Introduction to Email Authentication

An explanation of how SPF, DKIM, and DMARC function

2Q 2015

Introduction to DMARC.org

DMARC.org is an initiative of the non-profit Trusted Domain Project (TDP).

The mission of DMARC.org is to promote the use of DMARC and related email authentication technologies to reduce fraudulent email, in a way that can be sustained at Internet scale. This overall goal is met by educating individuals and organizations through a combination of articles, tutorials, presentations, and webinars.

For more information, please visit https://dmarc.org

For more about TDP, please visit http://trusteddomain.org

Introduction to DMARC.org

DMARC

The work of DMARC.org is made possible through the generous support of these sponsors:

AGARI

Comcast

Google

PayPal

Return Path

TDP Trusted Domain Project

FARSIGHT SECURITY

Overview

- Background
- What is Email Authentication
- SPF Basics
- DKIM Basics
- DMARC Concepts
- DMARC Mechanics

Who Sent This Message?

- It says it's from BigBank…
- It shows a BigBank address…
- It has a BigBank logo…
- I can't tell it isn't from BigBank…
- I have a BigBank account…
- I don't want my account closed…

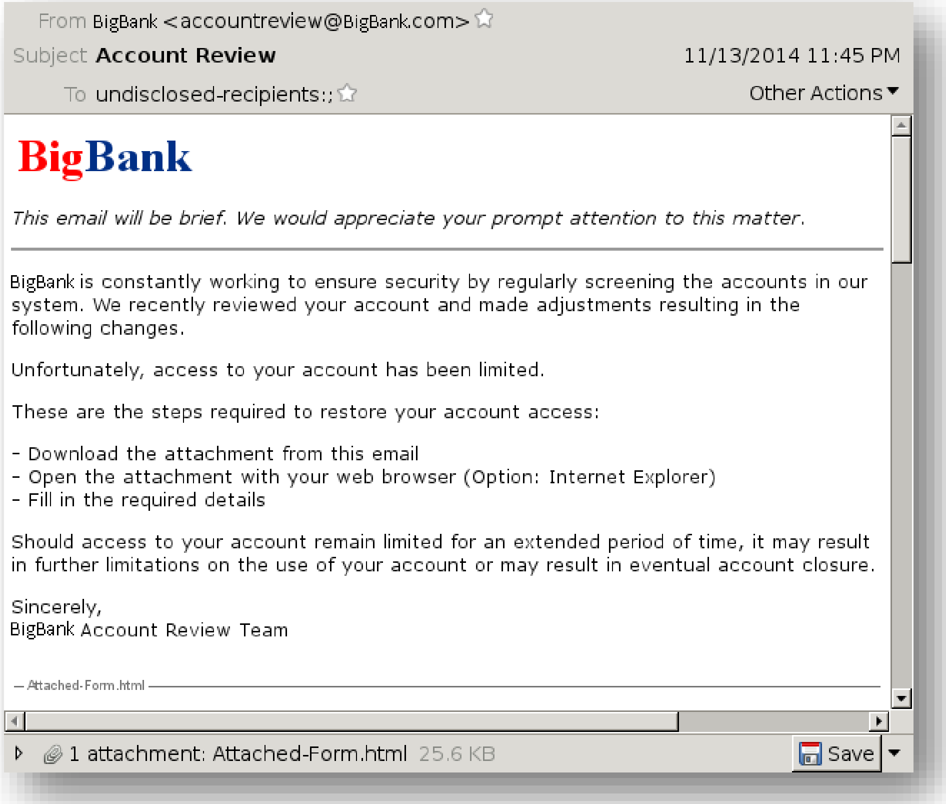

From BigBank <accountreview@BigBank.com>

Subject **Account Review** 11/13/2014 11:45 PM

To undisclosed-recipients:;

BigBank

This email will be brief. We would appreciate your prompt attention to this matter.

BigBank is constantly working to ensure security by regularly screening the accounts in our system. We recently reviewed your account and made adjustments resulting in the following changes.

Unfortunately, access to your account has been limited.

These are the steps required to restore your account access:

- Download the attachment from this email
- Open the attachment with your web browser (Option: Internet Explorer)
- Fill in the required details

Should access to your account remain limited for an extended period of time, it may result in further limitations on the use of your account or may result in eventual account closure.

Sincerely,
BigBank Account Review Team

— Attached-Form.html

1 attachment: Attached-Form.html 25.6 KB

But it was sent by an imposter

Email Messages Can Be Faked

- Email evolved in a controlled and trusted environment, it wasn’t designed to exclude fakery and bad actors
- Now both spam and anti-spam are multi-billion dollar industries
- This “arms race” breeds better, more effective spam
- Phishing borrows these techniques to compromise users/systems rather than sell them something

Why Is That Even Possible?

- “Email” first appeared on timesharing systems in 1960s
- Network email appears on ARPANET in the early 1970s
- ARPANET: Closed community of academics, researchers, and government contractors
- Some abuse happened, but was addressed within the community
- Commercial use was largely illegal
- There was no money to be made by abusing email
- Priority in design of email through the early 1990s was on reliability and deliverability between different networks, operating systems, etc

Background: ARPANET → Internet

- DARPA/DISA to National Science Foundation (NSF) through the 1980s (CSNET, NSFNET, FIX/NAPs)
- UUCP/Usenet spreads, providing email without restrictions
- ARPANET finally decommissioned 1990
- Remaining restrictions on commercial use steadily removed 1990-95
- “Spamming” is coined on Usenet/NetNews circa 1993
- Blatant commercial spamming begins in 1994 with Cantor & Siegel on Usenet
- Practice quickly spreads to email; 90% of all email in 2009

