# New developments in Email Security: DMARC/ARC and MTA-STS

Alexey Melnikov
alexey.melnikov@isode.com
Isode Ltd and also IETF ART Area Director

# Summary

- Protecting incoming email from phishing/spam
  - Background on SPF, DKIM, DMARC
  - Stats on DMARC uptake
  - Problems with DMARC
  - Experience doing DMARC workarounds in IETF
  - Introduction of ARC
  - What DMARC/ARC don't solve?
- DKIM crypto update
- Protecting mail transfer between organizations: MTA-STS

# How email works?

- RFC 5321 (SMTP) and RFC 5322 (Email Message Format)
- SMTP Envelope: who should receive bounces (Envelope FROM), who are the recipients?
- Messages contain headers, with From header field (who authored the email)
- Envelope FROM and From header field don't have to be the same
  - There are legitimate cases when a message is authored by one user and sent by another
  - Can be abused by spammers

# Protection from phishing/spam/fraud

- “phishing” ***- the fraudulent practice of sending emails purporting to be from reputable companies in order to induce individuals to reveal personal information, such as passwords and credit card numbers.***
- Phishing emails look like the real thing
- Might be hard for recipients to spot, especially if they are not technical
- Traditional anti-spam (like use of “spammy” words) doesn't work that great
  - SPF, DKIM, DMARC help to combat phishing

# SPF (1 of 2)

- Sender Policy Framework (RFC 7208)
- Sort of “reverse MX”: “Which SMTP servers can send email on behalf of a domain?”
- Published as DNS TXT records for <domain>, e.g.
  - "v=spf1 include:_spf.google.com ~all"
  - "v=spf1 ip4:64.233.160.0/19 ip6:2001:4860:4000::/36 mx ~all"

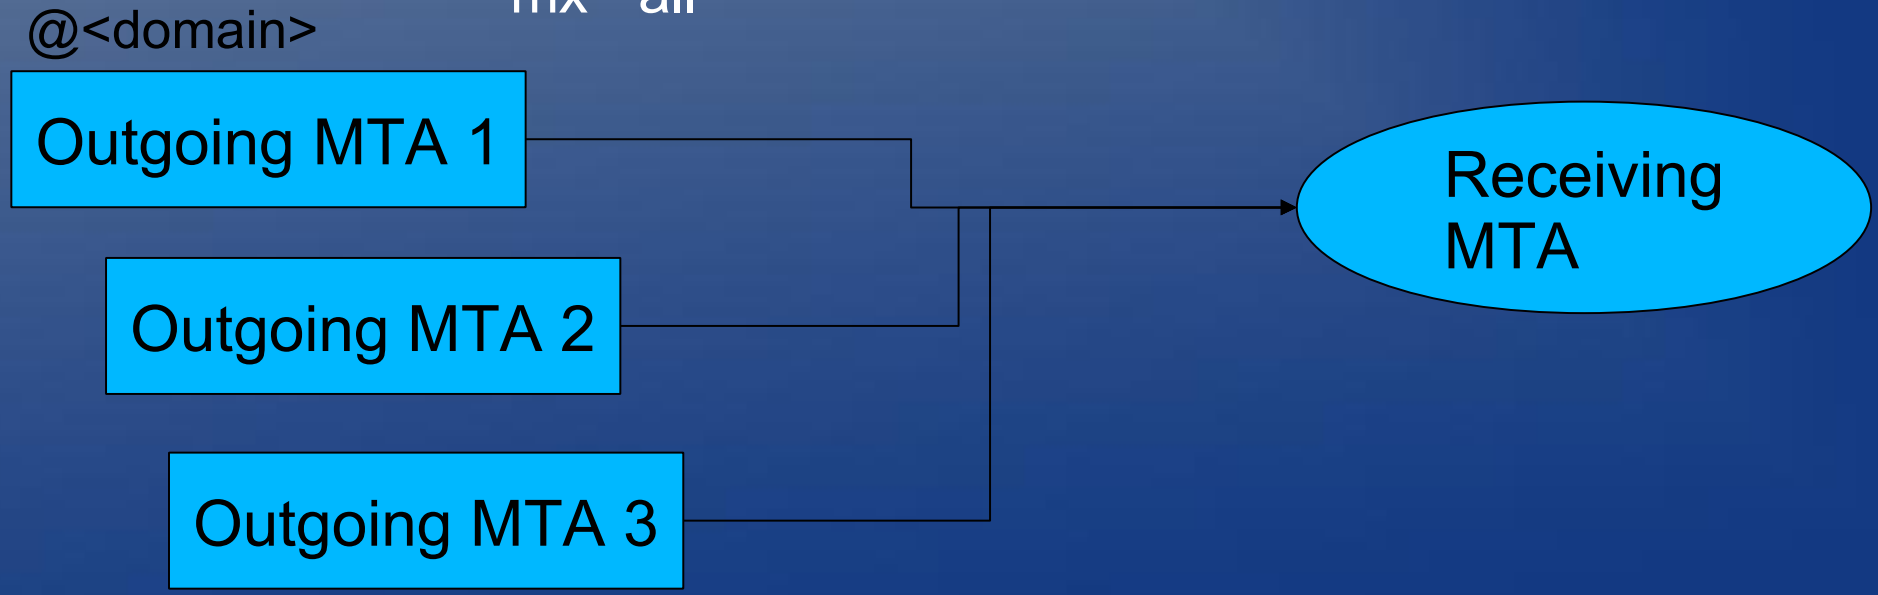

# SPF (2 of 2)

- When an SMTP server receives an email, it can lookup the SPF record and verify whether the message was sent by an authorized SMTP server.
- **Doesn't work with mailing lists/forwarders**

