port25
A Message Systems Company

WHITE PAPER

FAQ: DMARC for Email Service Providers

Email authentication is becoming a big deal.
Here’s a guide to the key standards from ValiMail.

port25
A Message Systems Company

WHITE PAPER

FAQ: DMARC for Email Service Providers

Port25 **Evaluation Requests**
By Hourly Email Volume

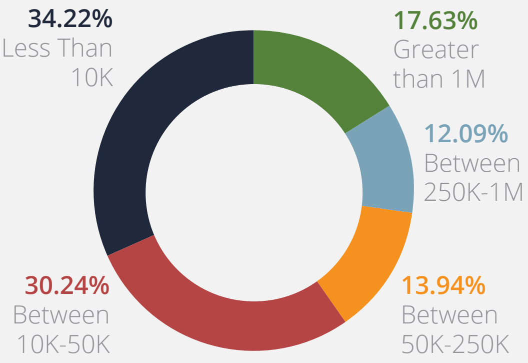

Port25 **LinkedIn Engagement Data**
By Company Size

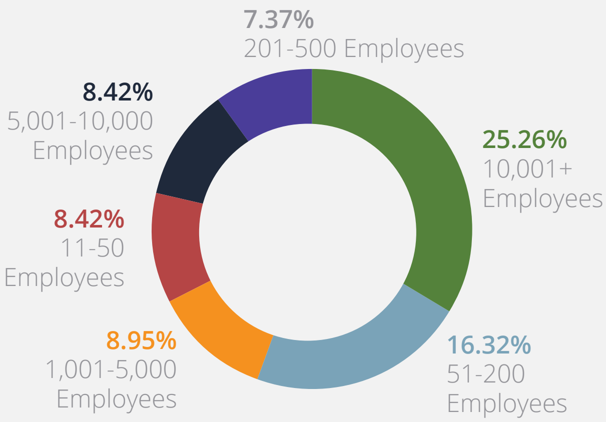

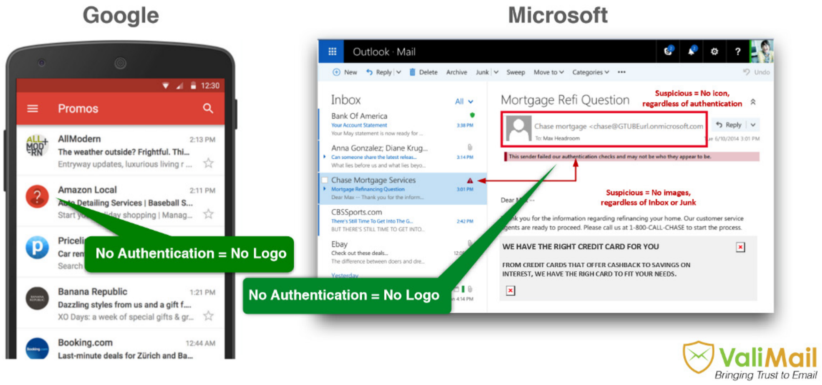

This sender failed our authentication checks and may not be who they appear to be.

WE HAVE THE RIGHT CREDIT CARD FOR YOU

FROM CREDIT CARDS THAT OFFER CASHBACK TO SAVINGS ON INTEREST, WE HAVE THE RIGH CARD TO FIT YOUR NEEDS.

Why email authentication matters: This is how non-authenticated email will look in Gmail and Outlook, starting sometime in 2016.

WHAT IS DMARC?

DMARC—Domain-based Message Authentication, Reporting and Conformance—is an open email authentication standard that sending domains use to block fraudulent emails. DMARC is built on top of two earlier standards—SPF and DKIM—and adds additional features like reporting, policy definition, and the notion of identity alignment.

DOES DMARC STOP PHISHING?

When configured correctly, DMARC can completely stop phishing attacks in which the attacker sends an email with a 'From' address that appears to originate from a protected domain. As this is the primary form of phishing attack, DMARC is a very effective tool to defend customers, employees, partners, and others from phishing.

DOES DMARC HELP DELIVERABILITY?

Large-scale email receivers, such as Google, Microsoft, and Yahoo!, are increasingly requiring that email messages be properly authenticated in a DMARC-compliant way. So adding a DMARC record for a domain, in conjunction with properly configured SPF and/or DKIM records, will help ensure proper delivery.