Cantor & Siegel’s First Spam Campaign

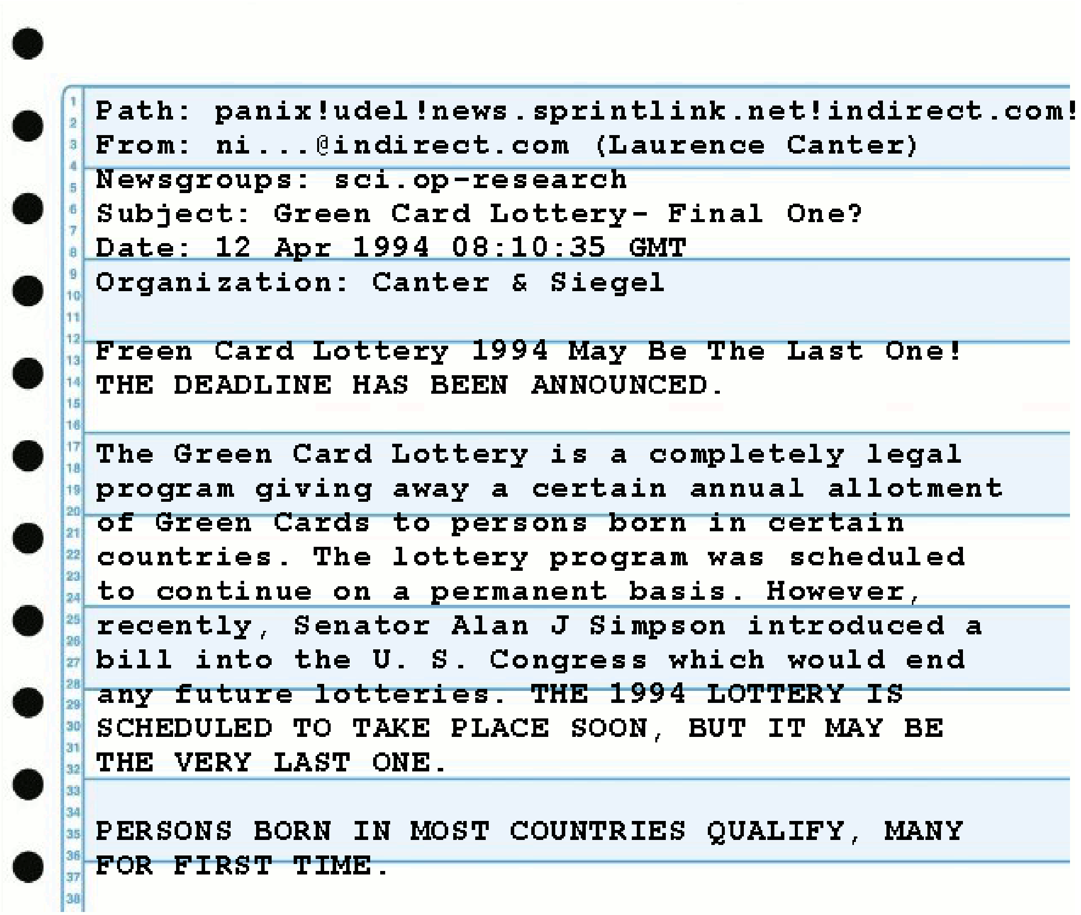

```
Path: panix!udel!news.sprintlink.net!indirect.com!
From: ni...@indirect.com (Laurence Canter)
Newsgroups: sci.op-research
Subject: Green Card Lottery- Final One?
Date: 12 Apr 1994 08:10:35 GMT
Organization: Canter & Siegel

Freen Card Lottery 1994 May Be The Last One!
THE DEADLINE HAS BEEN ANNOUNCED.

The Green Card Lottery is a completely legal
program giving away a certain annual allotment
of Green Cards to persons born in certain
countries. The lottery program was scheduled
to continue on a permanent basis. However,
recently, Senator Alan J Simpson introduced a
bill into the U. S. Congress which would end
any future lotteries. THE 1994 LOTTERY IS
SCHEDULED TO TAKE PLACE SOON, BUT IT MAY BE
THE VERY LAST ONE.

PERSONS BORN IN MOST COUNTRIES QUALIFY, MANY
FOR FIRST TIME.
```

What is Email Authentication?

What is Email Authentication?

Technologies that let you determine whether the sender you see really sent that message

- Most of these technologies are implemented in the infrastructure, and end-users don't see them
- Some end-user clients (MUAs) may have options to show a gold key, or similar icon
- All of these technologies have strengths and weaknesses
- Presently no 100% solutions

Three Primary Protocols

- **SPF** - Sender Policy Framework (2003)
  - IETF Status: Standards Track RFC
  - http://www.openspf.org

- **DKIM** – Domain Keys Identified Message (2007)
  - IETF Status: Standards Track RFC
  - http://opendkim.org

- **DMARC** – Domain-based Message Authentication, Reporting, and Conformance (2012)
  - IETF Status: Informational RFC, Working Group
  - http://dmarc.org

Other Protocols

- Sender-ID – Combination of SPF and Caller ID proposals
  - SPF doesn't deal with message headers – addresses that
  - Promoted by Microsoft, who asserted patent interests
  - Not used in new deployments
  - IETF Status: Experimental (2006)
- ADSP - Author Domain Signing Practices
  - Extension to DKIM allowing domain owner to specify whether or not they signed all outgoing mail
  - Specification itself discourages use of any positive assertion about signing
  - Actual use of the "discardable" policy, indicating all messages without a signature from the sending domain should be discarded, was highly controversial
  - IETF Status: Historic (2014)

Background: Envelope vs. Header

- RFC5321 defines the host-to-host protocol
- RFC5322 governs the contents of messages
- RFC5322.From is usually what the end-user sees

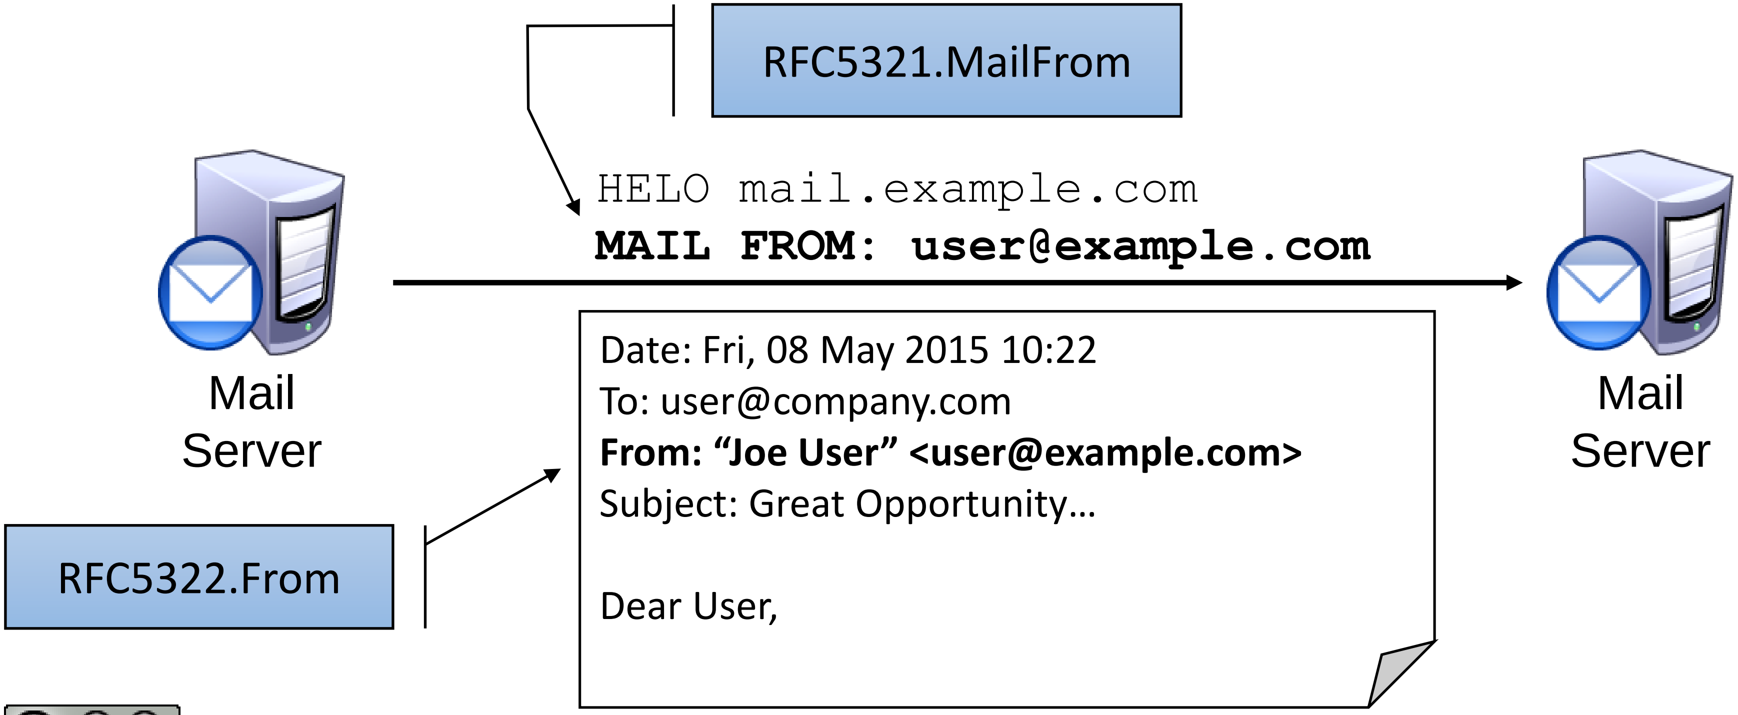

```
HELO mail.example.com
MAIL FROM: user@example.com
```