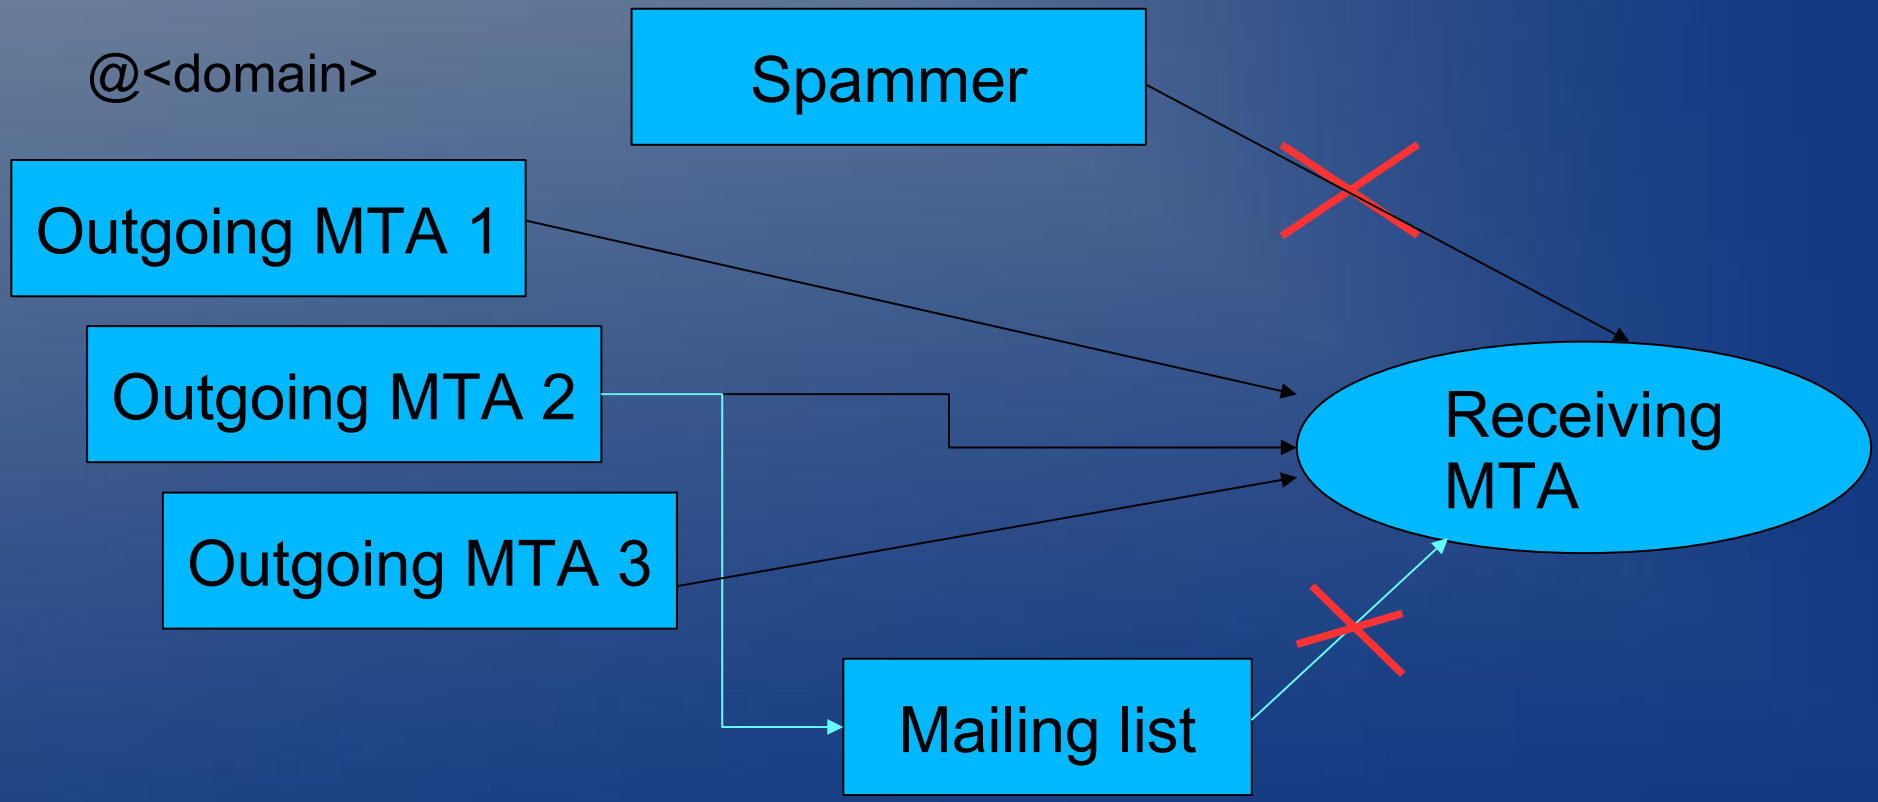

# DKIM (1 of 3)

- DomainKeys Identified Mail (RFC 6376)
- DKIM “permits a person, role, or organization that owns the signing domain to claim some responsibility for a message by associating the domain with the message. This can be an author's organization, an operational relay, or one of their agents.
- Specifies how to construct cryptographic signatures on selected email header fields
  - Prepended to the message itself
- Public keys for signatures are published in DNS
  - <selector>._domainkey.<domain> TXT records
  - Selector can be used for the whole domain or some specific users

# DKIM (2 of 3)

| | |
|---|---|
| From: alexey@example.com | ★ |
| To: boris@example.net | ★ |
| Accept-Language: en-GB, en-US | |
| Subject: Meeting to discuss project progress | ★ |
| Date: Fri, 1 Jun 2018 12:42:47 +0100 | ★ |
| Message-Id: <AD40307B-76A6-44B9-A1C8-6DFCECF7F5D1@example.com> | ★ |
| Content-Type: multipart/mixed | ★ |
| X-Mailer: iPhone Mail (15E302) | |
| Cc: boss@example.com | ★ |
| **Message Body** | ★ |

- Doesn't work with mailing lists/forwarders which change messages (e.g. if they add subject prefix)

# DKIM (3 of 3)

- Example DKIM-Signature header field:
  - **DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/simple; d=ietf.org; s=ietf1; t=1527437781; bh=KXuPpheci+050ZL55lsicVrBMnUO6NQNXRNExvYfh4A=;**

    **h=From:Date:To:Subject:List-Id:List-Unsubscribe:List-Archive:**

    **List-Post:List-Help:List-Subscribe;**

    **b=ZDTzQ66ll...**
- The corresponding DNS TXT record would be:
  - ietf1._domainkey.ietf.org
    - "k=rsa; p=MIGfMA0GCSqGSIb3DQ..."
- Doesn't work with mailing lists/forwarders which change messages (e.g. if they add subject prefix)

# DMARC

- DMARC (Domain-based Message Authentication, Reporting and Conformance)
  - DMARC policy is published as DNS TXT records
  - Authentication is done based on SPF and DKIM
  - (Independent piece) Reports are sent to the sending domain

# DMARC: policy

- Is published in DNS TXT record _dmarc.<domain>, e.g.
  - "v=DMARC1;p=reject;rua=mailto:d@rua.example.net, mailto:dmarc_rua@corp.example.com;ruf=mailto:d@ruf.example.net;fo=1;"

# DMARC: policy

| Policy type | Meaning |
| --- | --- |
| p=reject | Messages that fail DMARC policy get rejected (bounced) |
| p=quarantine | Messages that fail DMARC policy get quarantined. They don't get delivered to user's INBOX. |
| p=none | All messages gets delivered as usual. (Useful for getting DMARC reports) |

# DMARC: identifier alignment

- Alignment is how domain parts of Envelope FROM and From: header field identifiers are compared.
  - In the simplest case they should be the same

