

Sender Authentication Technology Update

Steven M Jones

Executive Director, DMARC.org

Senior Software Engineer, LinkedIn

<http://linkedin.com/in/stevenmjones>

JPAAWG 2nd General Meeting

ベルサール飯田橋ファースト B1F

2019.11.14

Session A7, Hall A



Topics

- Why Do We Focus On Sender Authentication?
- Standards and Protocols
- DMARC Use Update
- Common Problems With DMARC Records

Why Do We Focus On Sender Authentication?

- Easier To Identify Legitimate Email
- Best Practices = Better Delivery
- Undelivered Mail = Wasted ¥
- Criminals Exploit Email Effectively
 - Phishing is #1 Cause - Data Breach
 - Business Email Compromise

BEC Scam Costs Media Giant Nikkei \$29 Million



Author:
Lindsey O'Donnell

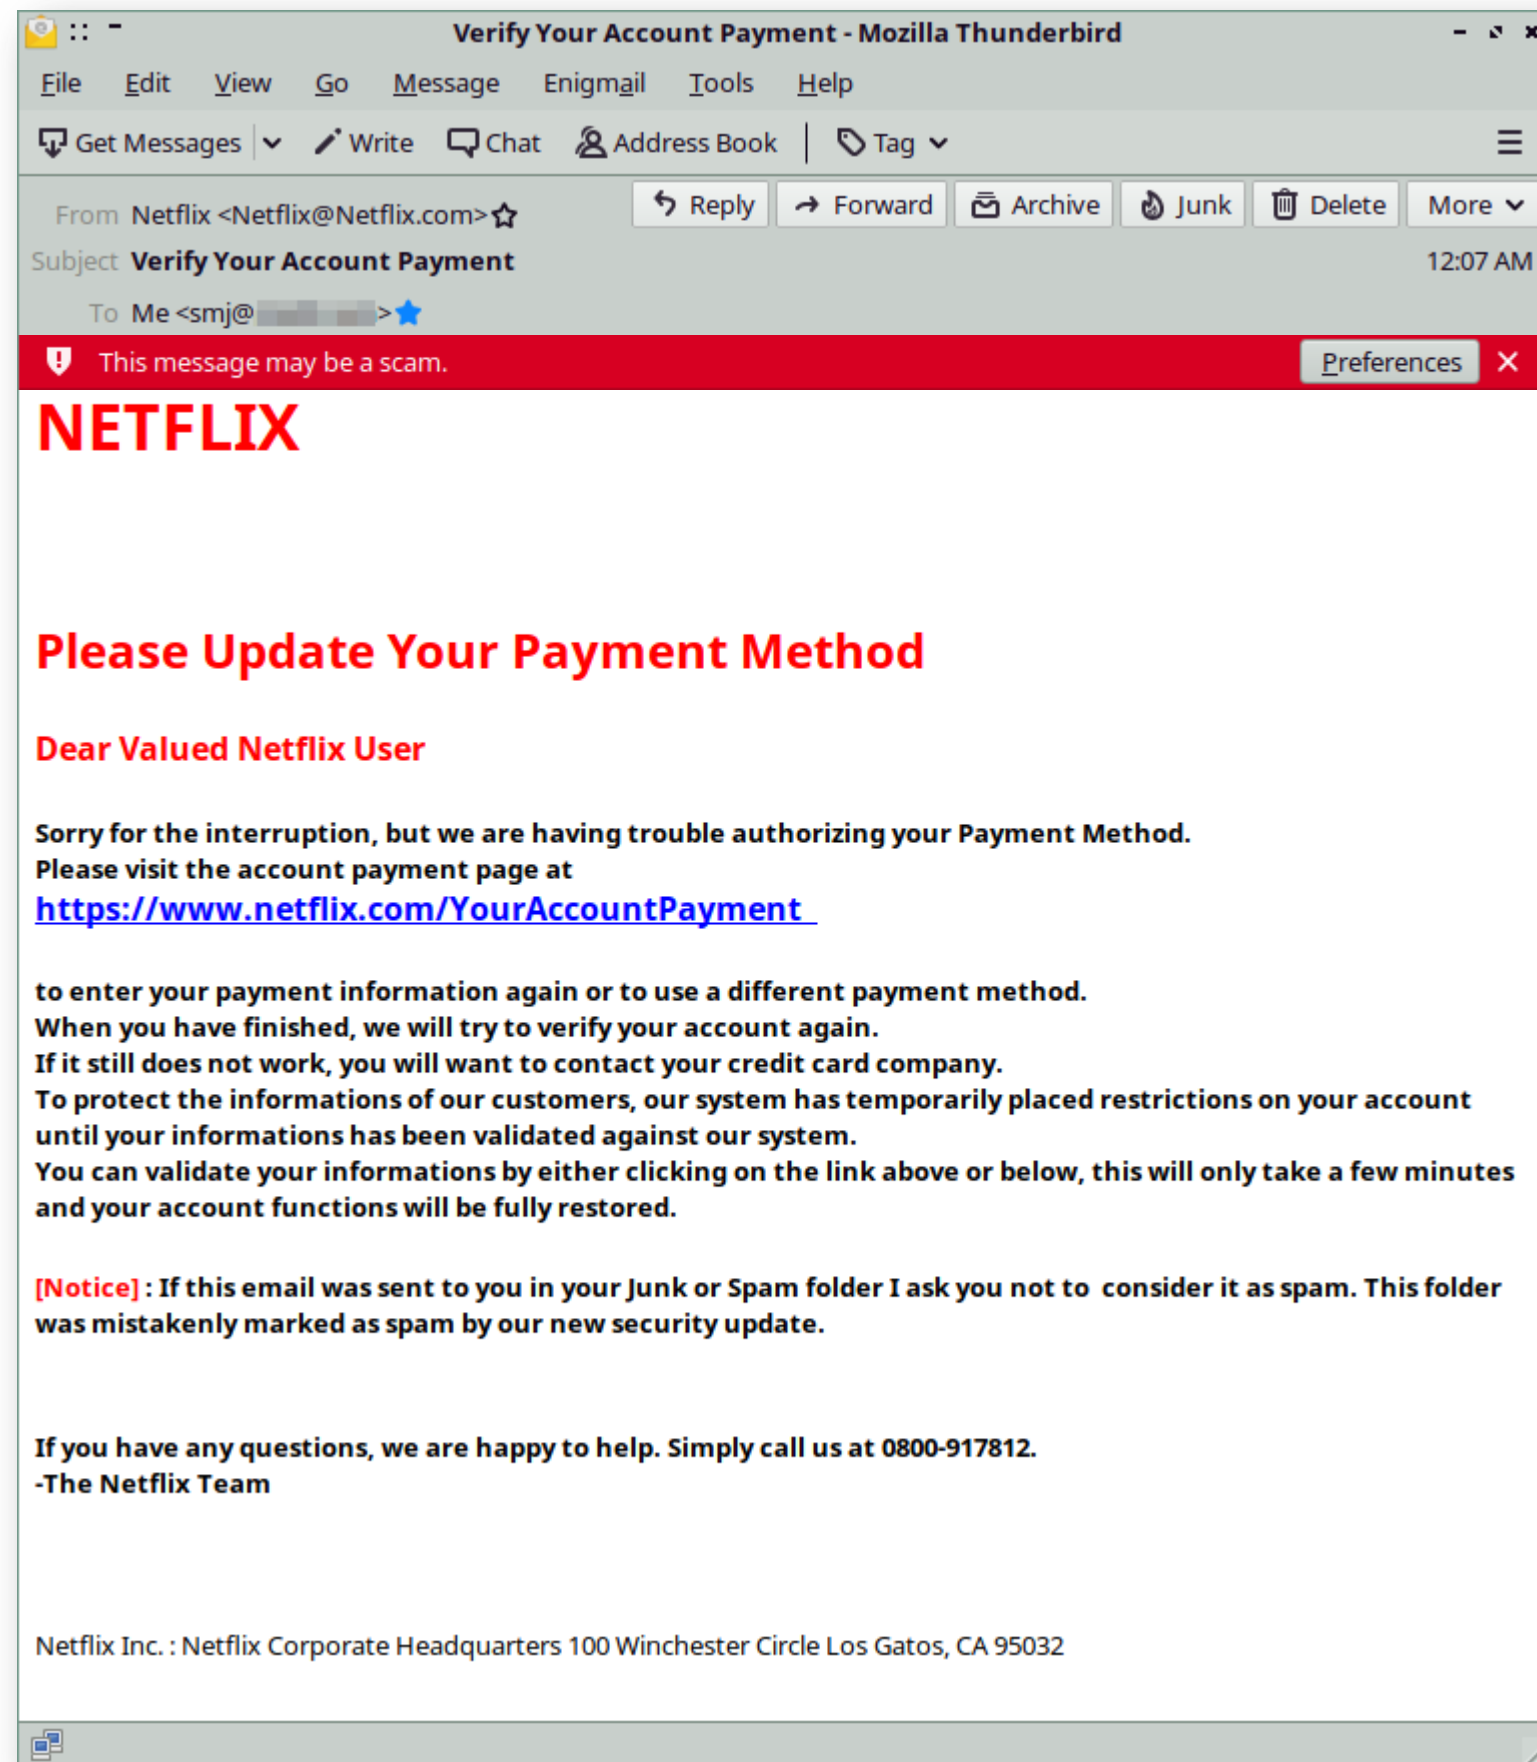
November 4, 2019
/ 10:34 am

3 minute read

In September, a Nikkei America employee transferred \$29 million to BEC scammers who were purporting to be a Nikkei executive.

Media conglomerate Nikkei Inc. has fallen victim to a business email compromise (BEC) scam that fleeced the company out of \$29 million.