Furthermore, the proper use of DMARC ensures that messages sent by spammers using a sender's domain will not negatively impact the domain's overall reputation. Such spam will be blocked and the sender's brand will be protected.

WHAT DOES DMARC ADD ON TOP OF EXISTING EMAIL AUTHENTICATION STANDARDS LIKE SPF AND DKIM?

DMARC contributes three major elements to previously existing email authentication standards:

Reporting:
DMARC-participating receivers agree to provide email authentication reports to sending domains. This allows the owners of these sending domains to understand the current state of email authentication for their domain, see legitimate services that may not be properly authenticating, and identify sources of domain abuse.

Policy:
With DMARC, sending domains can recommend how a receiver should treat an email that fails authentication, rather than leaving it to the discretion of the receiver. This allows sending domains to authenticate all sources of legitimate email over time, rather than requiring domain owners to fix all authentication issues immediately. A report-only policy of 'p=none' can be useful during this investigation phase, but domain owners should strive to reach an enforcement level of 'p=quarantine' or 'p=reject'.

Identity Alignment:
There are multiple sources of identity in an email message (including the From address, DKIM signature identity, and Return-Path address). DMARC prioritizes the human-readable From address as a source of identity, and only considers authentication results for identities that are aligned with this From address. SPF and DKIM use different sources of identity, and so the authentication they provide will only prevent fraud if their source of identity matches the human-readable From address in some way.

WHAT ARE THE DIFFERENT POLICY LEVELS DEFINED BY DMARC?

DMARC defines three policy levels that describe how receivers are supposed to handle email failing authentication.

These levels are 'p=none', 'p=quarantine', and 'p=reject'.

none:
Receivers are instructed to not change how they deliver email based on email authentication failures. The 'none' level is typically used when a domain owner is in the initial process of authenticating their email services; moving beyond this level is key to enable DMARC to stop fraud.

quarantine:
Receivers are asked to mark messages failing authentication as spam.

reject:
Receivers are requested to block messages failing authentication entirely, and not deliver them to their intended recipients.

In all cases the policy is enforced by the system receiving the email, and the receiving system may choose to handle email delivery differently that prescribed by the DMARC policy. For example, Microsoft Office 365 treats 'quarantine' and 'reject' identically.

WHAT IS IDENTITY ALIGNMENT?

Email messages typically include multiple sources of identity, each of which may be associated with a different domain. Some of these include:

- **The human-readable From address (RFC 5322.from, pra)**
- **The Return-Path address (RFC 5321.from, mfrom)**
- **Zero, one, or more identities associated with DKIM signatures of the message**
- **The Reply-To address**
- **The Sender address**

DMARC is primarily an anti-fraud standard, so it focuses on the identity that is displayed to the recipient: the human-readable From address. The domain in the From address is the one authenticated by DMARC.

SPF and DKIM each authenticate a different source of identity than the human-readable From address. So SPF or DKIM authentication by themselves will not provide the level of anti-fraud protection to which DMARC aspires.

Instead, DMARC will only consider a SPF or DKIM authentication result if the domain authenticated by SPF or DKIM matches the domain in the human-readable From address. In this situation the identities are considered 'aligned'.

In cases where SPF or DKIM authenticates with an identity whose domain doesn't match the domain in the human-readable. From address, the non-matching authentication result is simply discarded. DMARC does not cause any change in the underlying SPF or DKIM behavior, and the corresponding identities against which they authenticate.

SO PASSING SPF OR DKIM ISN'T ENOUGH TO ENSURE THAT A MESSAGE PASSES DMARC?

Right. It is very common to see messages that pass SPF and/or DKIM, but fail DMARC because of identity misalignment.

This is a common point of confusion, and many resources on the web get this wrong. Understanding how identity alignment works with SPF and DKIM is crucial to making DMARC work.

WHAT ARE ORGANIZATIONAL DOMAINS?

To support flexible domain matching, DMARC introduces the concept of an organizational domain. Typically, this is the top level domain (TLD) plus one more label. For some countries that use second-level domains to categorize organizations, the organizational domain is this second-level domain plus one more label.

So these are all organizational domains:

- **nytimes.com**
- **bbc.co.uk**
- **wikipedia.org**

And the following domains all share the same organizational domain (bbc.co.uk):

- **mail123.bbc.co.uk**
- **abc.def.x-y.bbc.co.uk**
- **bbc.co.uk**

HOW DOES DMARC USE ORGANIZATIONAL DOMAINS FOR IDENTITY ALIGNMENT?