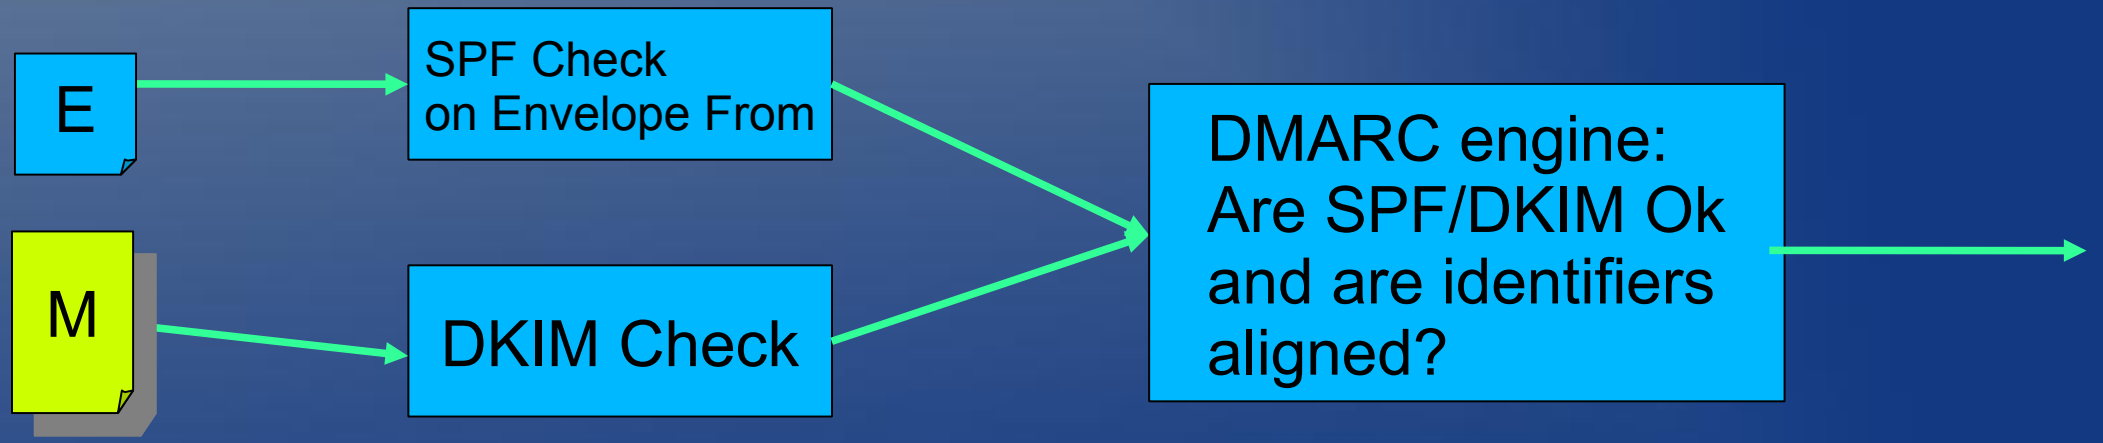

# DMARC: policy attributes

| Attribute | Description | Examples |
|---|---|---|
| v= | Version of DMARC policy | v=DMARC1 |
| p= | Policy (what to do with messages which fail the policy) | p=none<br>p=reject<br>p=quarantine |
| pct= | Percentage of messages subject to the DMARC policy | pct=0;<br>pct=10;<br>pct=100; |
| rua= | Where to send aggregated reports | rua=mailto:dmarc-aggr@example.com |
| ruf= | Where to send failure reports | ruf=mailto:dmarc-fail@example.net |
| adkim= | Alignment mode for DKIM | adkim=s<br>adkim=r |
| aspf= | Alignment mode for SPF | aspf=s<br>aspf=r |
| rf= | Reporting format | |
| | | |

# DMARC: reporting

- Aggregated reports, controlled by “rua” attribute
  - Delivered daily. XML or ZIPed XML
  - Help to spot SPF/DKIM/DMARC misconfigurations
  - Also help to know who is spoofing emails from your domain. Can be used for blocking them.
- Failure reports, controlled by the “ruf” attribute
  - Sent for each message that fails validation.
  - Can be lots of traffic!

# DMARC uptake

From 2016 to 2017, the number of DMARC records increased 3x, from 80K → 240K.