Easier Detection, Better Protection



```
From: Netflix <Netflix@Netflix.com>  
Authentication-Results: XXX.XXXXXX.com/xACJ8inv058374;  
    dmarc=fail (p=reject dis=none) header.from=Netflix.com  
Authentication-Results: XXX.XXXXXX.com; spf=fail smtp.mailfrom=Netflix@Netflix.com  
Authentication-Results: XXX.XXXXXX.com; dkim=pass (2048-bit key; unprotected)  
    header.d=uttarauniversity.edu.bd header.i=@uttarauniversity.edu.bd  
header.b=kU8F/hqO
```

- Consistent authentication makes your legitimate email stand out, easy to model
- Machine Learning leverages this to detect cousin domains / “display name” attacks

Standards and Protocols



Overview of Common Protocols

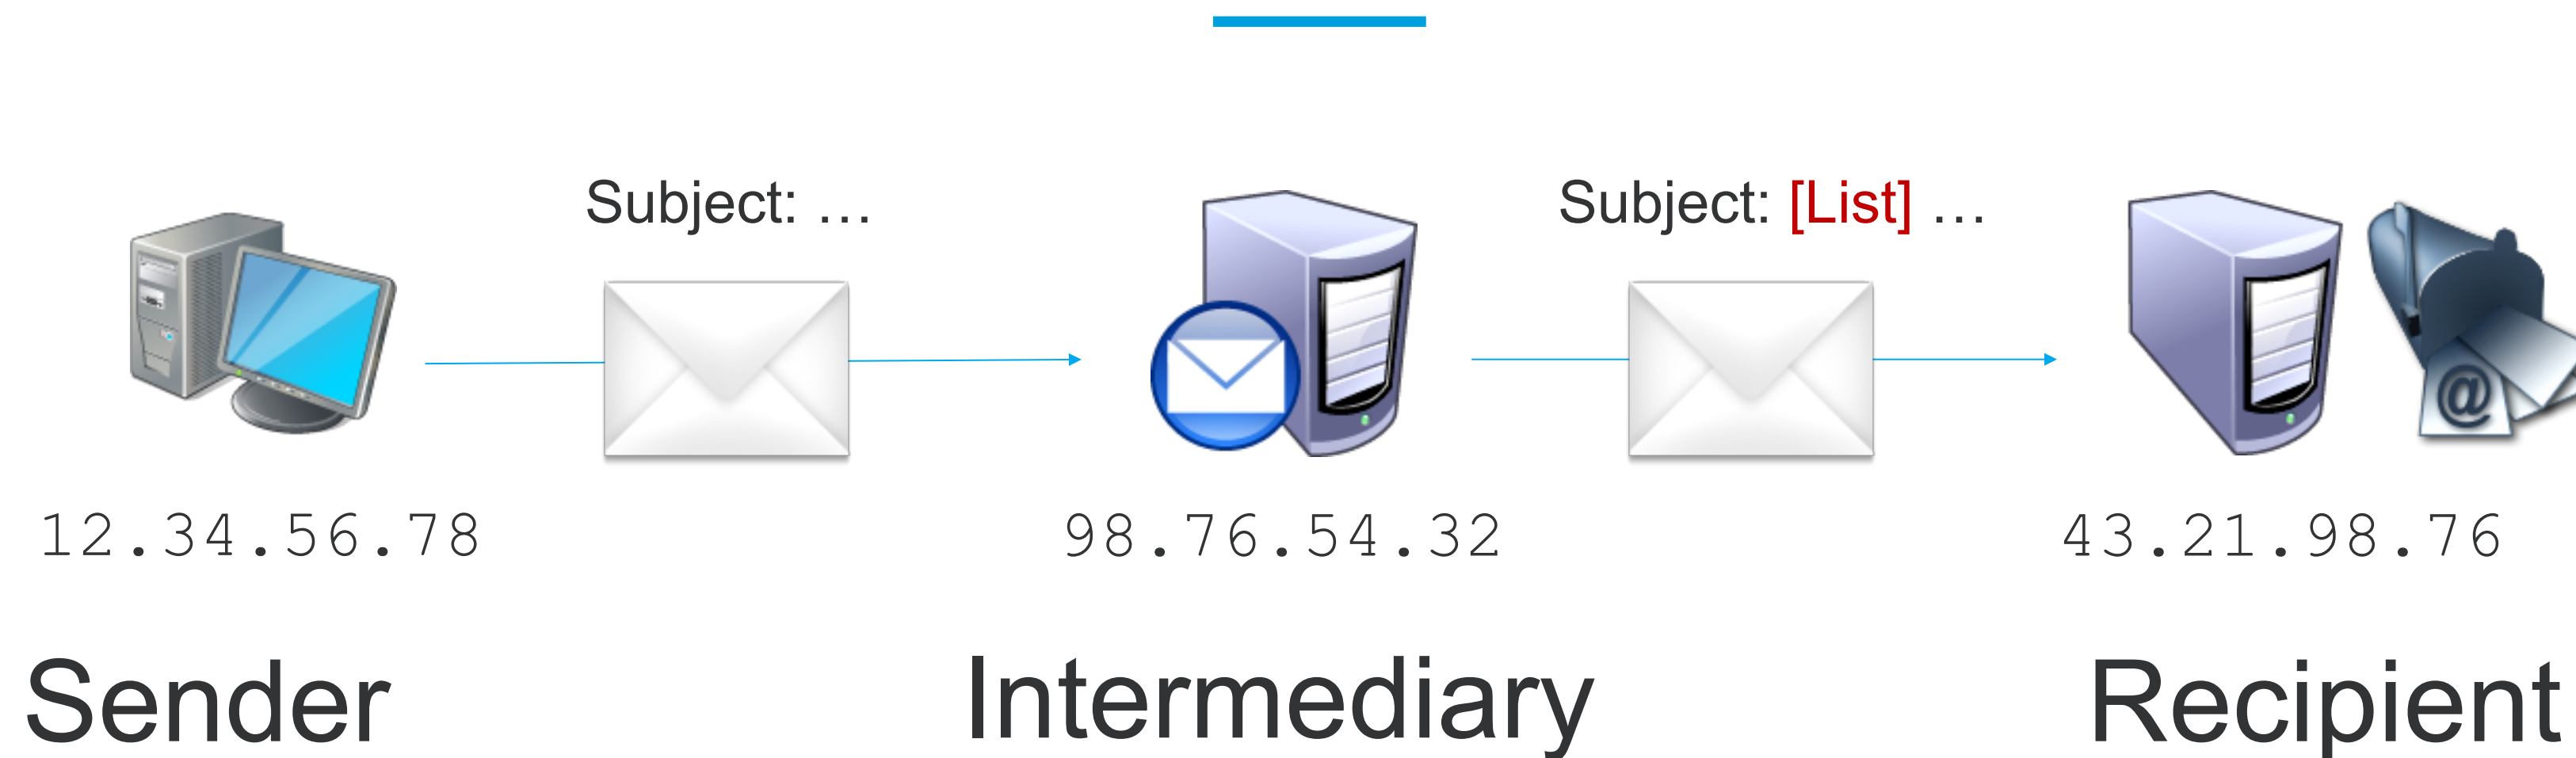


- Sender Policy Framework (SPF)
RFC 7208
- Domain Keys Identified Message (DKIM)
RFC 6376
- Domain-based Message Authentication, Reporting & Conformance (DMARC)
RFC 7489
- Authenticated Received Chain (ARC)