In the most common ('relaxed') configurations, DMARC authentication does not require an exact match between the domain from the human-readable From address and the domain used in the SPF or DKIM identity. In this mode DMARC only requires that the organizational domains match.

While it is possible to set DMARC records to require exact matches ('strict'), this is generally not necessary for security reasons. Because it is extremely common for organizations to use multiple subdomains for sending email, configuring DMARC to require exact matches is usually ill-advised.

HOW DO DOMAIN OWNERS CONFIGURE DMARC FOR THEIR DOMAINS?

DMARC is configured through the use of DNS TXT records. A domain owner adds a special DNS TXT record on the '_dmarc' subdomain of the domain for which they'd like to configure DMARC. This DNS TXT record defines the policy for the domain.

DO DMARC POLICIES GET INHERITED BY SUBDOMAINS?

Yes, but only from the organizational domain.

The DMARC standard defines a lookup rule for DNS records, which explains how the relevant DMARC policy DNS record is determined. The rule is as follows:

1. **Extract the email domain from the human-readable From address**
2. **Look up a DMARC record on the '_dmarc' subdomain of the email domain**
3. **If a record is found, use that record to determine the DMARC policy and terminate the lookup process**
4. **If no record is found, and the email domain is not an organizational domain, then look up the '_dmarc' subdomain of the corresponding organizational domain. If a record is found on this subdomain, use that record to determine the DMARC policy.**
5. **If no record is found, then the process terminates and DMARC is not enforced for the message.**

A key takeaway from this process is that it is generally sufficient to define a single DMARC record on the organizational domain. Even if an email service provider or domain owner is using a subdomain to send email, they don't need to create separate DMARC records for each subdomain.

In certain circumstances this lookup rule can lead to non intuitive behavior. For example, consider email sent from the domain 'abcd.xyz.example.com'. The lookup rule means that the receiver will potentially look for a DMARC record on the following domains:

1. **_dmarc.abcd.xyz.example.com**
2. **_dmarc.example.com**

A DMARC record on _dmarc.xyz.example.com will be ignored, and won't apply to this email.

WHAT IS THE 'SP' ATTRIBUTE IN A DMARC RECORD, AND DO I NEED TO SET IT?

The 'sp' attribute is the subdomain policy, and it defines the way DMARC behaves for subdomains of the organizational domain, provided that the subdomain doesn't have a DMARC record explicitly defined.

Because of how DMARC records are looked up from subdomains, the 'sp' attribute only affects behavior when it is defined on the organizational domain. If this attribute is defined on a subdomain's DMARC record it will have no effect.

In most cases domain owners should not set an 'sp' attribute.

If you are an email service provider managing a dedicated subdomain then you should not set the 'sp' attribute on the subdomain record.

WHAT DOES A DMARC DNS RECORD LOOK LIKE?

There are only four attributes found in most DMARC DNS records. These are:

v—'DMARC1' for the current DMARC revision. This attribute must appear first.

p—Specifies the enforcement level requested by the sender. Allowed values are 'none', 'quarantine', and 'reject'. This attribute is required, and must be the second attribute in the record.

rua—A comm a-separated list of URLs for aggregate report delivery. These are typically 'mailto' URLs. This attribute is optional.

ruf—A comma-separated list of URLs for forensic/failure report delivery. These are typically 'mailto' URLs. This attribute is optional.

So a sample DMARC record for example.com might be:

v=DMARC1; p=quarantine; rua=mailto:dmarc_agg@vali.email; ruf=mailto:dmarc-reports@example.com

If you'd like a more complete discussion of the DMARC record syntax, please consult the RFC here.

WHAT IS SENDER POLICY FRAMEWORK (SPF)?

Sender Policy Framework (SPF) is an open, DNS-based email authentication system that allows sending domains to define which IP addresses are allowed to deliver email to receiving mail servers on behalf of the domain.

ON WHICH DOMAINS DO RECEIVING MAIL SERVERS LOOK UP THE SPF RECORD?

An email message contains a number of domains that could be used to look up the SPF DNS record. The SPF standard specifies that only two of these possible domains should be used to look up the SPF record. These domains are:

- **The domain specified in the Return-Path address**
- **The domain used specified in the EHLO / HELO SMTP command by the delivering mail server**

Typically, only the first of these two options is used.