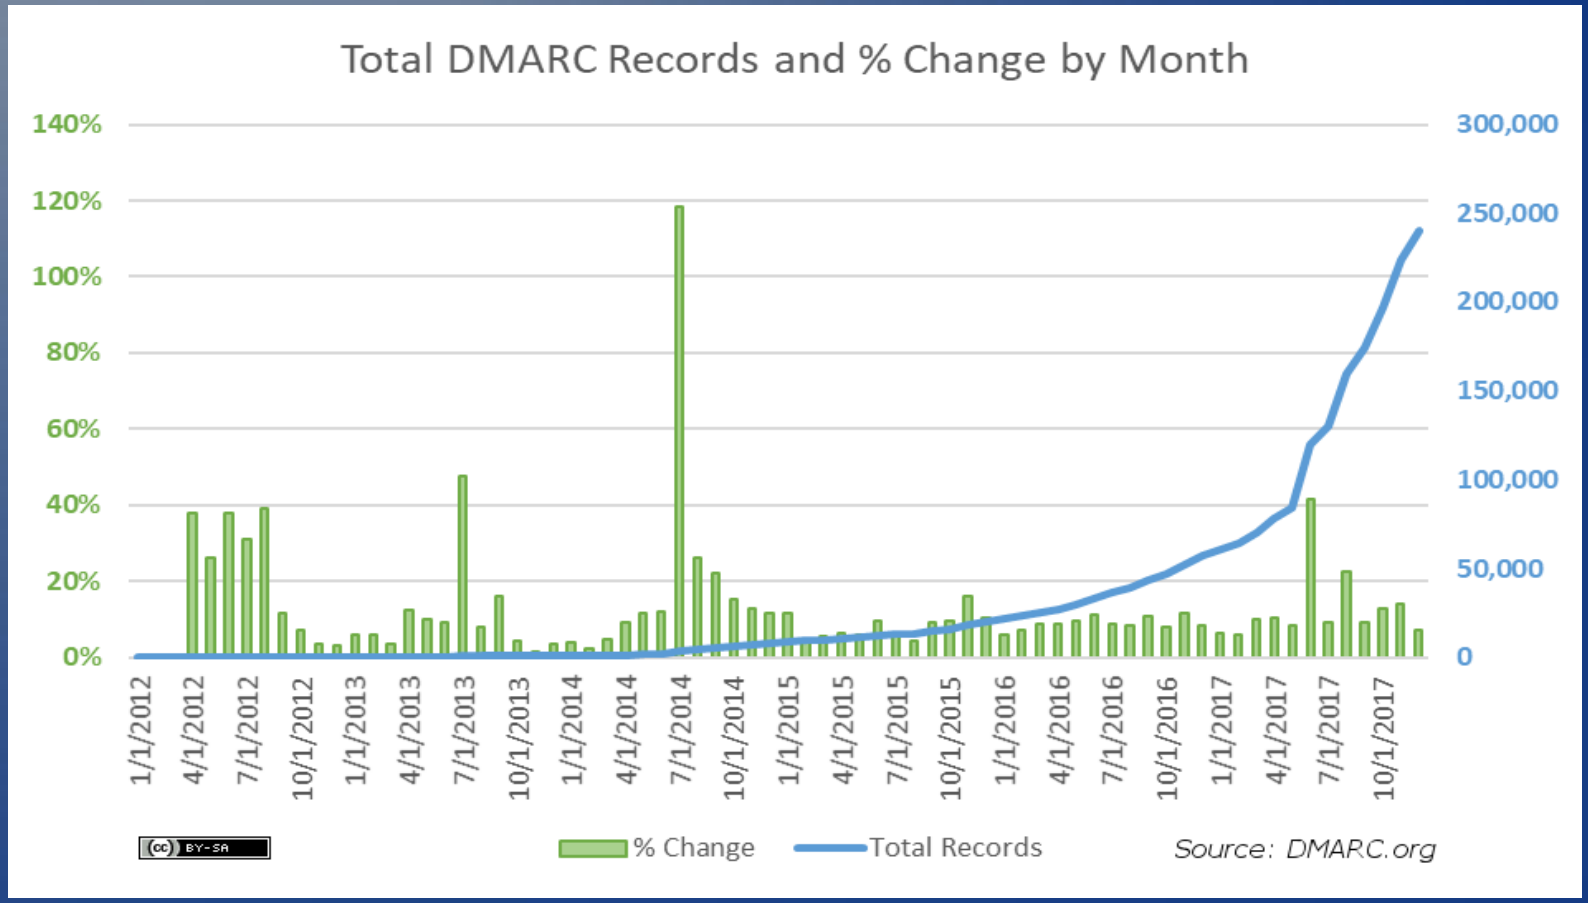

# DMARC uptake by country

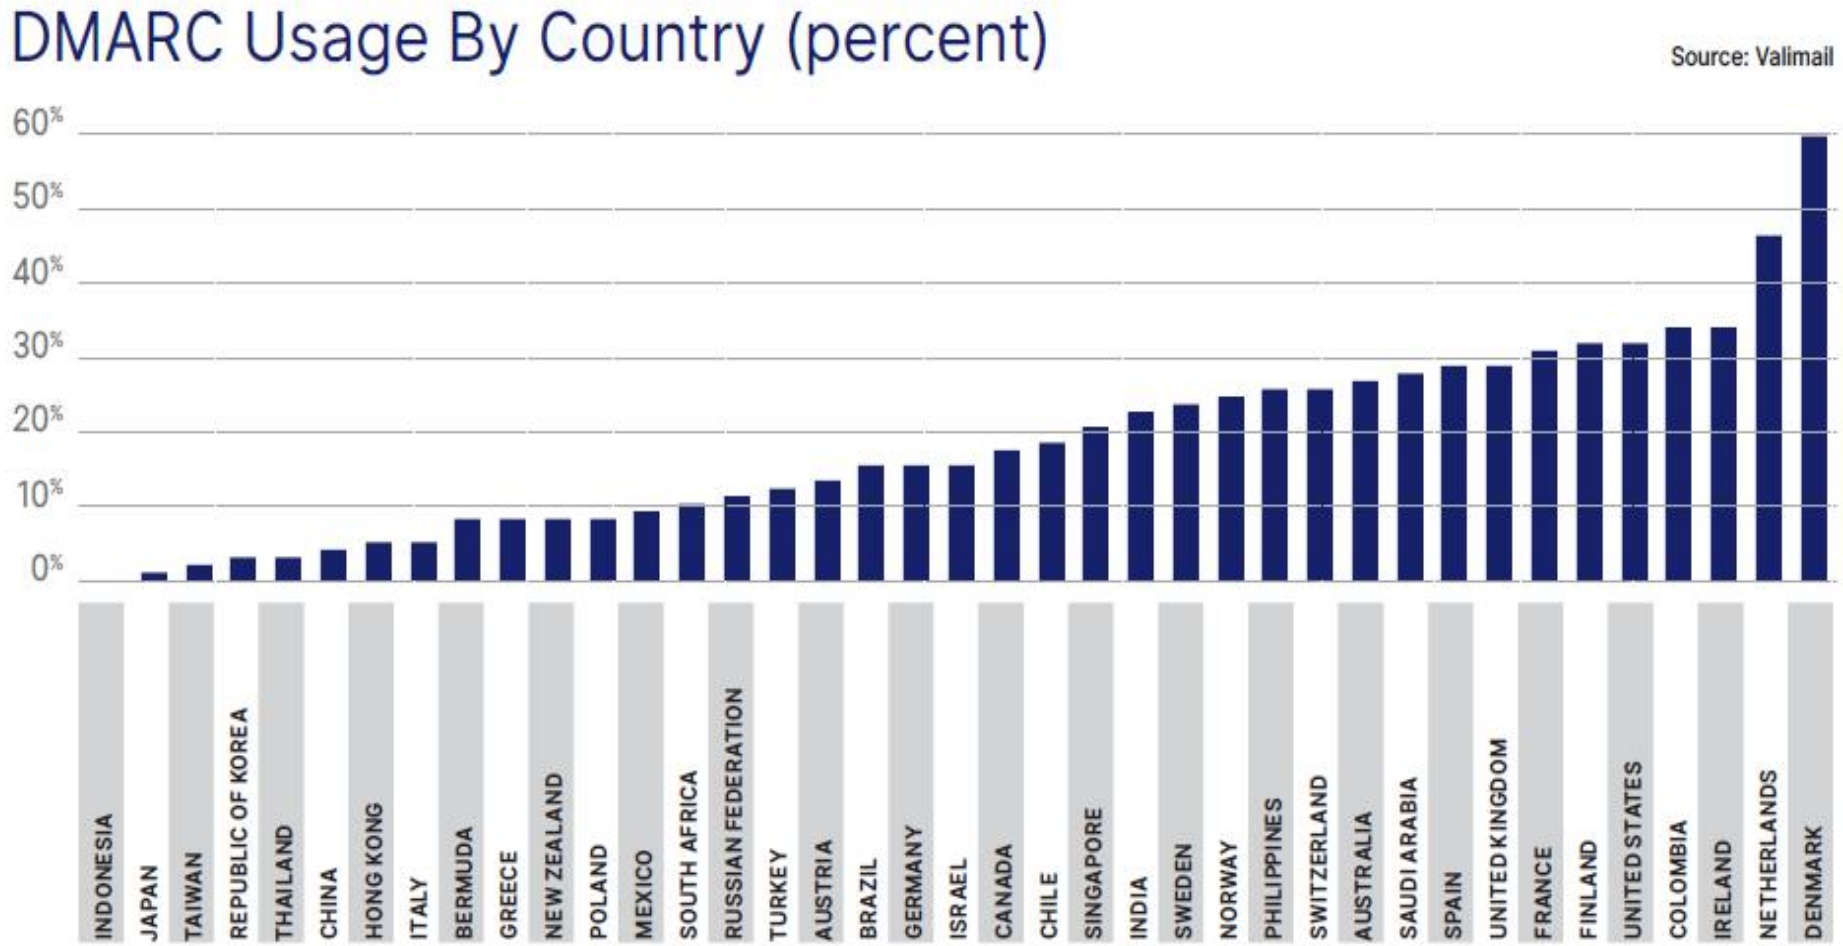

# DMARC uptake

- Statistics by selected countries (DMARC increase in 2017):
  - Europe: 2.25x increase overall
  - Australia → 2.4x
  - China → 2.8x
  - India → 3x
  - Russia → 2x (maybe more!)
- Trends
  - More DMARC use in consumer space (enterprises are lagging)
  - More DMARC use from big companies (e.g. big email providers), banks, government organization
    - “Brand” protection

# DMARC: How to deploy?

- Start with “p=none”
  - Start getting reports and look for misconfigurations
- Move to “p=quarantine”. Can start with small pct value (e.g. “p=quarantine; pct=10”) and increase it until it reaches 100
- Optional: switch to “p=reject”
  - Beware of indirect mail flows problem!