Date: Fri, 08 May 2015 10:22
To: user@company.com
From: “Joe User” <user@example.com>
Subject: Great Opportunity…

Dear User,

SPF Basics

SPF – Sender Policy Framework

- Allows domain owner to specify which servers may use addresses in that domain in the RFC5321.MailFrom
  - A “path based” approach
  - Fallback to the RFC5321.HELO domain for a “null sender”
- Indirect mailflows (forwarders, mailing lists) cause SPF to either fail, or lookup against a rewritten RFC5321.MailFrom
  - The latter may provide a pass against a different domain than the original author – which is something spammers often do…
- No link required between RFC5321.MailFrom and RFC5322.From
- Mail receivers declined to filter mail based solely on SPF results due to a combination of indirect mailflows, widespread deployment errors, and other issues

SPF - Limitations

- SPF typically fails after the first relay or "hop"
  - Forwarding, mailing lists, etc
- Mailing lists and other indirect flows can rewrite the RFC5321.MailFrom to generate an SPF pass.
  - Unfortunately something spammers like to do, too
- Many receivers do not act on SPF's policy assertions
  - Widespread misconfiguration, historically & presently
  - Problematic indirect mailflows
  - What am I supposed to do with a "softfail?"

SPF Records - Contents

- DNS TXT records located at the name of the domain in question
  - `example.com` or `mail.example.com`
- Sample:

```
example.com    IN    TXT    "v=spf1 a:mail.example.com -all"
```

- Identifier tag: `v=spf1`
- Mechanisms:
  - `a` check host against this hostname
  - `mx` check host against this DNS MX record
  - `ip4` check host against this IPv4 address specification
  - `ip6` check host against this IPv6 address specification
  - `exists` check host against a (complex) macro
  - `ptr` officially deprecated in RFC7208 – do not use
- CIDR address blocks are common (`ip4:192.168.1.0/24`)

SPF Records - Mechanisms

- Result for mechanism matches:

+	Pass (implicit – “+a:mail.example.com”)
−	Fail
~	Softfail
?	Neutral

- Many macros can be used, but 90+% of records include:

+all	Pass all matches
-all	Fail all matches
~all	Softfail all matches
?all	Neutral result for all matches

SPF Records – Common Examples

- `v=spf1 a:mail.example.com ip4:192.168.1.0/29 ~all`
  - Allow mail.example.com
  - Allow any host in IPv4 address block
  - Any others are probably unauthorized, but not 100% - softfail
- `v=spf1 mx include:spf.example.net include:[...] -all`
  - Allow any host that appears in the SPF record at spf.example.net
  - Allow the host if it appears in the MX records for this domain
  - Note: the [...] above is just for slide formatting – not legal
- `v=spf1 -all`
  - Fail everything – deployed for “parked” or unused domains

SPF In Action – Common Case

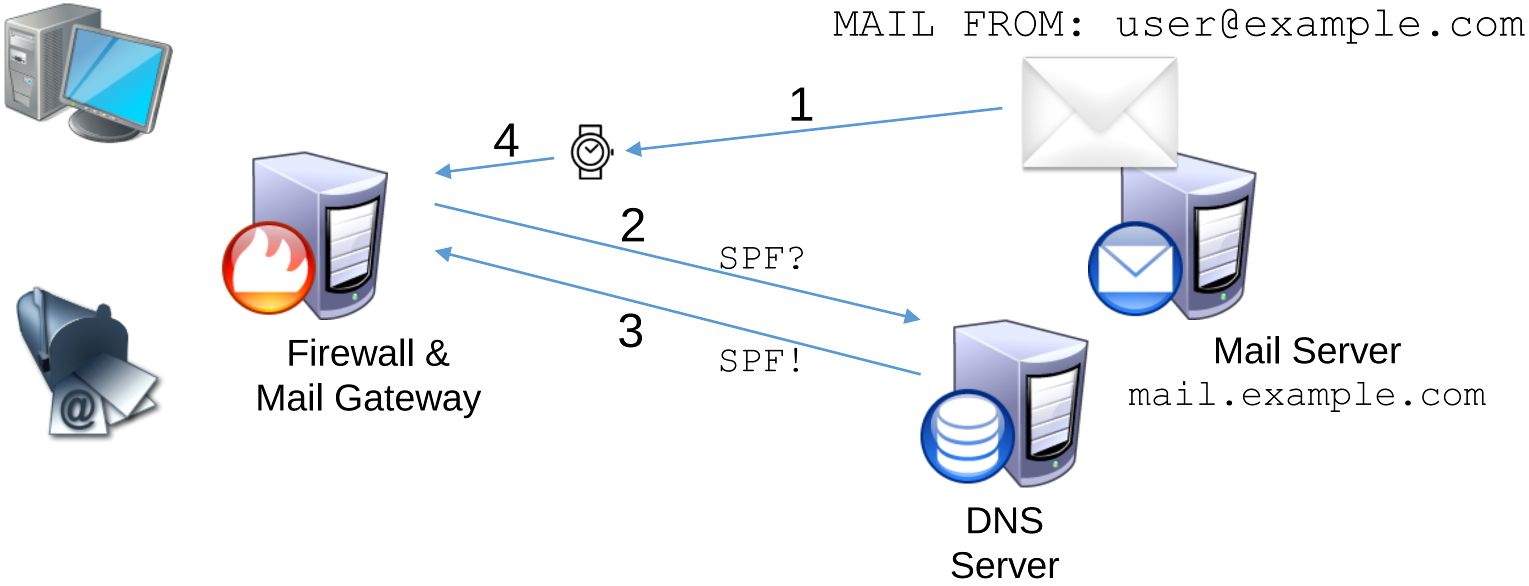

```
example.com   IN   TXT   “v=spf1 a:mail.example.com -all”
```