It is important to understand that the domain in the human readable From address will not be used for SPF record lookup unless it matches one of the two domains listed above. This has important implications when SPF is being used for DMARC.

WHAT DOES AN SPF DNS RECORD LOOK LIKE?

Here's an example of a simple SPF record authenticating a single service:

v=spf1 include:_spf.google.com –all

Here's an example of a more complex record, with additional services and some dedicated IP addresses:

v=spf1 include:spf.protection.outlook.com include:mail.zendesk.com ip6:2001:db8::/32 ip4:203.0.113.6 –all

WHAT IS THE SPF DOMAIN LOOKUP LIMIT?

As part of evaluating whether an email message passes SPF authentication, a receiving mail server may have to make one or more DNS lookups. Typical situations where such a lookup might be required include:

- When evaluating an 'include' directive to pull in the SPF rule defined on another domain
- When checking an IP address against an 'a', 'mx', or 'ptr' directive—which require a A, MX, or PTR DNS lookup respectively to evaluate

To protect receiving mail servers from denial of service attacks the SPF standard includes a hard limit on the number of domain lookups such a server is permitted to make when evaluating whether an email message passes SPF. That limit is 10 lookups.

CAN YOU PROVIDE SOME EXAMPLES OF HOW TO CALCULATE THE DOMAIN LOOKUP COUNT FOR AN SPF RECORD?

Sure. Let's start with the SPF records published by Google for use by its customers.

Google asks its customers to include the SPF record _spf.google.com in their individual domain SPF records. Looking up the contents of this record in DNS we find:

v=spf1 include:_netblocks.google.com include:_netblocks2.google.com include:_netblocks3.google.com ~all

with each of the included _netblocks SPF records simply consisting of a list of IP addresses, and hence requiring no additional domain lookups.

Now consider a domain with an SPF record that only includes the Google-recommended record:

v=spf1 include:_spf.google.com ~all

will have a domain lookup count of four—one for the _spf.google.com lookup, and one for each of the netblocks.google.com, netblocks2.google.com, and netblocks3.google.com domains.

Now imagine you want to include the records provided by a number of other email service providers. The contribution to the domain lookup from each included record can be calculated in a manner analogous to the _spf.google.com. The total domain count of the resulting record is just the sum of the individual contributions.

So, for example, if mailgun.org includes two other domains in its SPF record, then the following record:

v=spf1 include:_spf.google.com include:mailgun.org ~all

will have a total domain lookup count of 7.

If a domain owner were then to add an 'a' and an 'mx' directive to the SPF, like so:

v=spf1 a mx include:_spf.google.com include:mailgun.org ~all

then each of these directives would contribute one to the domain lookup count, and the total count for this SPF record would be 9.

WHEN DOES THE SPF DOMAIN LOOKUP LIMIT MATTER?

Most email service providers publish an SPF record in DNS for use by their customers. These customer-centric SPF records frequently include other domains, or have additional directives that require domain lookups, so that their total contribution to the lookup count may be much larger than 1. Customers are then instructed to ensure that the SPF record published on their own domain has an 'include' directive referencing the domain on which the email service provider published their SPF record.

Unfortunately, this makes it very easy for domain owners to run into the SPF domain lookup limit. It is not uncommon, even when using 2–3 services, for sender domains to reach a domain lookup count that exceeds the SPF lookup limit.

It's also not necessarily obvious when the SPF domain lookup limit has been exceeded. SPF records consist of a set of rules, each one evaluated sequentially. So if a message is validated by one of the rules defined early in the SPF record, the message will authenticate even though the SPF record as a whole is broken. Messages which are intended to be authenticated by rules that appear later in the SPF record will fail, because the receiver will stop evaluating the record before it reaches those rules.

Domain owners should endeavor to limit the number of includes in the SPF records for their domains, to help ensure that they don't exceed this limit. And email service providers should not recommend or require that domain owners add an 'include' directive unless it is absolutely necessary.

WHAT DO I NEED TO DO TO AUTHENTICATE MESSAGES I SEND WITH SPF IN A DMARC-COMPATIBLE WAY?

What you need to do depends on how you are handling bounce messages.

I don't need to intercept bounce messages
In this case, just use the message's From address for the Return-Path address. Publish an SPF record with your mail servers that your customers can include in their domain's SPF record.