# Problem with DMARC

- Indirect mail flows: mailing lists, forwarders or filtering services
  - When a message from p=reject domain goes through a mailing list, it might not get delivered to some mailboxes who enforce DMARC policy, because SPF and possibly DKIM validation fails
    - Some emails get blackholed. People see partial conversations
  - Mailing list managers get DMARC related bounces from mailing list recipients that enforce DMARC policy.
    - Such recipient can get unsubscribed, if many emails from p=reject domain get sent in a short period of time. This happens because mailing list software can't distinguish between DMARC bounces versa other types of bounces

# DMARC and indirect mail flows

- Long term fix: ARC + reputation systems
- Short term fixes: updates to mailing list software to “mangle” emails so that they don't fail DMARC
  - Change emails from p=reject/p=quarantine domains so that their From header field comes from a domain with more relaxed DMARC policy.

# Experience doing DMARC workarounds in IETF

- Short term fix
- After discussing with IETF community, we settled on 2 possible solutions to be applied to email coming from p=reject domains
  - Emails from non p=reject/p=quarantine domains are not affected
  - Proposal 1: Replace From with a mapped @dmarc.ietf.org address
  - Proposal 2: Wrap messages inside message/rfc822 wrapper or multipart/mixed wrapper with From address that doesn't have p=reject policy. E.g. a mailing list related email address.

# Experience doing DMARC workarounds in IETF (proposal 1)

- p=reject From header field rewriting
  - Replace From with a mapped @dmarc.ietf.org address, e.g. alexey@example.com becomes alexey=40example.com@dmarc.ietf.org
  - dmarc.ietf.org domain publishes p=none policy
- Cons:
  - Addressbook “pollution” - hard to measure!
  - Need to maintain infrastructure for forwarding emails sent to mapped addresses, so that messages can get delivered to original recipients.

# Experience doing DMARC workarounds in IETF (proposal 2)

- Wrap messages inside message/rfc822 wrapper or multipart/mixed wrapper with From address that doesn't have p=reject policy. E.g. a mailing list related email address
  - Such messages appear as if they were “forwarded as attachments”
- Cons:
  - Messages from p=reject domains might appear as if they are forwarded (which might be ugly)
    - Broken email clients! Such messages are not always displayed correctly and sometimes can't be replied to.
    - Hard to measure how well this is supported in email clients

# ARC

- Longer term fix for the “indirect mail flows” problem
- ARC (Authenticated Received Chain): draft-ietf-dmarc-arc-protocol-14
- ARC allows each intermediary (e.g. mailing list or forwarder) to record state of DKIM/SPF verification on received messages and allow adding additional signatures
  - For example, a mailing list can re-sign with its own ARC signature

# ARC: How it works

- Each participating ARC intermediary adds a block of 3 header fields:
  - **ARC-Authentication-Results** (AAR) – results of SPF/DKIM/DMARC verification as observed by the intermidiary
  - **ARC-Message-Signature** (AMS) – similar to DKIM-Signature header field. Covers major header fields, whether or not they were modified by the intermediary
  - **ARC-Message-Signature** (AS) – simplified version of DKIM-Signature header field, which covers the newly added AAR and AMS header fields, as well as all AAR/AMS/AS added by previous hops
- DKIM code can be adopted for generation of AMS/AS

# ARC: Example

## Initial message header and header fields added by 1st MSA/MTA

```
ARC-Seal: i=1; a=rsa-sha256; t=1421363107;
    s=origin2015; d=d1.example; cv=none;
    b=pCw3Qxgfs9E1qnyNZ+cTTF3KHgAjWwZz++Rju0BceSiuwIg0Pkk+3RZH/kaiz61
     TX6RVT6E4gs49Sstp41K7muj1OR5R6Q6IlahLlQJZ/YfDZ3NImCU52gFWLUD7L69
     EU8TzypfkUhscqXjOJgDwjIceBNNOfh3Jy+V8hQZrVFCw0A=
ARC-Message-Signature: i=1; a=rsa-sha256; c=relaxed/relaxed;
    d=d1.example; s=20130426; t=1421363082;
    bh=EoJqaaRvhrngQxmQ3VnRIIMRBgecuKf1pdkxtfGyWaU=;
    h=MIME-Version:CC:Content-Type:Content-Transfer-Encoding;
    b=HxsvPubDE+R96v9dM9Y7V3dJUXvajd6rvF5ec5BPe/vpVBRJnD4I2weEIyYij
     rvQwbv9uUA1t94kMN0Q+haFo6hiQPnkuDxku5+oxyZWOqtNH7CTMgcBWWTp4QD
       4Gd3TRJlgotsX4RkbNcUhlfnoQ0p+CywWjieI8aR6eof6WDQ=
Received: ...
ARC-Authentication-Results: i=1; d1.example;
    spf=pass smtp.mfrom=jqd@d1.example;
    dkim=pass (1024-bit key) header.i=@d1.example;
    dmarc=pass
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/simple; d=d1.example; s=20130426; t=1421363082;
     bh=EoJqaa...
```

```
Message-ID: <54B84785.1060301@d1.example>
Date: Thu, 14 Jan 2015 15:00:01 -0800
From: John Q Doe <jqd@d1.example>
To: arc@example.org
Subject: [Lists] Example 1
Content-Type: text/plain
...
```

# ARC: Example

## Message goes through an MTA that doesn't support ARC

```
Received: from [10.10.10.131] (w-x-y-z.dsl.static.isp.com [w.x.y.z])
    (authenticated bits=0) by segv.d1.example with ESMTP id t0FN4a8O084569;
    Thu, 14 Jan 2015 15:00:01 -0800 (PST) (envelope-from jqd@d1.example)
ARC-Seal: i=1; a=rsa-sha256; t=1421363107;
    s=origin2015; d=d1.example; cv=none;
    b=pCw3Qxgfs9E1qnyNZ+cTTF3KHgAjWwZz++Rju0BceSiuwIg0Pkk+3RZH/kaiz61
     TX6RVT6E4gs49Sstp41K7muj1OR5R6Q6IlahLlQJZ/YfDZ3NImCU52gFWLUD7L69
     EU8TzypfkUhscqXjOJgDwjIceBNNOfh3Jy+V8hQZrVFCw0A=
ARC-Message-Signature: i=1; a=rsa-sha256; c=relaxed/relaxed;
    d=d1.example; s=20130426; t=1421363082;
    bh=EoJqaaRvhrngQxmQ3VnRIIMRBgecuKf1pdkxtfGyWaU=;
    h=MIME-Version:CC:Content-Type:Content-Transfer-Encoding;
    b=HxsvPubDE+R96v9dM9Y7V3dJUXvajd6rvF5ec5BPe/vpVBRJnD4I2weEIyYij
     rvQwbv9uUA1t94kMN0Q+haFo6hiQPnkuDxku5+oxyZWOqtNH7CTMgcBWWTp4QD
       4Gd3TRJlgotsX4RkbNcUhlfnoQ0p+CywWjieI8aR6eof6WDQ=
Received: ...
ARC-Authentication-Results: i=1; d1.example;
    spf=pass smtp.mfrom=jqd@d1.example;
    dkim=pass (1024-bit key) header.i=@d1.example;
    dmarc=pass
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/simple; d=d1.example; s=20130426; t=1421363082;
     bh=EoJqaa...
```