Refining Protections Over Time

SPF: Combat “backscatter” from spamming	2002 – 2004
<ul style="list-style-type: none">- Left header From: unprotected- Easily misconfigured, rarely enforced	
DKIM: Protect header From:, message forgery	2004 – 2007
<ul style="list-style-type: none">- No accepted policy mechanism- Third-party signatures problematic	
DMARC: Has policy mechanism, enforced at ISP	2009 - 2015
<ul style="list-style-type: none">- Cousin domains and “display name” attacks- Problems with mailing lists, forwarding	

Example of an Indirect Mail Flow




- Intermediary sends the message from a new IP address, causing SPF to fail to verify for Sender's domain
- Intermediary changes the message contents (subject:), causing Sender's DKIM signature to fail to verify

ARC Protocol

- ARC assists with authentication of “indirect mailflows”
- Under development since 2014
- Part of the IETF DMARC Working Group since 2016
- 8+ Interoperability testing sessions
- Draft of Usage Guide / FAQ available



ARC Published As RFC 8617 on 2019.07.09



ARC Implementations

- FastMail, Google, Microsoft – hosted email services
- Cloudmark, Halon, MailerQ and MessageSystems (SparkPost) – Mail Transfer Agent (MTA)
- Mailman and Sympa - Mailing List Manager (MLM)
- Free Software – dkimpy, Mail::DKIM, OpenARC
- More at arc-spec.org → Resources



Microsoft and ARC

- 2019.10.24 – Announces ARC support on Microsoft 365 Roadmap
- Testing in May 2019
- Began using production key in July
- Messages from many Office 365 tenants sent with ARC headers since July



Looking For Users of ARC

- ARC supports mailing lists – look there
- `arc-discuss@dmarc.org` mailing list
 - First message 2018.01.31 from an OpenARC user
 - 8.6% of posts have included an ARC Seal
- IETF's `ietf@ietf.org` mailing list
 - First message 2019.06.25 from Office 365 customer
 - 4.7% of posts have included an ARC Seal

RFC 8616: Email Authentication for Internationalized Mail

- Use of Unicode characters in domains and email addresses has been evolving
- RFC 8616 updates the core SPF, DKIM and DMARC specifications to clarify which form of Internationalized Domain Name (IDN) each uses
- Published on 2019.06.30

名がドメイン.co.jp


xn--v8jxj3d1dzdz08w.co.jp

DMARC and Public Suffix Domains

- Allow for DMARC to be applied at ccTLD, like `.uk` or `.jp`
- Also cover intermediate domains, ex. `gov.uk`
- Allow TLDs to have a DMARC policy for *non-existent* domains, ex. `nodomain.gov.uk`
- Proposed at M³AAWG 44 (Brooklyn) in 2018.10
- Several revisions in the IETF DMARC Working Group
- Nearing publication (as of November 2019)

Cryptography Changes From 2018

Changes in DKIM Cryptography (RFC 8463)



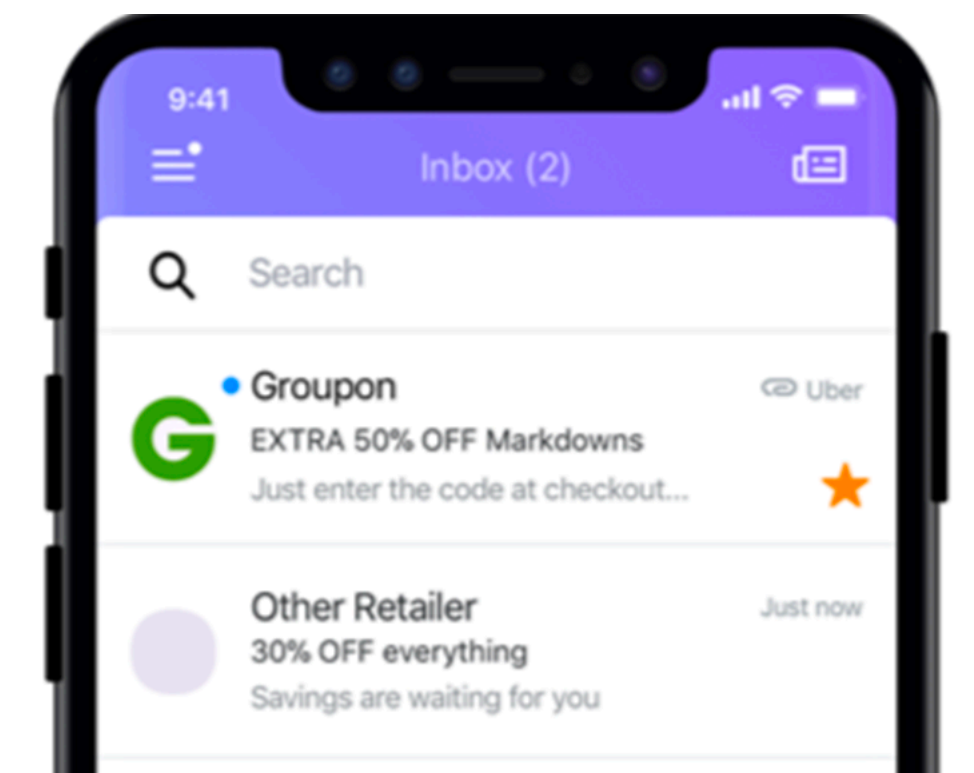
- RSA algorithm was deprecated under RFC 8017
- Elliptic Curve signing algorithm standardized under RFC 8031
- DKIM may now use PureEdDSA variant Ed25519
- Smaller keys for equivalent strength

Quantum Computing and Encryption

- 2019.10.23 – Google claims “Quantum Supremacy”
- What are the implications for traditional cryptography
- M3AAWG 46 (Montreal) had sessions on this topic
- Impacts most online activity, communications
- Directly impacts DKIM and ARC; indirectly DMARC
- How quickly can the IETF address this issue?

BIMI

- Brand Indicators for Message Identification (BIMI)
- Email clients would show sender's logo with messages
- Entrust Datacard issued first Verified Mark Certificate (VMC) in September 2019
- Yahoo US running a trial; Google in 2020
<https://www.brandindicators.org>



