomainKeys Identified Mail) to try to validate that a message was authorized by the sending domain and had not been altered in transit. While effective in isolation, these mechanisms lacked alignment and policy enforcement, meaning that even if a message failed authentication checks, receiving mail systems were not consistently instructed on how to handle it.

Recognizing this gap, industry leaders from Google, Microsoft, Yahoo, PayPal, and elsewhere created DMARC to provide domain owners with a means to authenticate their email and specify how receiving systems should treat messages that fail authentication checks. DMARC builds on

SPF and DKIM by introducing the concept of alignment, ensuring that the domain used in the visible “From” address matches the domain validated by authentication mechanisms.

It also allows domain owners to publish policies in DNS that instruct receiving mail servers to take specific actions such as monitoring, quarantining, or rejecting suspicious messages. In addition, DMARC introduces a reporting capability, enabling organizations to receive aggregate and forensic feedback about how their email is being handled across the internet. This visibility is critical for identifying abuse, misconfigurations, and unauthorized use of a domain.

Ultimately, DMARC emerged as a pragmatic solution to improve trust in email communications. By combining authentication, policy enforcement, and reporting into a single framework, the DMARC standard, and the many commercial vendors supporting its practical implementation, helped organizations better protect their brand, reduce their phishing risk, and improve the security email deliverability.

In this report, we provide an update on DMARC since the commercial landscape has shifted during the past few years, and as AI has emerged as an enabling factor in both cyber offense and defense. We believe DMARC remains an important aspect of enterprise security, and we are eager to share how we view its foundational approach as providing useful guidance for how trust can be implemented in emerging AI agentic communications.

How Has DMARC Been Supported in Enterprise?

As alluded to above, enterprise adoption of DMARC increased through the 2000’s and 2010’s as organizations began to recognize email as a primary attack vector and brand impersonation as a material business risk. While the standard itself is published openly, operationalizing DMARC at scale in a complex enterprise environment proved challenging, which is why our industry saw many excellent startup companies emerge to assist with the task.

The challenge then, as well as now, is that large organizations would typically be managing dozens or even hundreds of legitimate email sources, ranging from internal mail servers to SaaS platforms and third-party marketing tools, which must all be properly authenticated and aligned. The potential to be spoofed or a victim of an email-based fraud activity is obviously high in such situations.

What followed was a process where most enterprise CISOs adopted a phased approach to dealing with the issue starting with monitoring policies (p=none) to gather visibility, then gradually enforcing quarantine or reject policies as confidence improved. This transition required not just technical configuration, but also cross-functional coordination between security, IT, marketing, and third-party vendors.

To address these challenges, a new category of DMARC-focused vendors emerged to simplify deployment, provide visibility, and automate policy enforcement. Early pioneers such as Agari and ValiMail played a critical role in translating the DMARC specification into enterprise-ready solutions. These vendors developed platforms that ingest and analyze DMARC aggregate

reports, identify authorized and unauthorized senders, and guide organizations through the process of achieving full enforcement.

In addition, they introduced automation capabilities to help maintain accurate SPF and DKIM configurations, reduce manual effort, and minimize the risk of disrupting legitimate email flows. Their offerings often extended beyond basic compliance to include threat intelligence, phishing detection, and brand protection services.

Over time, DMARC support has become a standard feature across broader email security and cloud platforms, including native capabilities from providers like Microsoft and Google. However, the foundational work done by early vendors such as Agari and ValiMail helped establish best practices and demonstrated the business value of DMARC enforcement, particularly in reducing phishing success rates and protecting customer trust.

Today, enterprise DMARC programs are often integrated into a larger identity and email security strategy, with continuous monitoring, automated remediation, and alignment with Zero Trust principles. This evolution reflects a broader shift in cybersecurity, where identity-based controls, including domain authentication, are treated as essential components of enterprise defense.

What Are Some Operational Challenges with DMARC?

A central operational challenge with DMARC is the fear of unintended business disruption. Specifically, there is concern that legitimate emails will be blocked, quarantined, or silently dropped once enforcement policies are enabled. Moving from a monitoring posture (**`p=none`**) to enforcement (**`p=quarantine`** or **`p=reject`**) introduces real risk, particularly in large enterprises with complex email ecosystems.

IT and security teams often lack a complete inventory of all systems sending email on behalf of their domain, including shadow IT, legacy applications, and third-party SaaS platforms. If these sources are not properly configured with aligned SPF (Sender Policy Framework) and DKIM (DomainKeys Identified Mail), then legitimate business communications, such as invoices, customer notifications, or marketing campaigns, can fail authentication and be rejected.