**[...]**

# ARC: Example

## Message arrives to an ARC-aware mailing list

```
ARC-Seal: i=2; a=rsa-sha256; t=1421363107; s=seal2015; d=example.org; cv=pass; b=pCw3Qxgf...
ARC-Message-Signature: i=2; a=rsa-sha256; c=relaxed/relaxed; d=example.org; s=clochette;
t=1421363105; ...
Received: from segv.d1.example (segv.d1.example [72.52.75.15]) by lists.example.org
(8.14.5/8.14.5) with ESMTP id t0EKaNU9010123 for <arc@example.org>; Thu, 14 Jan 2015
15:01:30 -0800 (PST) (envelope-from jqd@d1.example)
ARC-Authentication-Results: i=2; lists.example.org; spf=pass smtp.mfrom=jqd@d1.example;
    dkim=pass (1024-bit key) header.i=@d1.example; dmarc=pass
```

```
Received: from [10.10.10.131] (w-x-y-z.dsl.static.isp.com [w.x.y.z])
    (authenticated bits=0) by segv.d1.example with ESMTP id t0FN4a8O084569;
    Thu, 14 Jan 2015 15:00:01 -0800 (PST) (envelope-from jqd@d1.example)
```

```
[...]
[...]
```

# ARC: Example

## Message gets delivered to one of recipients on Gmail

```
ARC-Seal: i=3; a=rsa-sha256; t=1421363253; s=notary01; d=gmail.com; cv=pass;
 b=sjHDMriRZ0Mui5e...
ARC-Message-Signature: i=3; a=rsa-sha256; c=relaxed/relaxed;
    d=gmail.com; s=20120806; h=mime-version:content-type:x-original-sender...
Received: by mail-yk0-f179.google.com with SMTP id 19so2728865ykq.10
    for <mailbox@gmail.com>; Thu, 14 Jan 2015 15:02:45 -0800 (PST)
ARC-Authentication-Results: i=3; gmail.com; spf=fail
    smtp.from=jqd@d1.example; dkim=pass (1024-bit key)
    header.i=@example.org; dmarc=fail; arc=pass
```

```
ARC-Seal: i=2; a=rsa-sha256; t=1421363107; s=seal2015; d=example.org; cv=pass; b=pCw3Qxgf...
ARC-Message-Signature: i=2; a=rsa-sha256; c=relaxed/relaxed; d=example.org; s=clochette;
t=1421363105; ...
Received: from segv.d1.example (segv.d1.example [72.52.75.15]) by lists.example.org
(8.14.5/8.14.5) with ESMTP id t0EKaNU9010123 for <arc@example.org>; Thu, 14 Jan 2015
15:01:30 -0800 (PST) (envelope-from jqd@d1.example)
ARC-Authentication-Results: i=2; lists.example.org; spf=pass smtp.mfrom=jqd@d1.example;
    dkim=pass (1024-bit key) header.i=@d1.example; dmarc=pass
```

```
[...]
[...]
[...]
```

# ARC: How it can be used?

- Presence of a valid ARC chain (when all blocks of ARC header fields are syntactically valid and their signatures verify) is extra input for anti-spam engines if DMARC policy enforcement fails
  - So messages that were failed to get deliver using DMARC policy might get delivered by ARC-aware MTA
- Failed ARC chain can help to debug/find out which intermediaries cause breakage

# What ARC doesn't do?

- ARC depends on reputation of intermediaries
  - Valid ARC chain doesn't mean much without knowing whether intermediaries recorded in the chain are trusted
  - There is currently no standard way of sharing reputation scores
- Some remaining open questions (need deployment experience!)
  - What does it mean to have an ARC signature by an unknown mailing list?
  - Denial-of-Service attacks by injecting long ARC chains that take time to validate?
  - Spammers will inject fake ARC chains

# What phishers/spammers might do next/already doing?

- Because messages without DMARC/ARC might be treated as “more suspicious” by anti-spam system and would result in non delivery to recipients, this will force phishers/spammers to use hacked accounts so that sent messages don't trigger DMARC/ARC validation failures

# Crypto upgrade to DKIM

- RFC 8301: Cryptographic Algorithm and Key Usage Update to DKIM
  - Recommendations to stop using SHA-1 hashing and migrate to SHA-256
  - RSA Keys should be >= 1024 bits, 2048 bit keys are recommended
  - ***What happens with DKIM DNS records if the RSA key size gets even bigger?***
- draft-ietf-dcrup-dkim-crypto-09
  - Edwards-Curve Digital Signature Algorithm using the Curve25519 curve (ed25519), which has much shorter keys than RSA for similar levels of security

# SMTP Strict Transport Security and TLS reporting

- SMTP TLS Reporting (draft-ietf-uta-smtp-tlsrpt-22, approved for publication as an RFC)
  - Describes how to publish STARTTLS use reporting policy in DNS and format of reports
- SMTP MTA Strict Transport Security (MTA-STS) (draft-ietf-uta-mta-sts-19)
  - DNS is used to signal to always use STARTTLS when sending to a particular domain
  - A policy document is published over HTTPS

# SMTP TLS use reporting

- STARTTLS use reporting policy: _smtp._tls.<domain> DNS TXT record
  - _smtp._tls.example.com. IN TXT "v=TLSRPTv1;rua=mailto:reports@example.com"
  - or
  - _smtp._tls.example.com. IN TXT "v=TLSRPTv1; rua=https://reporting.example.com/v1/tlsrpt"
- Report multipart/report email containing a JSON or GZIPed JSON document describing different types of STARTTLS failures by sending IP/receiving MX

# MTA STS

- Protecting integrity and confidentiality of inter organizational email transfer