DMARC Use Update



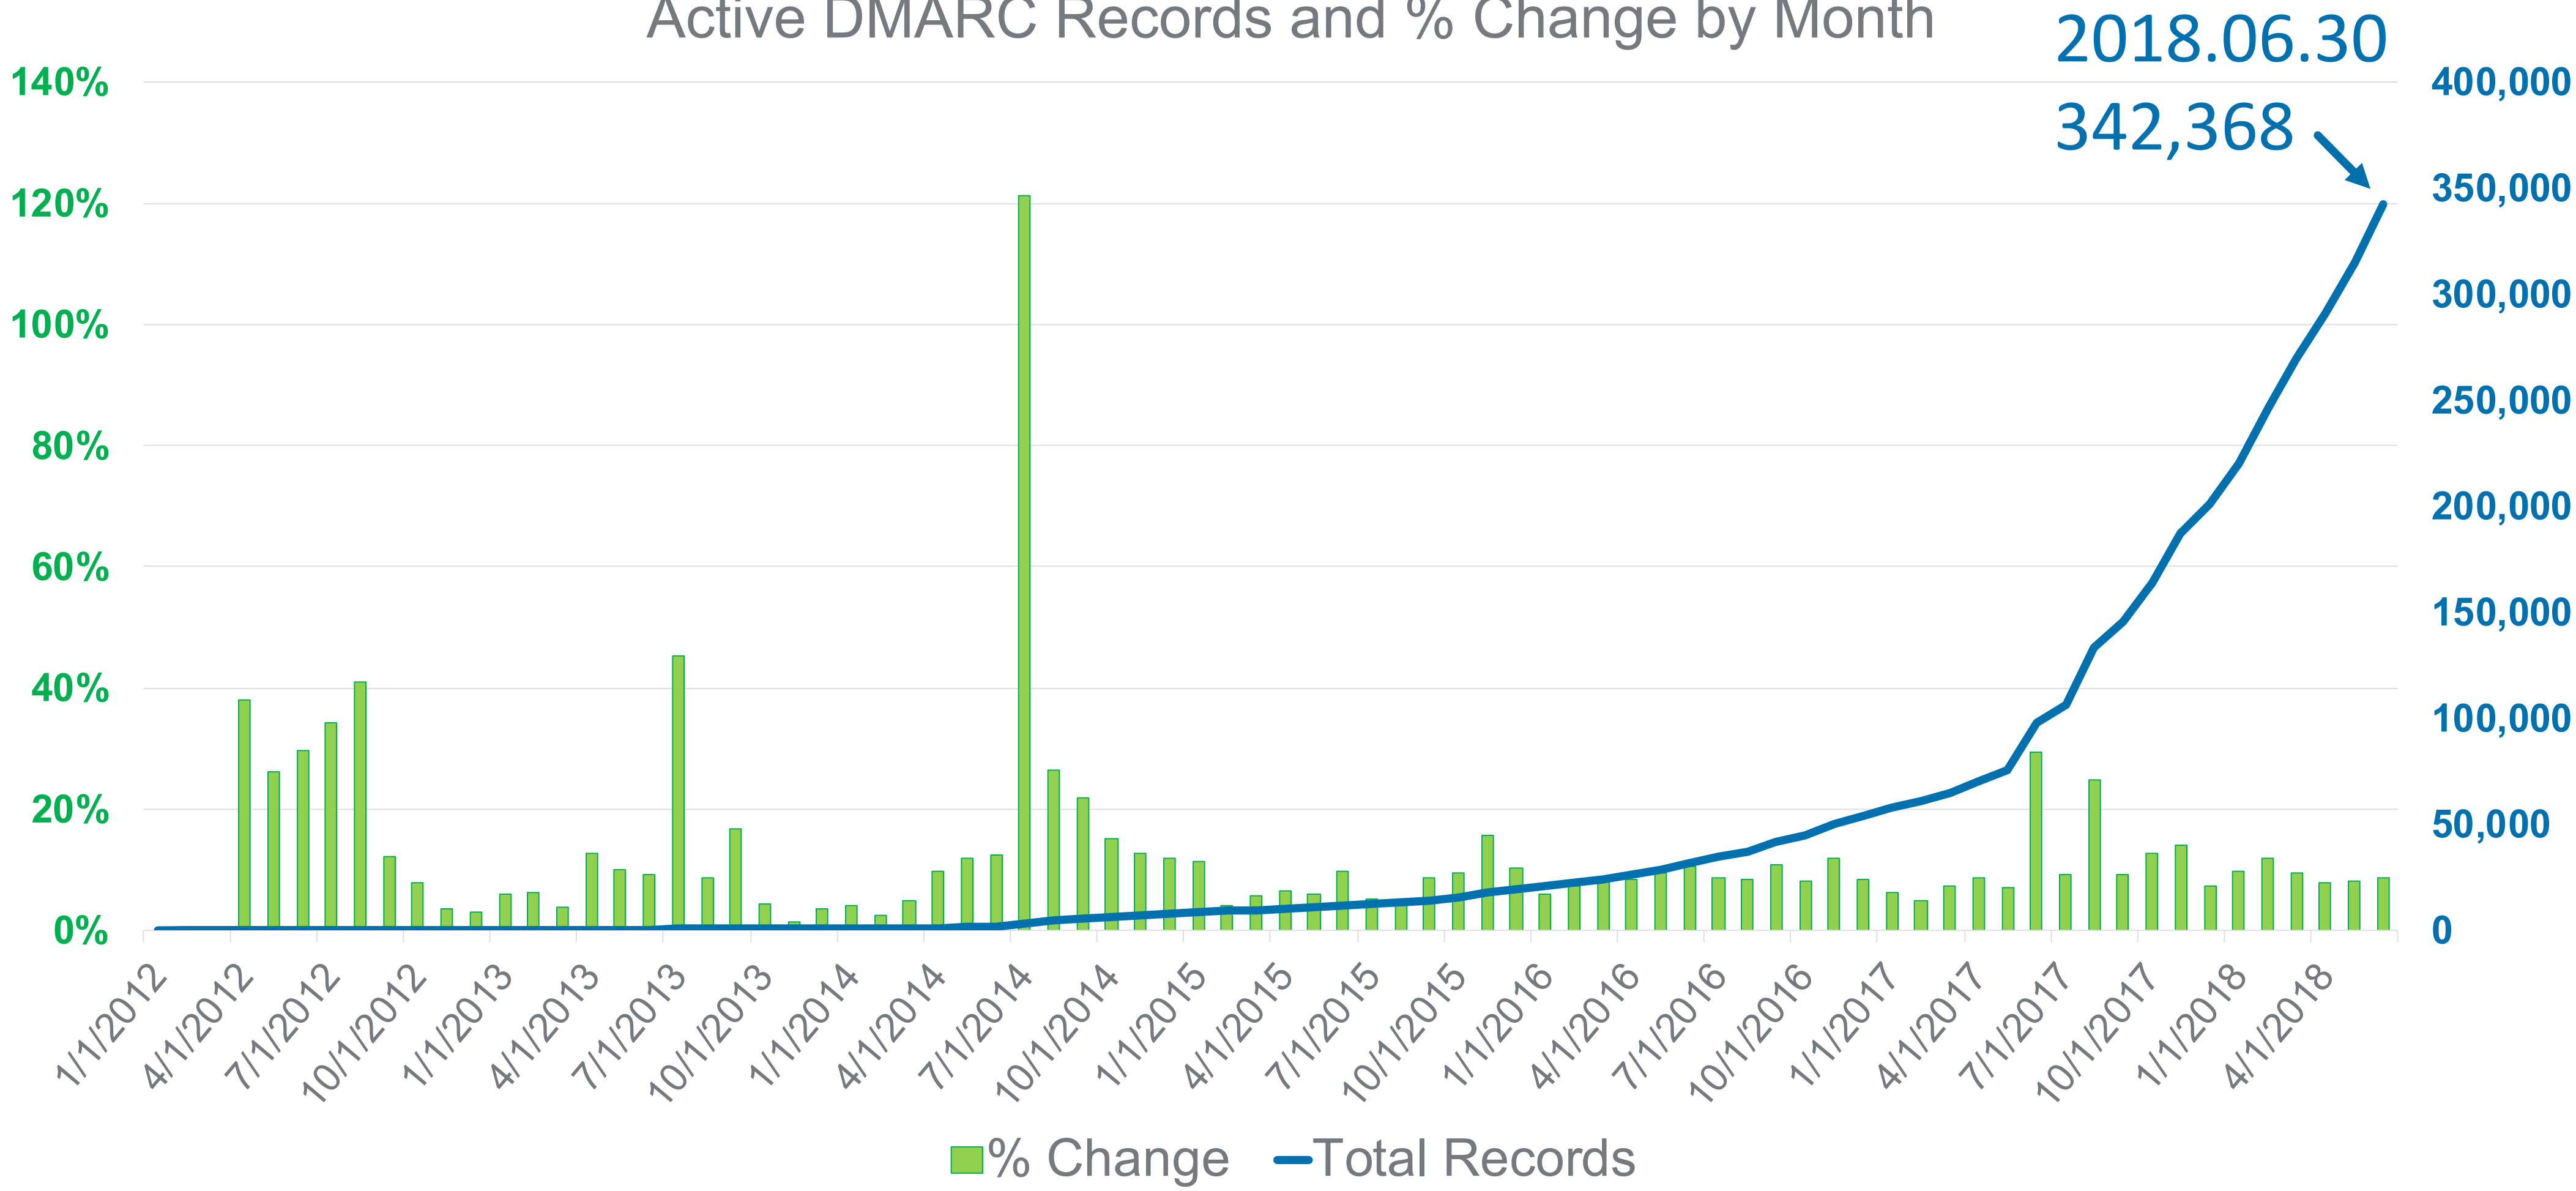
Farsight Security DNS Data

- Sensors located at network providers around the world
- Response data – the answer – is timestamped and stored
- Sensors only see records when somebody looks them up

- DMARC.org only includes valid records still published in DNS, and are tracked by when they were first published
 - The set of active records changes over time

Active DMARC Records – 2Q 2018


Active DMARC Records and % Change by Month



Data provided by Farsight Security
Graph © 2018 Trusted Domain Project

Active DMARC Records – 3Q 2019

Active DMARC Records and % Growth by Month



Data provided by Farsight Security
Graph © 2019 Trusted Domain Project


3 Million New DMARC Records?

- Millions of DMARC records with strange names
 - `_dmarc.mx.mx.mx.mx.mx.ichiban.example.com`
- Most appear to trace back to “X”
- Nobody was aware of “X” behaving badly
- Exclude these records for now...