Another persistent challenge is the complexity of analyzing and acting on DMARC reporting data. While DMARC provides valuable aggregate and forensic reports, the data is often voluminous, highly technical, and difficult to interpret without specialized tooling. Enterprises must parse XML-based reports from potentially hundreds of receiving mail systems, correlate sending sources, and determine which are legitimate versus malicious.

In addition, maintaining alignment over time is not a one-time effort. New services, vendors, and applications are constantly introduced, requiring ongoing updates to SPF records, DKIM keys, and DMARC policies. Without continuous monitoring and governance, organizations risk configuration drift that can either weaken protection or inadvertently disrupt valid communications.

Finally, process-related challenges often prove just as difficult as technical hurdles. DMARC implementation requires coordination across multiple stakeholders, including IT, security, marketing, legal, and external partners, each of whom may control different email systems and have varying risk tolerances. For example, marketing teams may prioritize email deliverability and campaign reach, while security teams focus on strict enforcement to prevent spoofing.

Additionally, third-party vendors may not fully support the objective of establishing DMARC alignment or may resist configuration changes, further complicating enforcement efforts. These combined concerns reinforce a cautious, incremental approach to DMARC, even as the threat landscape continues to push organizations or all sizes and sectors toward stronger identity-based email protections.

What is the Current Status of DMARC in Enterprise?

The current status of DMARC in the enterprise reflects broad awareness and growing adoption, but uneven maturity in enforcement. Most large organizations have at least published a DMARC record, often starting with a monitoring policy (`p=none`)[1] to gain visibility into who is sending email on their behalf. This baseline step has become common, driven in part by regulatory encouragement, cyber insurance expectations, and mandates in certain sectors.

However, a smaller subset of enterprises has progressed to full enforcement (**`p=reject`**), where unauthorized emails are actively blocked. The gap that emerges between awareness and enforcement is largely due to the operational complexities associated with aligning all legitimate senders and ensuring proper configuration of SPF and DKIM across diverse environments.

At the same time, DMARC has become increasingly embedded within broader enterprise email and identity security strategies. Native support from the two dominant email platforms from Microsoft and Google has simplified deployment, while specialized vendors continue to provide advanced analytics, automation, and managed services to help organizations reach enforcement safely.

There is also a shift toward continuous DMARC management rather than one-time deployment, reflecting the dynamic nature of enterprise email ecosystems. As phishing and business email compromise remain persistent threats, DMARC is now widely viewed as a foundational control. That is, it is viewed as necessary, but not sufficient, within a layered defense strategy centered on identity protection and Zero Trust principles.

How Can DMARC Serve as a Model to Secure AI Agent Communications?

As explained above, DMARC was designed to secure human-centric email communications, but its underlying principles, including domain authentication, policy enforcement, and reporting, translate naturally into the emerging world of AI agent communications. As AI agents begin to

[1] See the Appendix for an overview (with example) of how DMARC works.

interact with one another, invoke APIs, and exchange structured messages across domains, the need to verify the authenticity of the sending domain becomes just as critical as it is in email.

In this context, DMARC can serve as a conceptual model for establishing trust boundaries, thus ensuring that an AI agent claiming to represent a given enterprise domain is in fact authorized to do so. When combined with existing security controls such as SPF and DKIM, DMARC-like mechanisms could be adapted to validate signed messages or API requests exchanged between agents.

In environments leveraging Model Context Protocol, where AI agents dynamically share context, tools, and instructions across systems, the increased risk of spoofed or manipulated communications becomes significant. An adversary could attempt to impersonate a trusted agent, inject malicious prompts, or replay previously valid messages to influence downstream behavior.

While DMARC does not natively support these non-email protocols, its model of ensuring that the identity presented matches the authenticated origin offers a useful design pattern. Extending this concept, MCP-based systems could require cryptographic signing of agent messages, domain-based identity validation, and policy-driven handling of failed authentication (for example, rejecting or sandboxing untrusted inputs).

That said, DMARC alone is not sufficient to secure AI agent communications. AI interactions are more dynamic than email, requiring additional controls such as mutual authentication, token-based authorization, secure enclaves for execution, and continuous behavioral monitoring. Enterprises will need a new class of “DMARC for agents” frameworks that combine identity verification with runtime policy enforcement and observability.

What are the Next Steps for DMARC in an AI-Enabled Era?