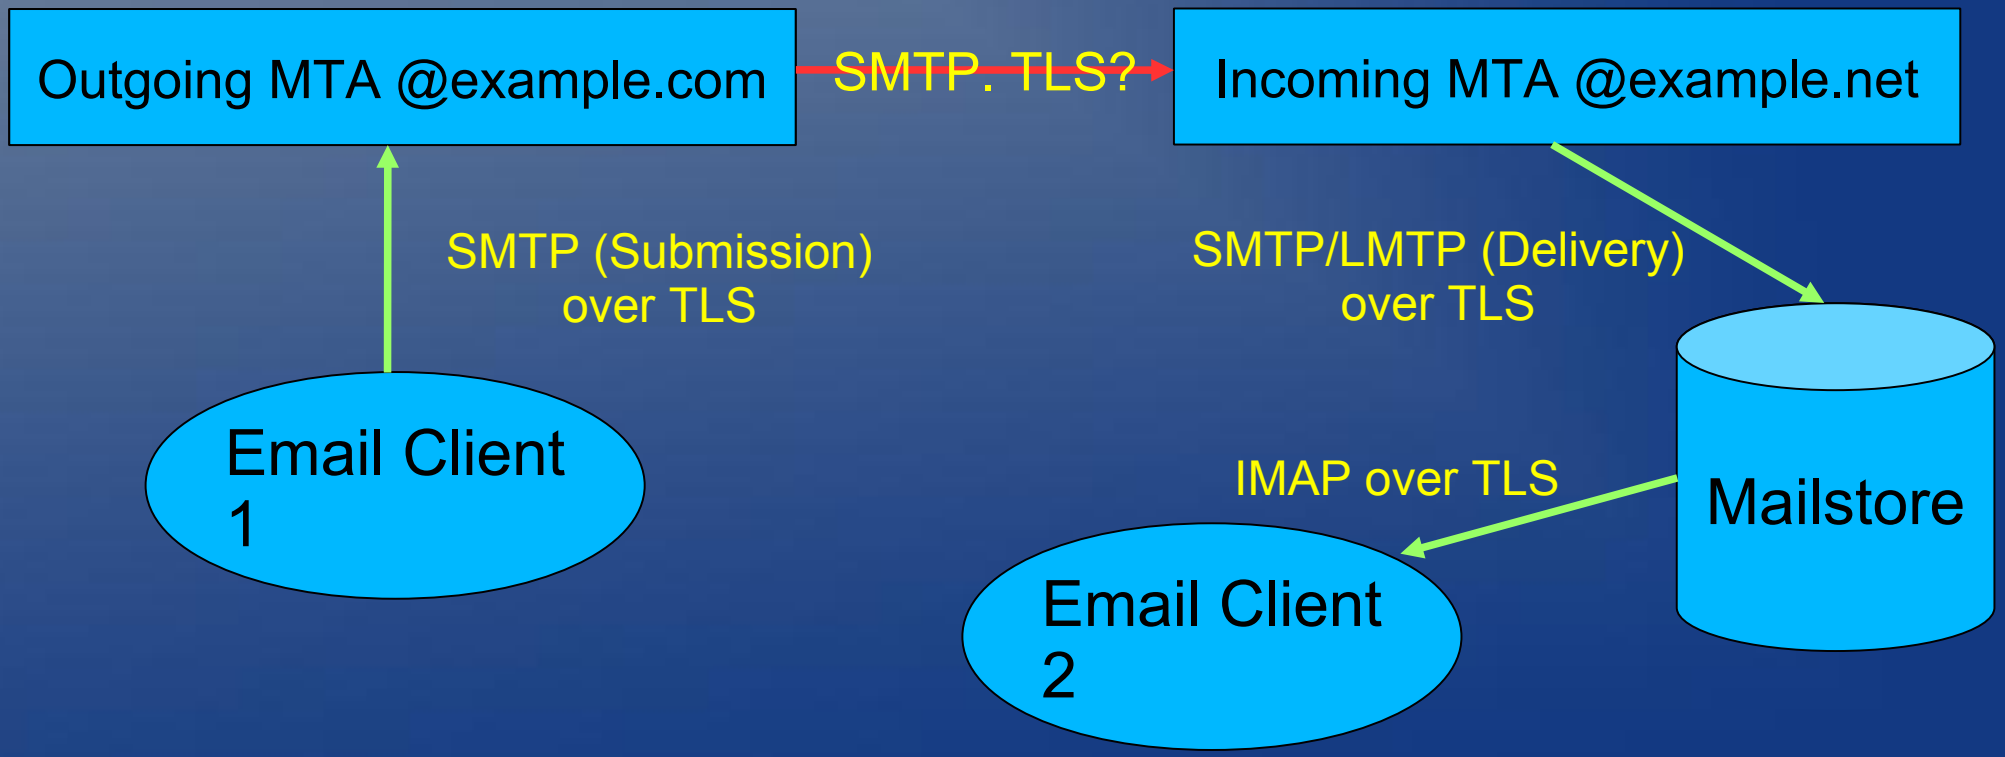

# MTA STS

- Signal in DNS
- Policy in HTTPS

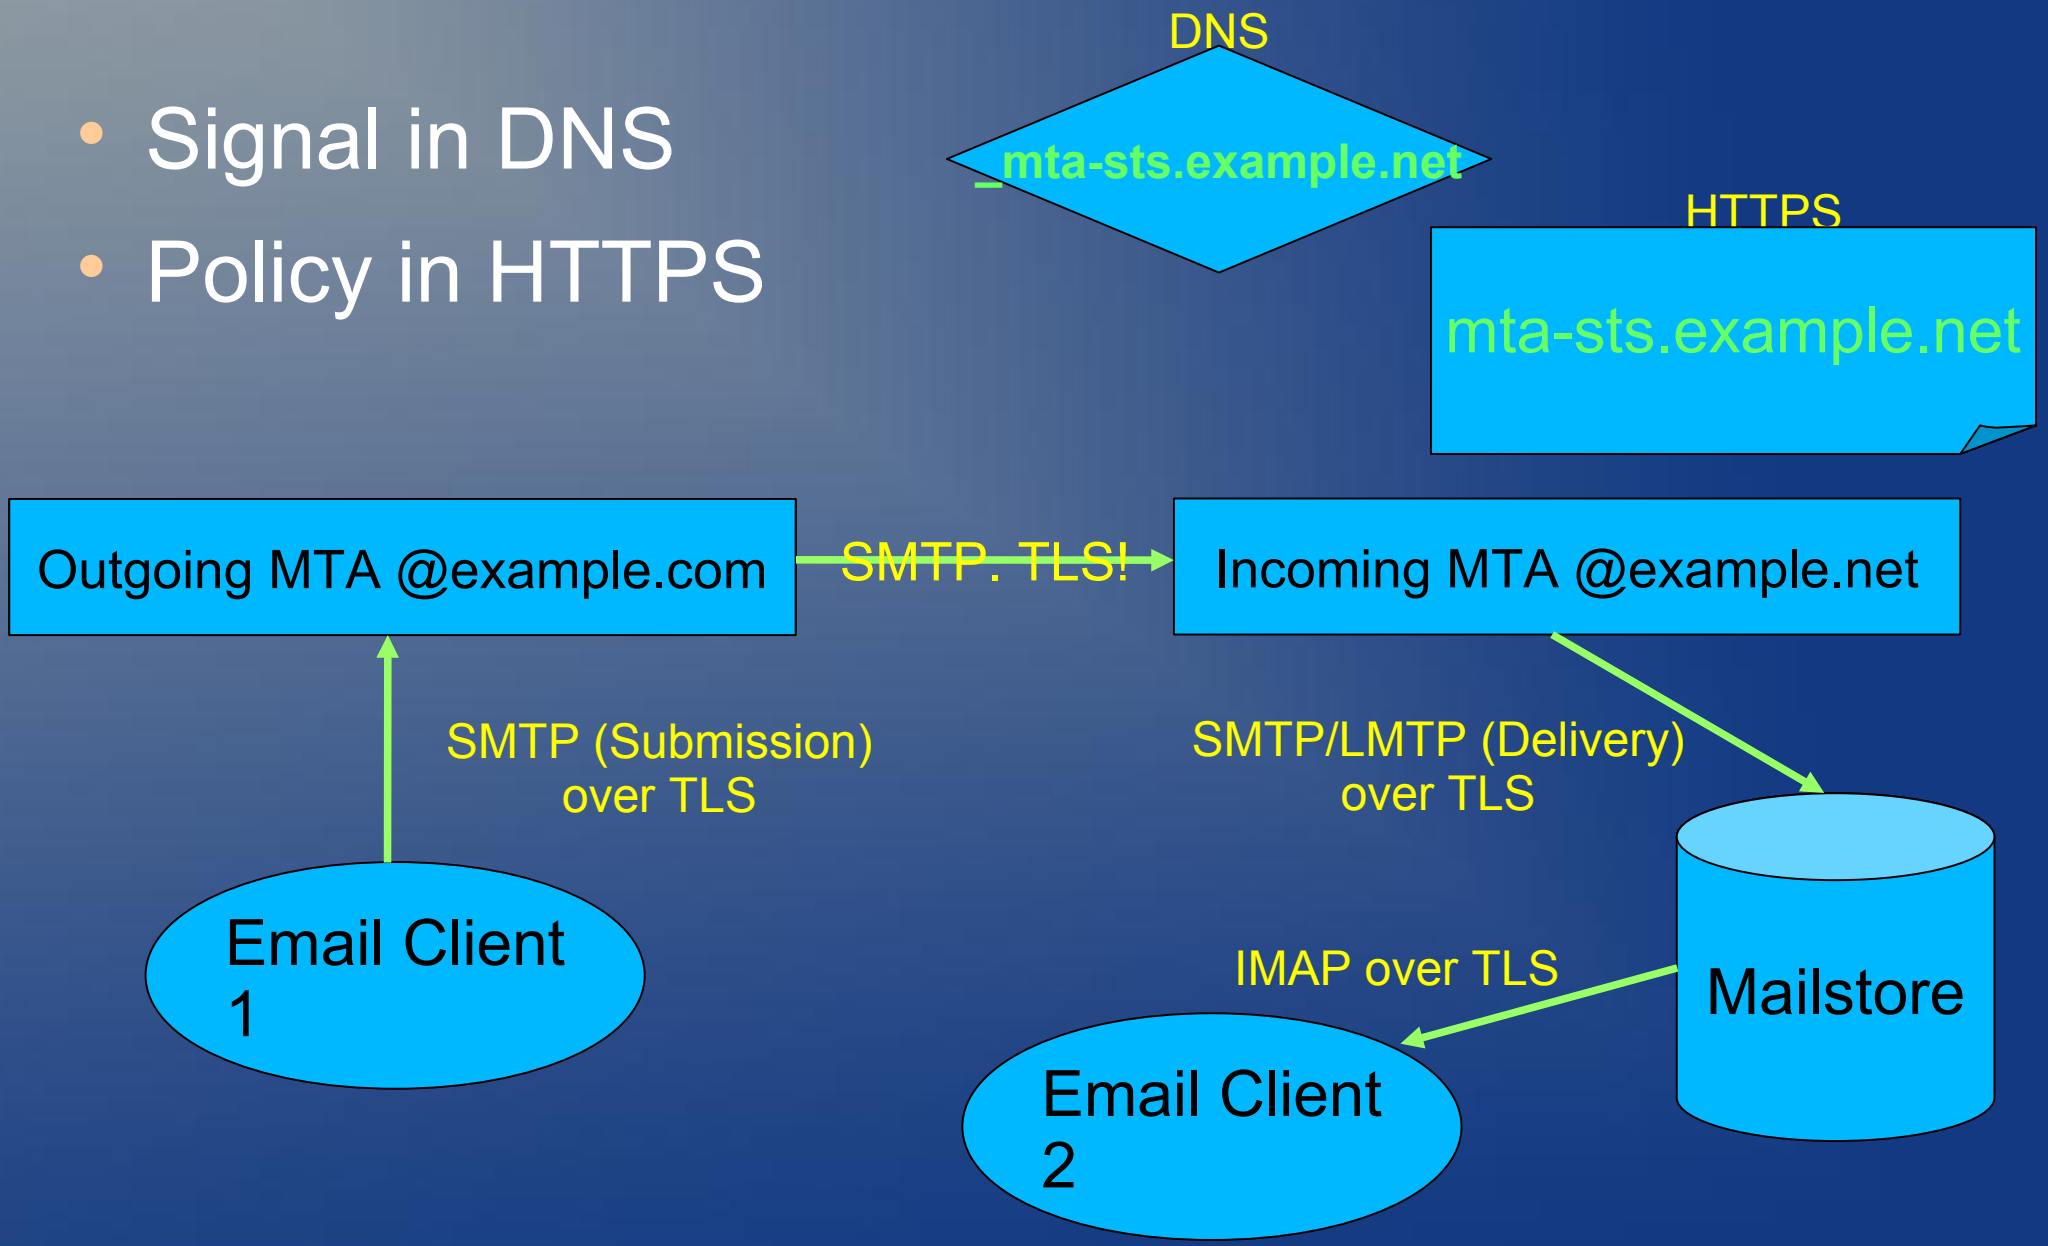

# How MTA STS works

- DNS TXT record
  - _mta-sts.<domain> TXT record, e.g.
  - _mta-sts.example.com.  IN TXT "v=STSv1; id=20160831085700Z;"
- Policy published on the web:
  - "https://mta-sts.<domain>/.well-known/mta-sts.txt"
  - For example:
  - version: STSv1
  - mode: enforce
  - mx: mail.example.com
  - mx: *.example.net
  - mx: backupmx.example.com
  - max_age: 604800

# Summary

- DMARC
  - Builds upon SPF and DKIM
  - Lets you see who sends email using your domain, and track/block unauthorized senders
  - With some policies helps to block all unauthorized messages from reaching your

    customers, partners, and employees
  - Doesn't work for indirect mail flows
- ARC
  - Helps to address indirect mail flow problem
- MTA STS
  - Helps to protect (with TLS) domain-to-domain email traffic
  - Helps to detect attacks redirecting email traffic

# Acknowledgements


- Valimail, in particular Seth Blank
- dmarc.org
- Participants of mailop@mailop.org mailing list

# Questions?

- Feel free to contact me at alexey.melnikov@isode.com