Active DMARC Records – 3Q 2019

Active DMARC Records and % Growth by Month

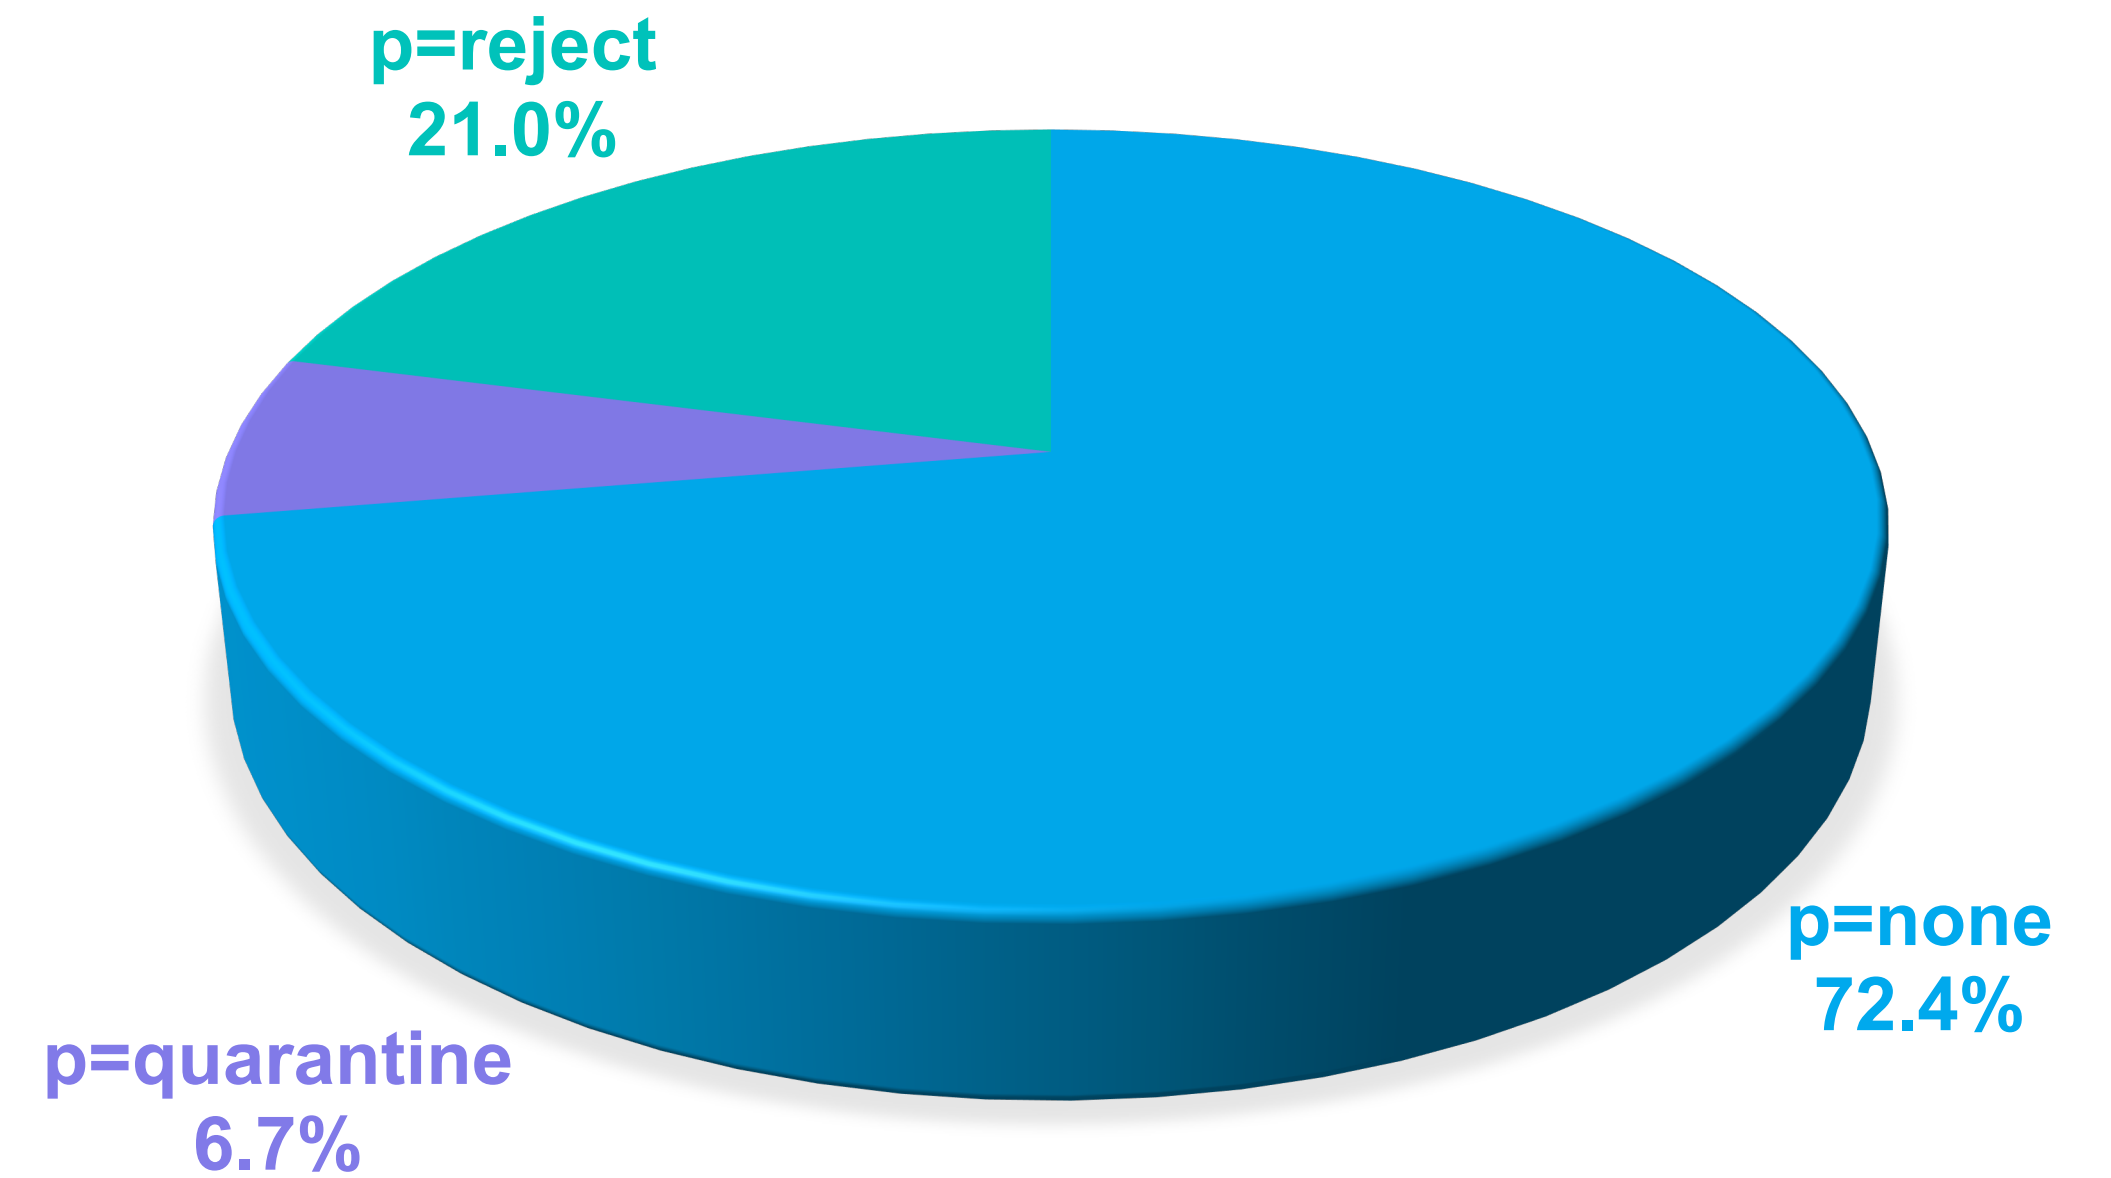


2019.09.30
1,366,656

Policy Breakdown of Active DMARC Records




2018.12.31



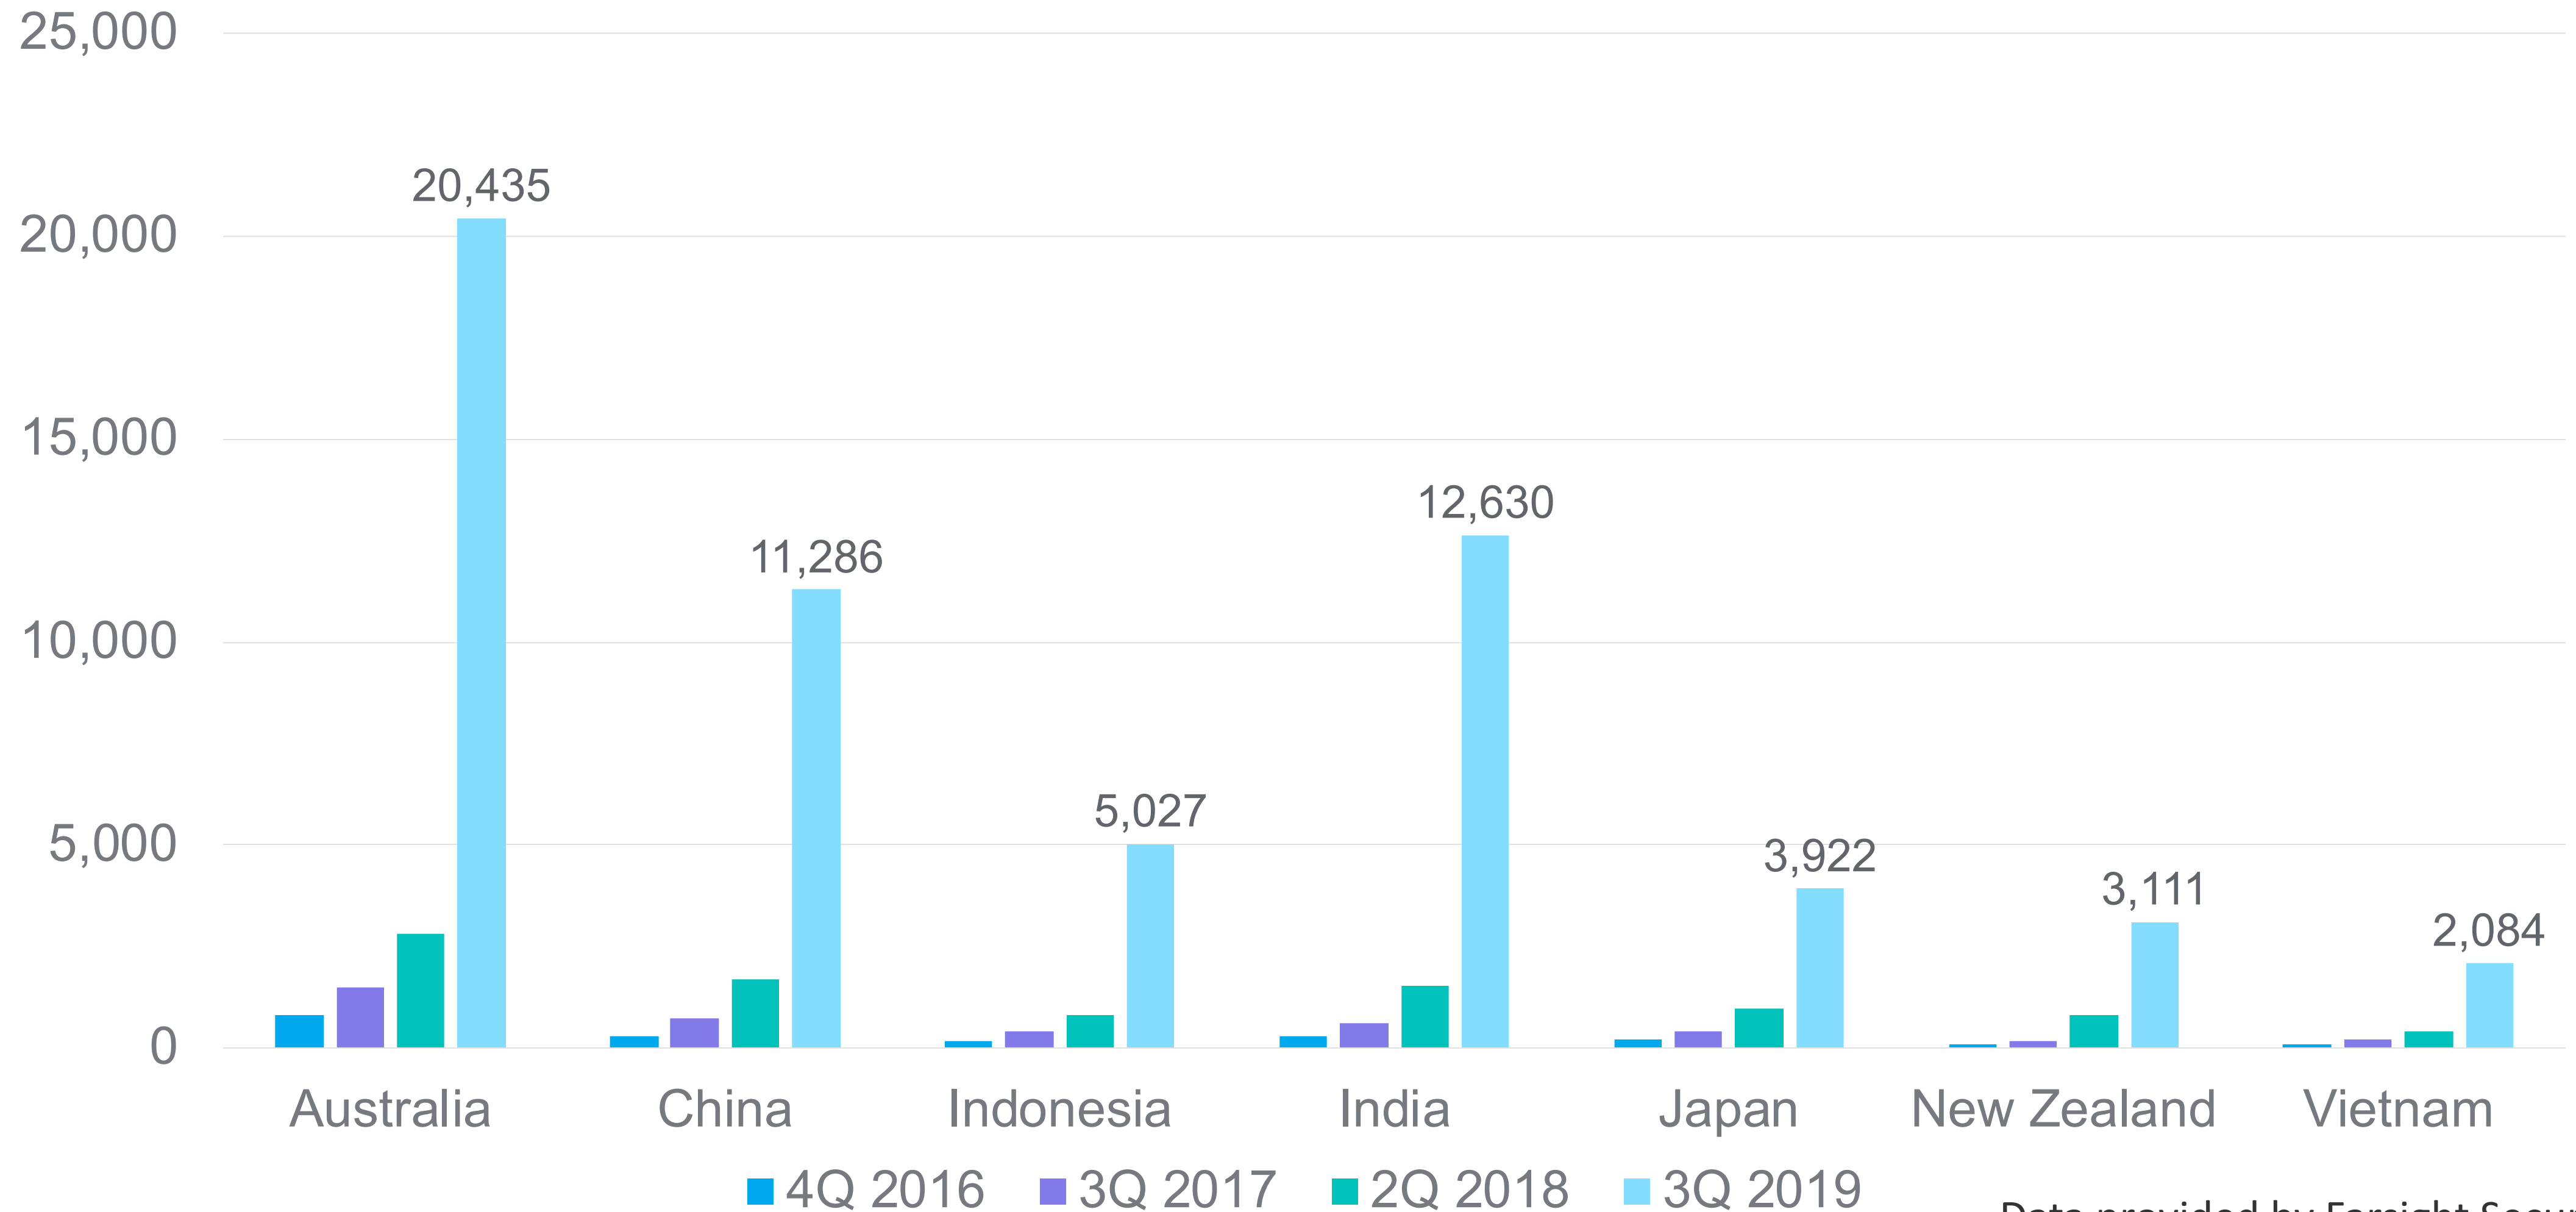
2019.09.30

Active DMARC Records in Euro ccTLDs



Data provided by Farsight Security
Graph © 2019 Trusted Domain Project

Active DMARC Records in Asia ccTLDs




Data provided by Farsight Security
Graph © 2019 Trusted Domain Project


DMARC Records Increase 2.5x Year-over-Year

- Cumulative counts confirmed in DNS for the periods ending
- Robust growth
- Nearly doubled in 2H2018 alone
- Excluding 5MM suspicious records created in 4Q2018