1: Message sent from example.com, invoking SPF check

2: Receiver looks up SPF record for RFC5321.MailFrom domain

3: SPF record returned

4: `mail.example.com` is authorized by the SPF record, message accepted

SPF and Mailing Lists

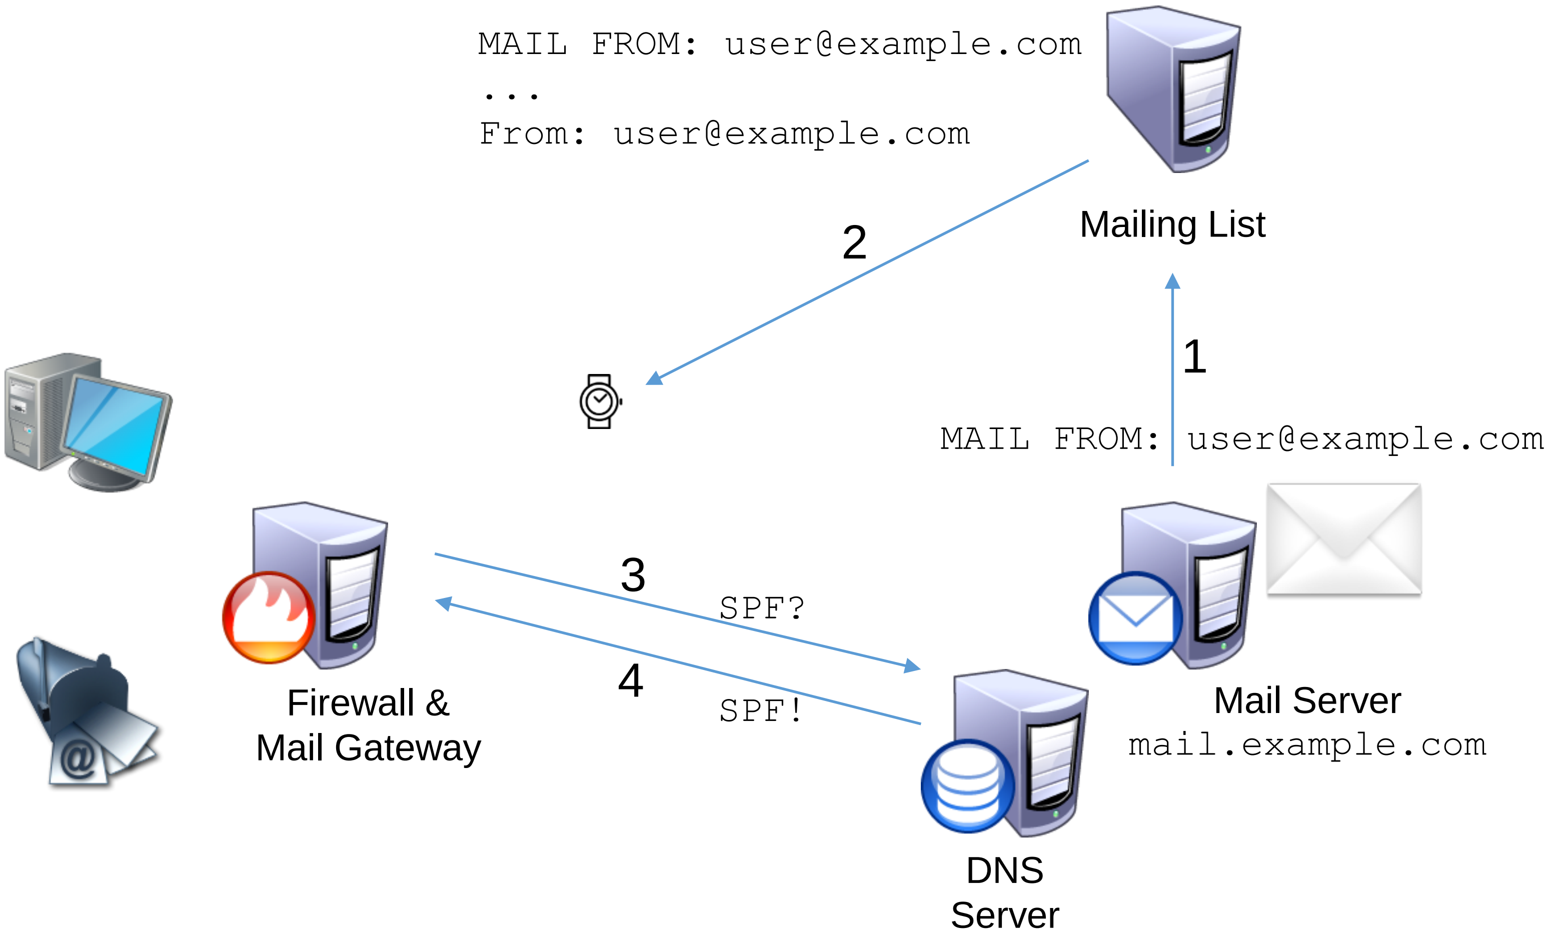

```
MAIL FROM: user@example.com
...
From: user@example.com
```

```
MAIL FROM: user@example.com
```

```
example.com   IN   TXT   "v=spf1 a:mail.example.com -all"
```

SPF and Mailing Lists

```
MAIL FROM: list-owner@listsRus.com
...
From: user@example.com
```