I need to intercept bounce messages, and I can use a customer's subdomain in my Return-Path address
In this case your customer has delegated a subdomain of their domain to you, using either a CNAME or NS record. You should use this subdomain in the Return-Path address for messages you send on their behalf. Publish an SPF record on the domain that includes the IP addresses of the mail servers sending email on the customer's behalf.

I need to intercept bounce messages, but I can't use a customer's subdomain in my Return-Path address
In this case SPF cannot be used to authenticate messages in a DMARC-compatible way. Instead you should use DKIM to authenticate your messages.

WHAT IS SENDER-ID, AND HOW IS IT DIFFERENT FROM SPF?

Sender-ID is a now obsolete alternative email authentication standard that is similar to the SPF standard but with a few crucial differences.

As part of its operation, Sender-ID reused DNS records originally defined for SPF and interpreted them somewhat differently. The most significant difference between SPF and Sender-ID is how the two standards determined which domain(s) to use to determine the SPF rule to apply to the message. Because the lookup process is different between the two standards, messages that authenticate with SPF may not authenticate with Sender-ID and vice versa, even though all of the DNS records are the same.

SPF, as described above, uses the domain from the Return-Path address to determine the DNS record containing the relevant SPF record. Sender-ID may use this address (called the mfrom in Sender-ID parlance), the human-readable From address (called the pra in Sender-ID parlance), or both to determine whether a given message is properly authenticated.

The implications of this difference in behavior has caused considerable confusion in the email authentication world.

DO I EVER NEED TO CONFIGURE SPF ON A DOMAIN THAT IS NOT USED BY THE RETURN-PATH?

No.

Because Sender-ID reused existing SPF records, there is a common belief that sending IPs needed to be authenticated by the domain in the From address, usually the customer's apex domain. As Sender-ID is obsolete, and no longer used by major receivers, this is no longer a good recommendation.

Many email service providers still recommend that their customers add an include directive to the apex domain SPF record, even if the email service provider is using another domain for the Return-Path. If you are such an email service provider you should revise your recommendations.

BUT ADDING A DIRECTIVE TO THE APEX DOMAIN SPF RECORD CAN'T HURT, SO WHY SHOULD I CHANGE OUR RECOMMENDATION?

Adding extraneous directives to a customer's apex domain SPF record can, and often does, break authentication for some of their email services.

Any domain-resolving directives in the apex domain SPF record, or pulled into the record via includes, will contribute to the number of domain lookups done by a receiver when evaluating SPF for a message. SPF has a hard limit of 10 lookups.

So adding an unneeded directive to the apex domain SPF record can push your customer's SPF record over the 10 domain lookup limit. This breaks authentication for the domain, and depending on the details, may break authentication for one of more services in use by the domain.

Inclusion in the apex domain SPF record should be reserved for internal systems and external services that use the apex domain in the Return-Path of the messages they generate.

DO I STILL NEED TO CONFIGURE DEDICATED SENDER-ID (V=SPF2.0) RECORDS?

Sender-ID records are DNS TXT records starting with 'v=spf2.0'. These should not be confused with SPF records, which start with 'v=spf1'.

As Sender-ID is now an obsolete standard, there is no longer any need to configure Sender-ID DNS records or to ask your customers to configure such records.

If you have any Sender-ID records currently set in DNS, you may wish to remove them.

WHAT IS DKIM?

DKIM is an open, DNS-based email authentication standard that uses public key encryption to authenticate email messages.

Message originators generate a hash of the email message, encrypt that hash using a private key, and include the encrypted hash (the cryptographic signature) as a header in the message. That signature is also associated with a DNS domain (included in the mail header with the signature).

The public key corresponding to the private key is published under a specific subdomain of the DNS domain used in the signature. Because only the domain owner (or someone they authorize) should be able to publish a DNS record for the subdomain, this associates the public key with the domain owner.

When a mail server receives a DKIM-signed message, it looks up the public key in DNS. The mail server uses the public key to decrypt the signature included in the email header. The mail server then computes a hash of the incoming message using the same algorithm as the message originator. If the computed hash matches the decrypted signature, then the message is authenticated. Otherwise, it fails DKIM authentication.

WHAT IS A DKIM SELECTOR?

To support many DKIM key records for a single domain, the DKIM standard introduced the notion of a selector. A selector is simply a sequence of one or more DNS domain labels. The selector is included as a field in the DKIM message header.