Raw Data: Farsight Security
Analysis: DMARC.org



Active DMARC Record Growth



Common Problems with DMARC Records

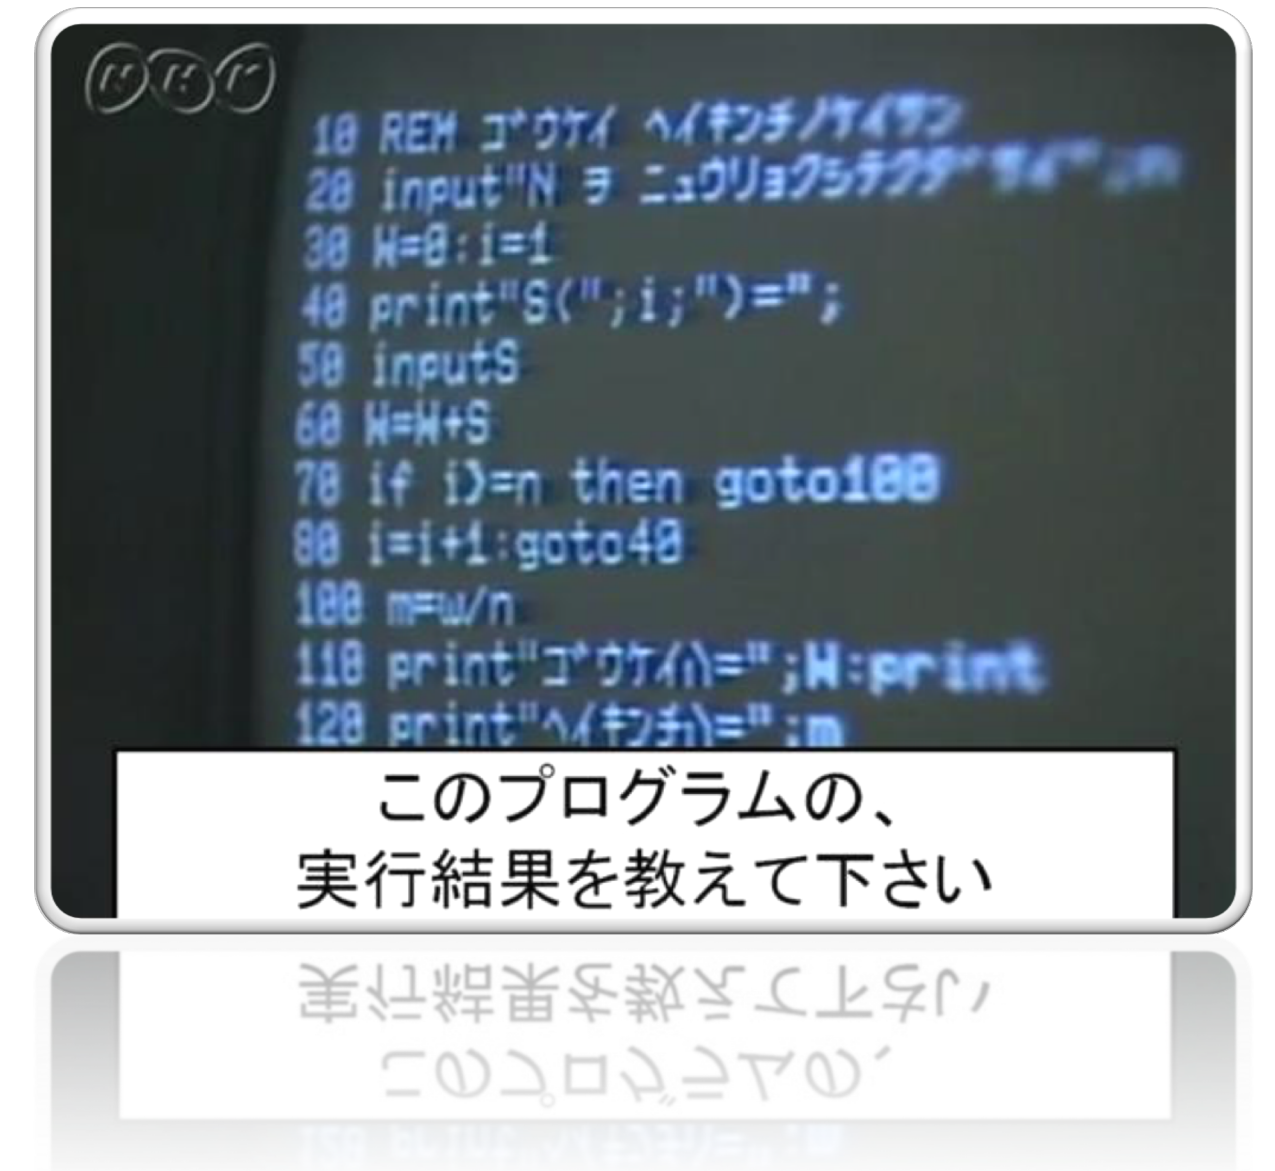


Problems with DMARC Records

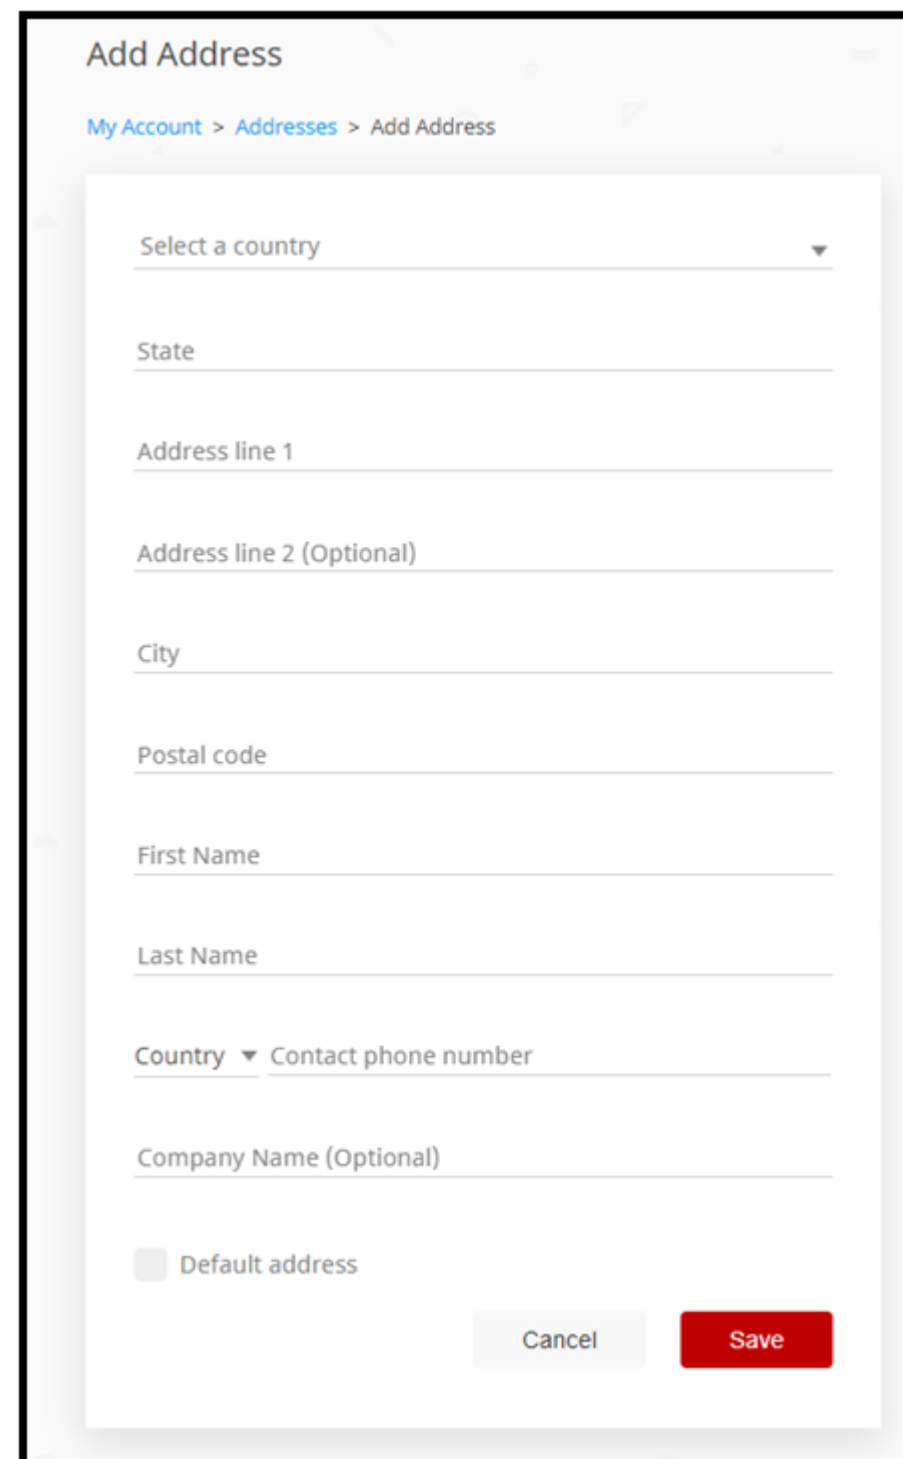
- 2012-2016: 489,000 bad TXT records (`_dmarc...`)
- 2017-2019: 446,000 bad TXT records
- Many are non-DMARC “wildcard” records
 - 76,000 `bio=<base64>`
 - 42,000 `google-site-verification`
 - 25,000 `v=dmarc1` (must be `v=DMARC1`)
 - 11,000 `MS=ms [0-9]*`

Problems with DMARC Records

- Many bad records are formatting issues in `rdata`
 - `\"v=DMARC1`
 - `v= DMARC1 ...`
 - `V-DMARC`
 - `Value: V=DMARC1; ...`
 - `_dmarc... IN TXT \"v=DMARC1 ...`



Problems (?) with DMARC Records



The screenshot shows a web form titled "Add Address" with a breadcrumb trail "My Account > Addresses > Add Address". The form contains the following fields: "Select a country" (dropdown), "State" (text), "Address line 1" (text), "Address line 2 (Optional)" (text), "City" (text), "Postal code" (text), "First Name" (text), "Last Name" (text), "Country" (dropdown) and "Contact phone number" (text) on the same line, and "Company Name (Optional)" (text). At the bottom, there is a "Default address" checkbox, a "Cancel" button, and a red "Save" button.

- Policy records with no reporting address
 - "v=DMARC1; p=none"
 - p=reject and no reporting, may be intentional
 - p=none and no reporting...?
- p=none intended to generate reports
- Does this really qualify as deploying DMARC?

Problems with DMARC Records

- Bad mailto: URIs in published policy
 - `rua=mailto:devops`
 - `rua=mailto:rua [] example.com`
 - `rua=user@domain` not `rua=mailto:...`
- Not just missing reports, may harass reporter
- Potential privacy violations



- Not Deliverable As Addressed
- Unable To Forward
- Insufficient Address
- Moved, Left No Address
- Unclaimed Refused
- Attempted – Not Known
- Not Such Street Number
- Vacant Illegible
- No Mail Receptacle
- Box Closed – No Order
- Returned For Better Address