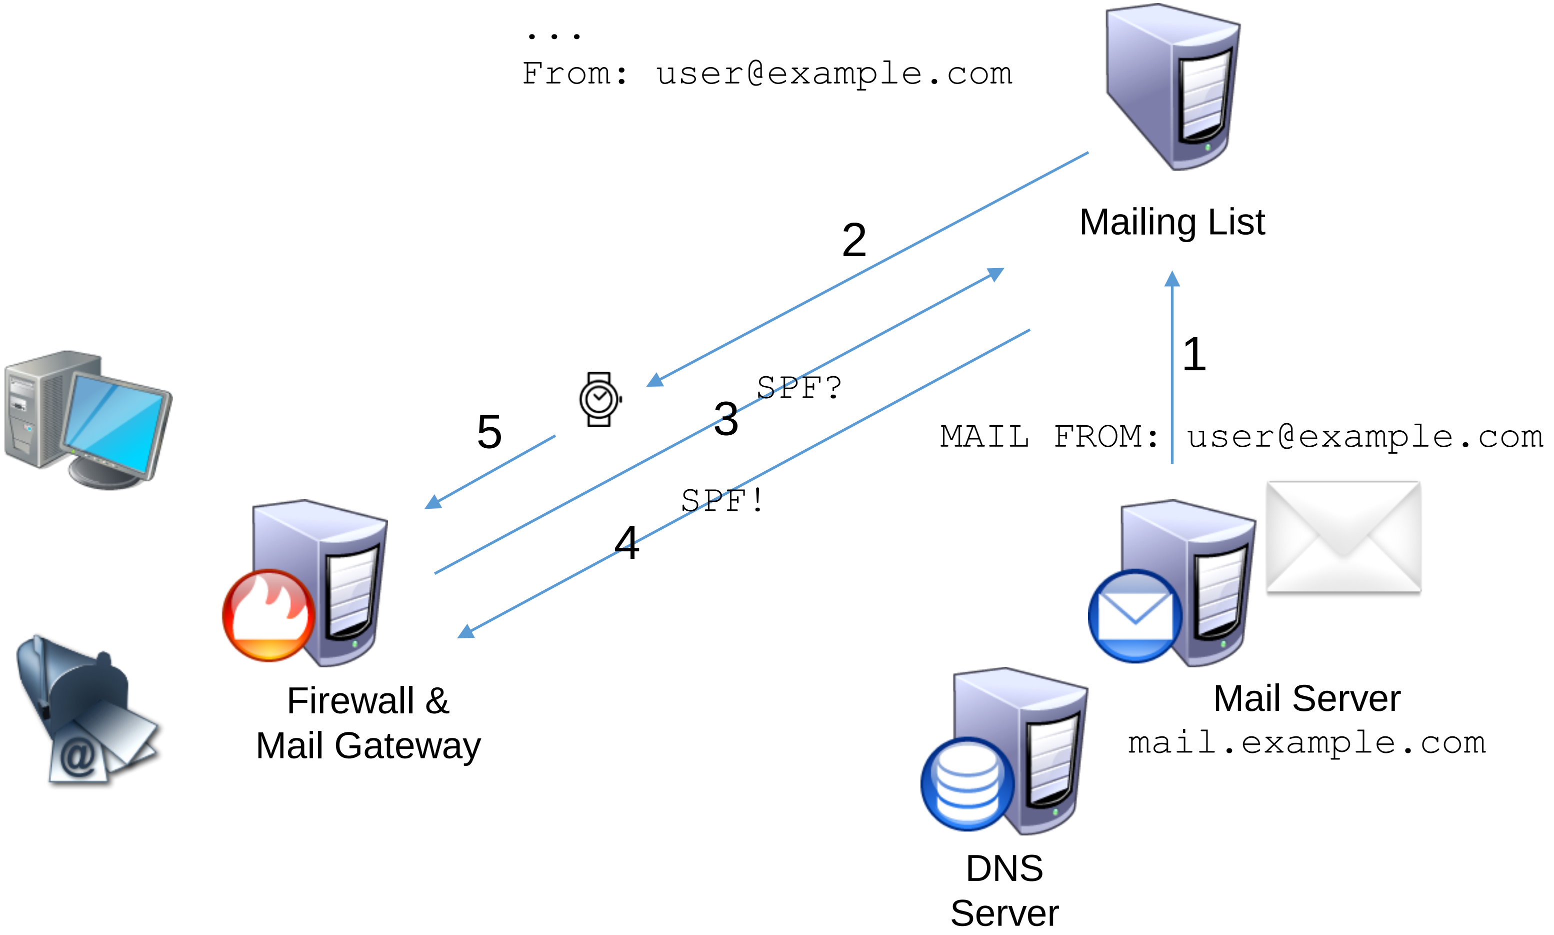

```
example.com   IN   TXT   "v=spf1 a:mail.example.com -all"
```