A receiving mail server uses the selector and the domain included in the DKIM message header to determine the exact domain to be used to retrieve the DNS record containing the public key. The DKIM record domain is constructed as follows: <selector>._domainkey.<header domain>.

For example, if a DKIM signature has a selector abcd.xyz and a domain of example.com then the corresponding DKIM record domain would be abcd.xyz._domainkey.example.com.

CAN I USE A SINGLE KEY, USING A SINGLE SHARED DOMAIN, TO AUTHENTICATE ALL OF MY CUSTOMERS?

No.

Using a single key on a shared domain can allow your messages to pass DKIM. But in today's world this is no longer sufficient, because DKIM authentication by itself does nothing to stop fraudulent emails. It's easy for phishers to use DKIM to authenticate messages against their own domains, while displaying a human-readable From address from another domain. This is precisely the abuse DMARC was designed to stop.

Because DMARC requires that a DKIM signature domain align with the message's From domain, if you are using a customer's domain in the From address, a shared domain DKIM signature will not allow the message to pass DMARC.

SO MY CUSTOMERS NEED TO INSERT ONE OR MORE DNS RECORDS TO ALLOW MY MESSAGES TO AUTHENTICATE WITH DKIM?

Yes.

If you wish to authenticate messages with DKIM then the public key record corresponding to the private (signing) key must be in DNS on an appropriate subdomain of your customer's domain. Your customer may insert this value directly as a TXT record, or can delegate to you or another third party via an NS or CNAME record.

CAN I SHARE DKIM KEYS BETWEEN CUSTOMERS, AS LONG AS THEY HAVE A DNS RECORD (TXT OR CNAME) WITH THE KEY IN THEIR DNS?

Yes, this is technically allowed. But it has serious security implications, and should probably be avoided.

When sharing a DKIM key between customers, an attacker who gets access to the signing key will be able to send fraudulent emails as any of your customers.

This decision requires even more consideration if your customers are using DNS TXT record for their DKIM key records, rather than CNAMEs to records that you are able to update. In this situation your customers will need to update or remove the DNS TXT records to block these fraudulent emails.

HOW FREQUENTLY DO I NEED TO CHANGE THE DKIM KEYS I'M USING TO SEND EMAIL ON BEHALF OF CUSTOMERS?

There is no hard and fast rule for how frequently you need to change (or 'rotate') the DKIM keys being used to sign messages. But, that said, frequent rotation of DKIM keys is the single most effective measure you can take to increase the security of DKIM authentication. So updating at least once a year, and preferably more frequently, is strongly advised.

IF I CHANGE MY DKIM KEYS, WON'T THAT BREAK MESSAGES THAT MAY BE IN TRANSIT AND HAVEN'T HIT THE RECEIVER'S INBOX?

Under some circumstances changing a DKIM key record can cause messages that were signed with the original DKIM key to fail DKIM authentication. This can happen if the DNS record with the public key is changed.

HOW CAN I AVOID BREAKING AUTHENTICATION FOR MY MESSAGES WHEN ROTATING MY DKIM KEYS?

The standard solution to this problem is to use multiple selectors for your DKIM signatures, and only change one at a time. This allows you to rotate DKIM keys without breaking any in-transit messages that may not have been processed by their receivers yet.

For example, you could have your customer set up two DKIM key DNS records in their domain, not just one, with selectors 's1' and 's2'. You sign messages with both 's1' and 's2' keys, generating two signatures on the message. When updating the 's1' DKIM key, that signature would fail to verify in some circumstances, but the 's2' DKIM signature would continue to verify, authenticating the message. At some later time, when updating the 's2' DKIM key, you would leave the 's1' DKIM key constant, and a similar process would occur.

WHAT KEY SIZE SHOULD I USE FOR MY DKIM KEYS?

At a minimum, you should be using 1024 bit keys. Google and some other receivers consider keys smaller than 1024 bits insecure, and will not use them for authentication.

Many cryptographers feel that a key size of 1024 bits will likely become insecure sometime in the next few years. You may wish to consider a larger key size (e.g. 2048), which is likely to be secure against cryptographic attacks for the foreseeable future.

WHAT ABOUT ADSP, SSP, AND OTHER EMAIL POLICY STANDARDS? ARE THEY RELEVANT FOR DMARC?

DMARC is largely intended to be a replacement for earlier standards that defined email authentication policy. As an email service provider you should not need to configure ADSP or SSP records, nor should you require your customers to do so.