- Postage Due _____

Verifying 3rd Party Report Receivers

- Domain owners publish *authorizing records* under RFC 7489 Section 7.1
 - `foo.com` wants DMARC reports sent to `bar.com`
 - `_dmarc.foo.com = "rua=mailto:foo@bar.com"`
 - **`foo.com._report._dmarc.bar.com = "v=DMARC1"`**
- Report generators are not checking
- Big privacy and legal implications

Q & A


—

“Why Do Your Numbers
Change?”

—

Growth of DMARC Adoption Globally – 3Q 2017

Total DMARC Records and % Change by Month




2016.12.31
66,321

Data provided by Farsight Security
Graph © 2017 Trusted Domain Project

Growth of DMARC Adoption Globally – 2Q 2018

Total DMARC Records and % Change by Month



Data provided by Farsight Security
Graph © 2018 Trusted Domain Project

Farsight Security DNS Data

- Sensors located at network providers around the world
- Response data – the answer – is timestamped and stored
- Sensors only see records when somebody looks them up
- DMARC.org only includes valid records still published in DNS, and are tracked by when they were first published
 - The set of active records changes over time

Why Do The Counts Change Over Years?

- `ichi.com` and `ni.com` publish DMARC records during 2015
- They are both still published as of 2015.12.31, so the total for 2015 as of 2015.12.31 is 2
- During 2016 `ni.com` removes its DMARC record, but `san.com` publishes a DMARC record
- The total for 2015 as of 2016.12.31 is 1, and the count for 2016 as of 2016.12.31 is 1.
- During 2016 `ichi.com` removes its DMARC record
- The count for 2015 as of 2017.12.31 is 0, and the count for 2016 is 1

Concrete Example

- As of 2017.09.30: We reported 66,321 DMARC records for 2016.12.31
- As of 2018.06.30: We reported 54,315 DMARC records for 2016.12.31
- 12,006 records that were active during the 2017.09.30 validation were no longer active during the 2018.06.30 validation
- Since they were no longer in DNS, they are not included in the 2016 total for the 2018 report

ありがとうございました

Thank you

—