SPF and Bad Actors

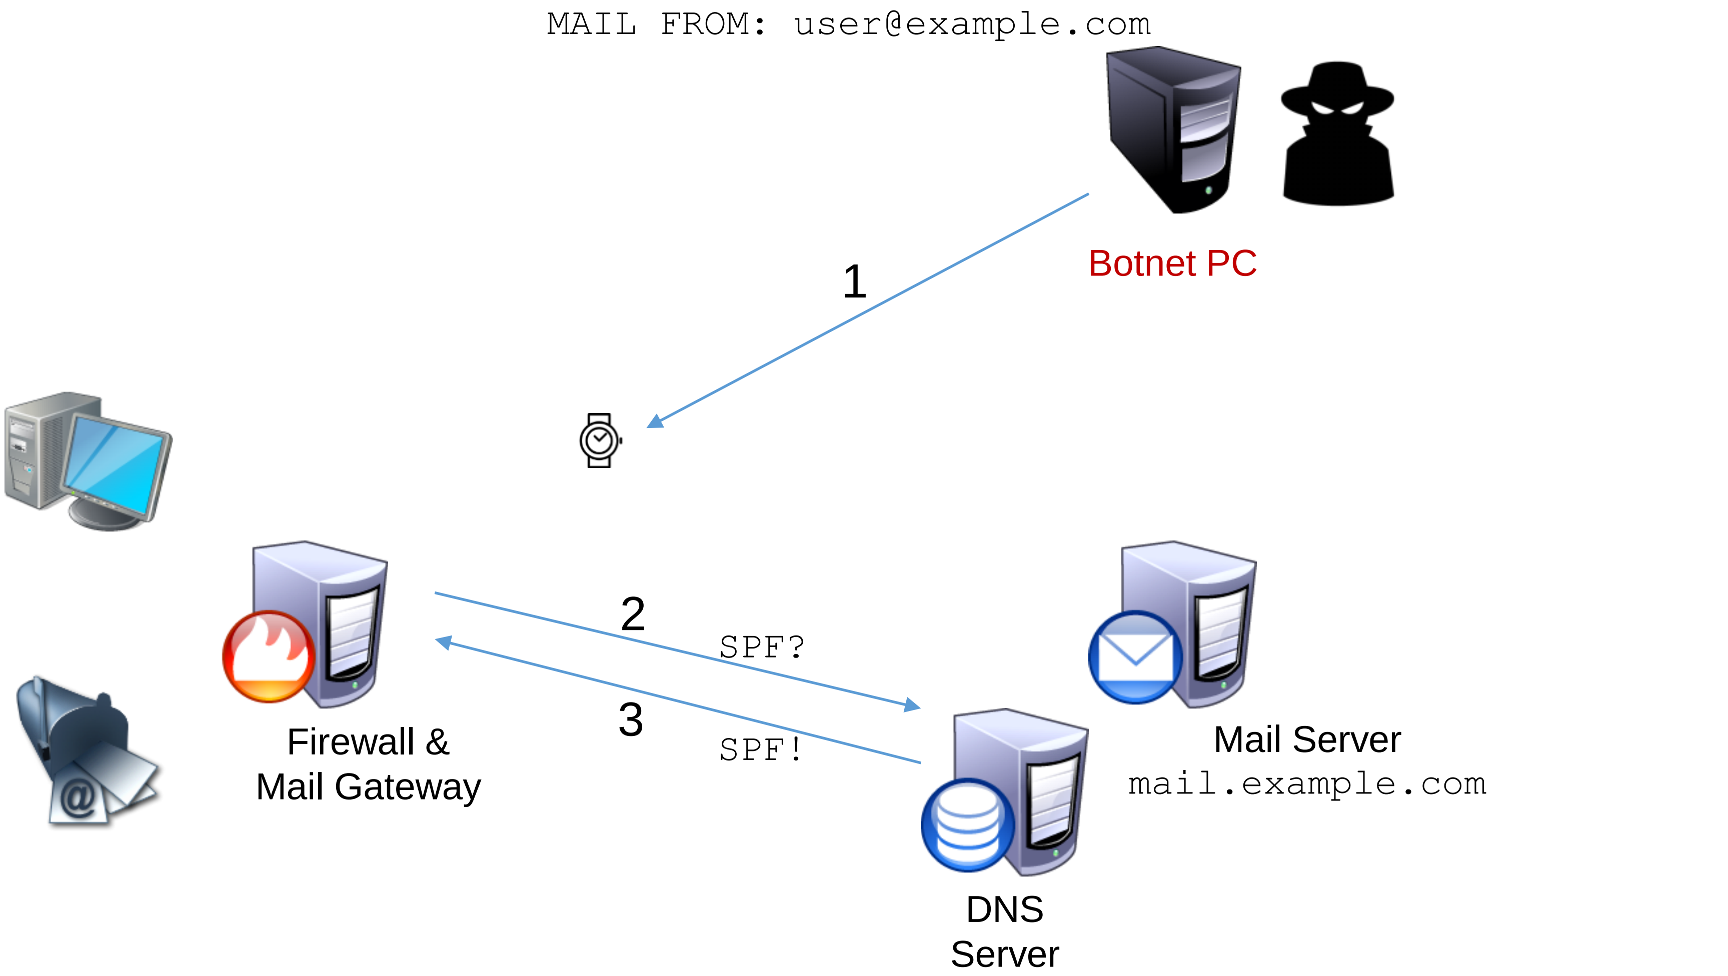

```
example.com   IN   TXT   "v=spf1 a:mail.example.com -all"
```

SPF and Bad Actors

```
MAIL FROM: badguy@evilspammer.com
…
From: user@example.com
```