The next phase of DMARC in an AI-enabled era should begin with leadership from standards bodies to evolve domain-based trust beyond traditional email. Organizations such as Internet Engineering Task Force and OASIS are well-positioned to define extensions or adjacent frameworks that apply DMARC-like principles to API-driven and agent-to-agent communications.

While the design decisions inherent in such standards-based interactions are tough to predict, one would expect that the resulting recommendations for extension would include formalizing mechanisms for domain-aligned identity in machine interactions, standardized cryptographic signing of messages, and policy frameworks that dictate how receiving systems should respond to failed authentication.

Importantly, these security standards must be designed for interoperability across cloud platforms and AI ecosystems, with support for emerging paradigms such as Model Context Protocol. Just as DMARC unified SPF and DKIM, the next generation of standards should unify identity, integrity, and provenance for AI communications.

Vendors, for their part, should operationalize these concepts into enterprise-ready solutions. Email security providers, identity platforms, and AI security startups alike have an opportunity to build “DMARC for agents” capabilities that extend domain authentication into API calls, prompt exchanges, and autonomous workflows. This includes embedding cryptographic identity into AI agents, providing validation of message origin, and more.

Vendors should also invest in usability, abstracting the complexity of configuration and providing clear visibility into agent communication flows, much like early DMARC pioneers did for email. Integration with existing enterprise infrastructure, including SIEM, identity providers, and Zero Trust architectures, will be critical to ensure adoption and avoid creating yet another siloed control plane.

Auditors and enterprise security teams will also play a pivotal role in driving accountability and adoption. Audit frameworks should begin to incorporate controls for authenticated machine-to-machine communication, extending beyond email into AI-driven workflows and automated decision systems. This includes validating that enterprises have visibility into agent interactions, enforce domain-aligned identity, and maintain logs for forensic analysis.

At the same time, enterprise security teams should begin preparing now by inventorying AI agents and services, establishing identity and access controls for machine actors, and piloting cryptographic signing and verification mechanisms where possible. They should treat AI communication channels as a new attack surface that is subject to spoofing, replay, and manipulation, and apply lessons learned from DMARC deployments.

Appendix A: Sample DMARC Record

In this appendix, let’s examine the DMARC Record for Microsoft.com, and use its settings to explain some of the more common choices made:

```
v=DMARC1; p=reject; pct=100; rua=mailto:itex-rua@microsoft.com;
ruf=mailto:itex-ruf@microsoft.com; fo=1
```

Let’s examine each element of the record. To start, **`v=DMARC1`** is known as a version identifier. It specifies that the record is using DMARC version 1, which is the current and only widely supported standard. This is a required field in every DMARC record, and it tells receiving mail systems how to interpret the policy.

The **`p=reject`** is known as a policy action, which is the core enforcement directive where **`reject`** means that if an email fails DMARC authentication (SPF and/or DKIM alignment), then the receiving server should reject it outright (not deliver it). This is maximum enforcement and it indicates that Microsoft has high confidence in its SPF/DKIM configuration. This represents strong protection against phishing and spoofing of microsoft.com.

The **pct=100** references percentage of messages to which the policy is applied. During rollout phases, some organizations might set this to 10%, 25%, or 50%. Obviously, the 100% implies that Microsoft is at full enforcement maturity and that all failed messages will be rejected.

The **rua=mailto:itex-rua@microsoft.com** defines the aggregate report URI, which defines where aggregate DMARC reports are sent. Reports are sent by receiving mail systems, containing summarized data (not individual emails) such as source IPs sending mail, pass/fail rates for SPF/DKIM, and volume statistics. This is used for monitoring domain health and abuse trends, thus enabling better detection of spoofing campaigns at scale.

The **ruf=mailto:itex-ruf@microsoft.com** is known as a forensic report URI. It defines where forensic (failure) reports are sent and is sent when individual messages fail DMARC. It may include message headers and/or portions of the email. This is useful for debugging misconfigurations or investigating attacks, but it is less widely supported due to privacy concerns.

Finally, the **fo=1**is a failure reporting option, which controls when forensic reports are generated. In this case, **fo=1** means to generate a report if either SPF or DKIM fails (not necessarily both). Other options include **fo=0** if both fail (default), **fo=d** for DKIM failures only, and **fo=s** for SPF failures only.

About TAG

Recognized by Fast Company, TAG is a trusted next generation research and advisory company that utilizes an AI-powered SaaS platform to deliver on-demand insights, guidance, and recommendations to enterprise teams, government agencies, and commercial vendors in cybersecurity and artificial intelligence.