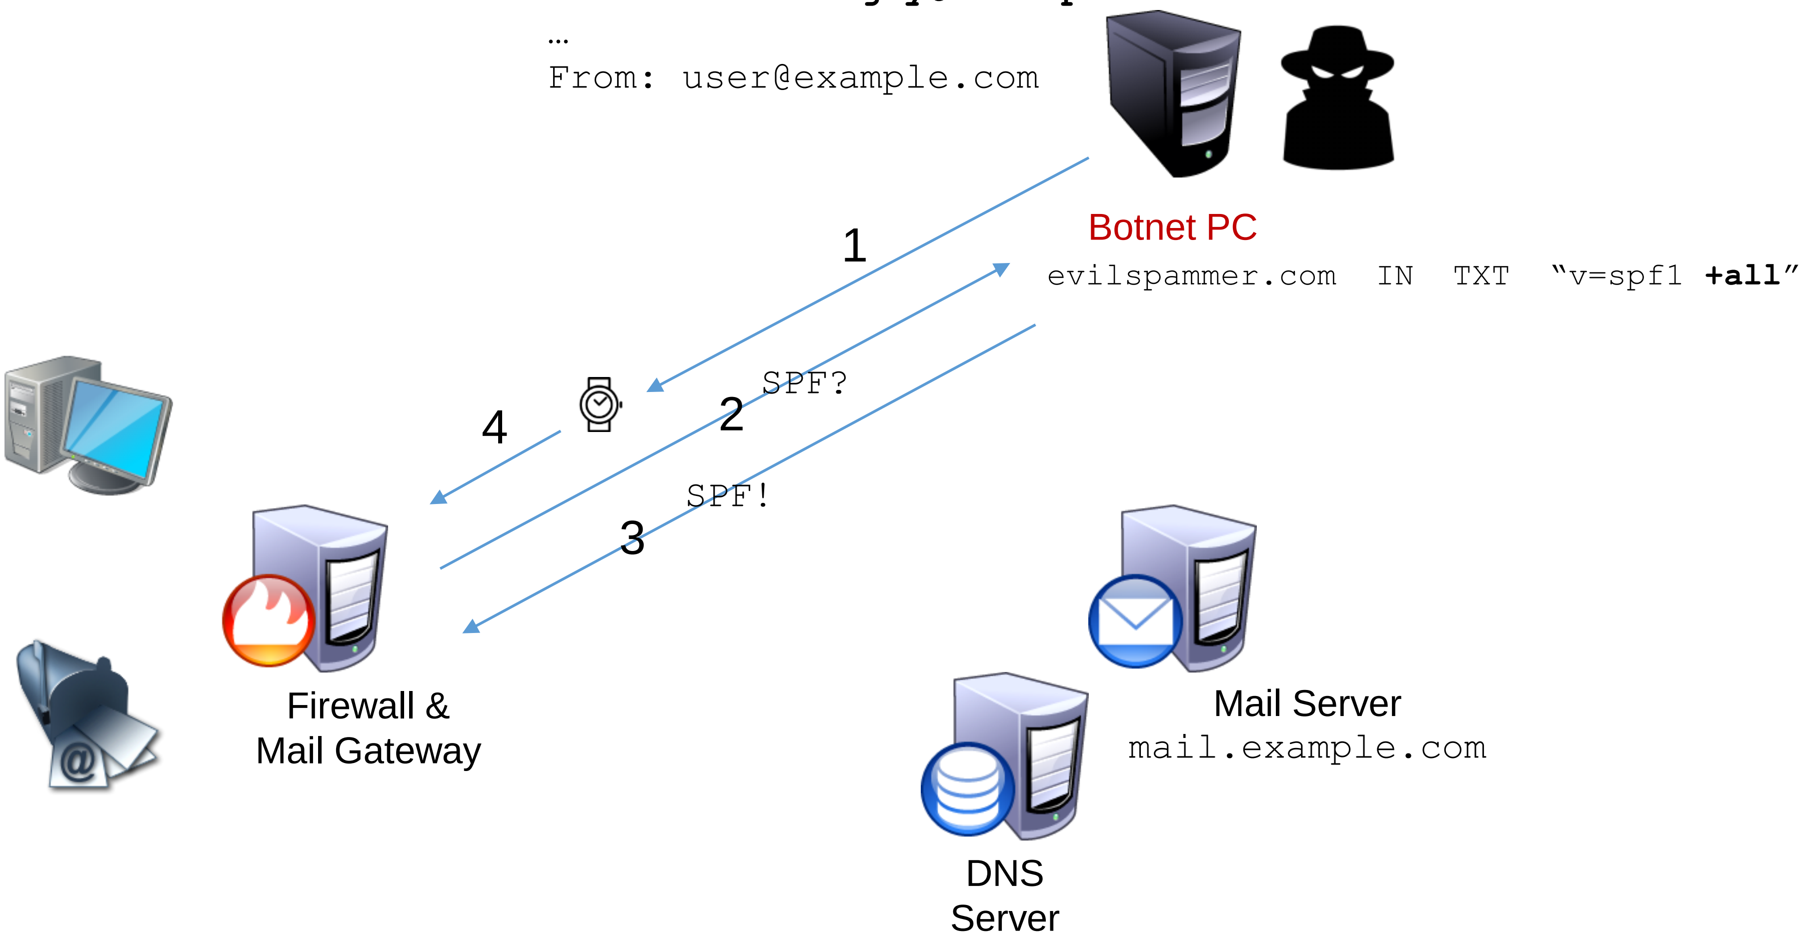

Botnet PC

```
evilspammer.com   IN   TXT   "v=spf1 +all"
```

1

SPF?

2

SPF!

3

4

Firewall &
Mail Gateway

Mail Server
mail.example.com

DNS
Server

```
example.com   IN   TXT   "v=spf1 a:mail.example.com -all"
```

DKIM Basics

DKIM - Domain Keys Identified Message

- DKIM uses a digital signature based on public key cryptography
- Sending organization uses private key to sign a hash or fingerprint of the message before it enters the Internet
- Receiver can retrieve the corresponding public key via DNS to verify the signature
- Signing domain does not have to have any relationship to the domains in the RFC5322.From or RFC5321.MailFrom
    - End-user typically never sees which domain asserted responsibility for a signed message

DKIM – Limitations

- More complicated to deploy than SPF
- Won’t verify if the signed parts of the message are altered
  - Mailing lists modifying Subject header
  - Corporate gateways adding a disclaimer/footer
  - Alumni services transcoding messages, e.g. ASCII to UTF8
  - Filtering services removing images or MIME parts
- Didn’t have a policy mechanism that was widely adopted
  - ADSP was made Historic by IETF in May 2014
- Crypto concerns need to be tracked and addressed
  - Key length for signatures
  - Strength of hashing algorithms

Anatomy of a DKIM Signature

```
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/simple; d=ietf.org; s=ietf1;
    t=1432264097; bh=DnGEIxFIoMduuUnbGf/ktbNUxOx7JkZRbjuQFmsr70M=;
    h=From:To:In-Reply-To:References:Date:Message-ID:MIME-Version:Cc:Subject:
    List-Id:List-Unsubscribe:List-Archive:List-Post:Content-Type:Sender;
    b=t6F/a3rYjOLKdEp8psEy2AfcIjxx0ibZsfRGHsGA7L4xOuwS9aGAwI/XxpxW0TcAY …
```

a=	Hashing algorithm used (SHA256)
b=	Signature data, a hash including the body hash and headers
bh=	Body hash, computed from the message body (up to l= bytes)
d=	Signing Domain Identifier (SDID)
h=	Headers included in signature
i=	Agent or User Identifier (AUID), optional
l=	Length limit of body included in body hash, optional
s=	Selector, identifies which public key to use to verify
t=	Time the signature was computed
x=	Expiration time of signature, optional

DKIM – Retrieving Public Keys

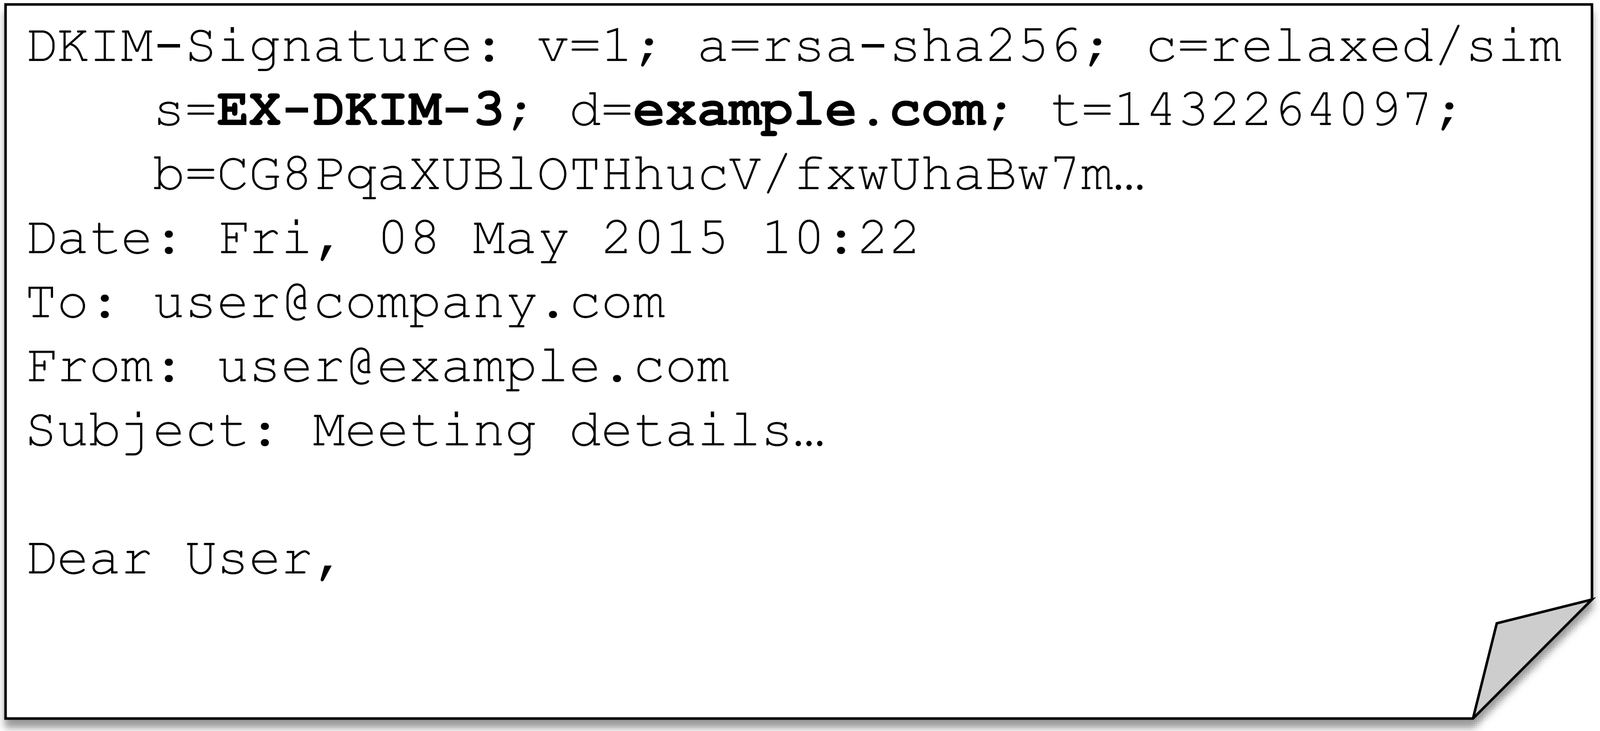

```
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/sim
    s=**EX-DKIM-3**; d=**example.com**; t=1432264097;
    b=CG8PqaXUBlOTHhucV/fxwUhaBw7m…
Date: Fri, 08 May 2015 10:22
To: user@company.com
From: user@example.com
Subject: Meeting details…

Dear User,
```

```
% dig +short txt EX-DKIM-3._domainkey.example.com
"v=DKIM1\; k=rsa\; h=sha1\; p=MIGfMA0GCSqGSIb3DQ…"
```

DKIM in Action

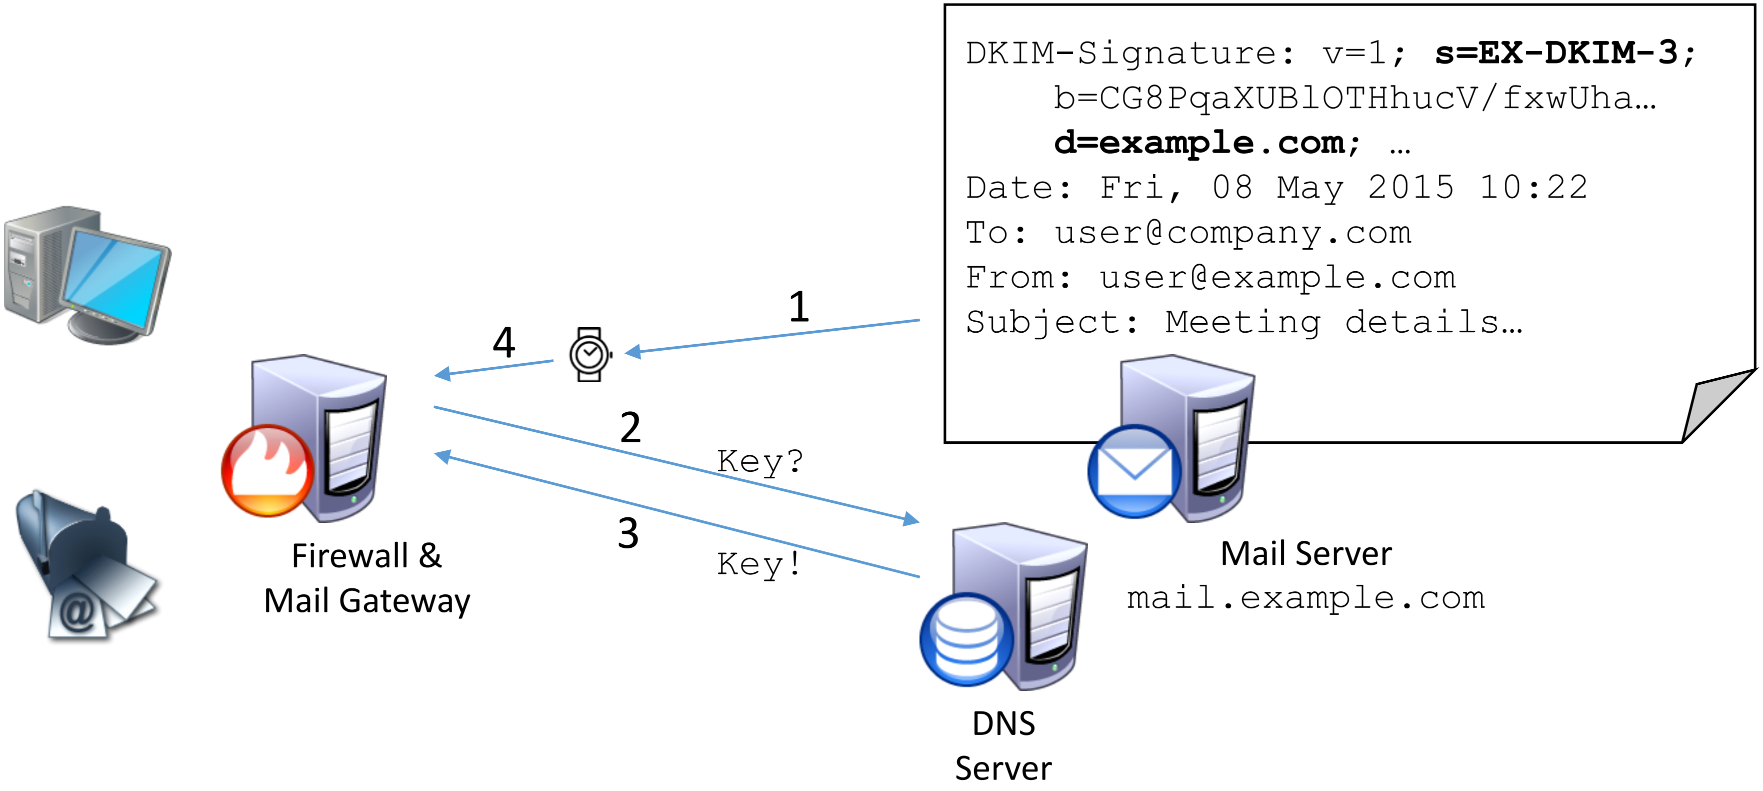

```
DKIM-Signature: v=1; s=EX-DKIM-3;
    b=CG8PqaXUBlOTHhucV/fxwUha…
    d=example.com; …
Date: Fri, 08 May 2015 10:22
To: user@company.com
From: user@example.com
Subject: Meeting details…
```

```
EX-DKIM-V3._domainkey.example.com    IN    TXT
"v=DKIM1; k=rsa; h=sha1; p=MIGfMA0GCSqGSIb3DQ…"
```

DKIM and Mailing Lists

Assuming the mailing list isn't changing the signed parts of the message…

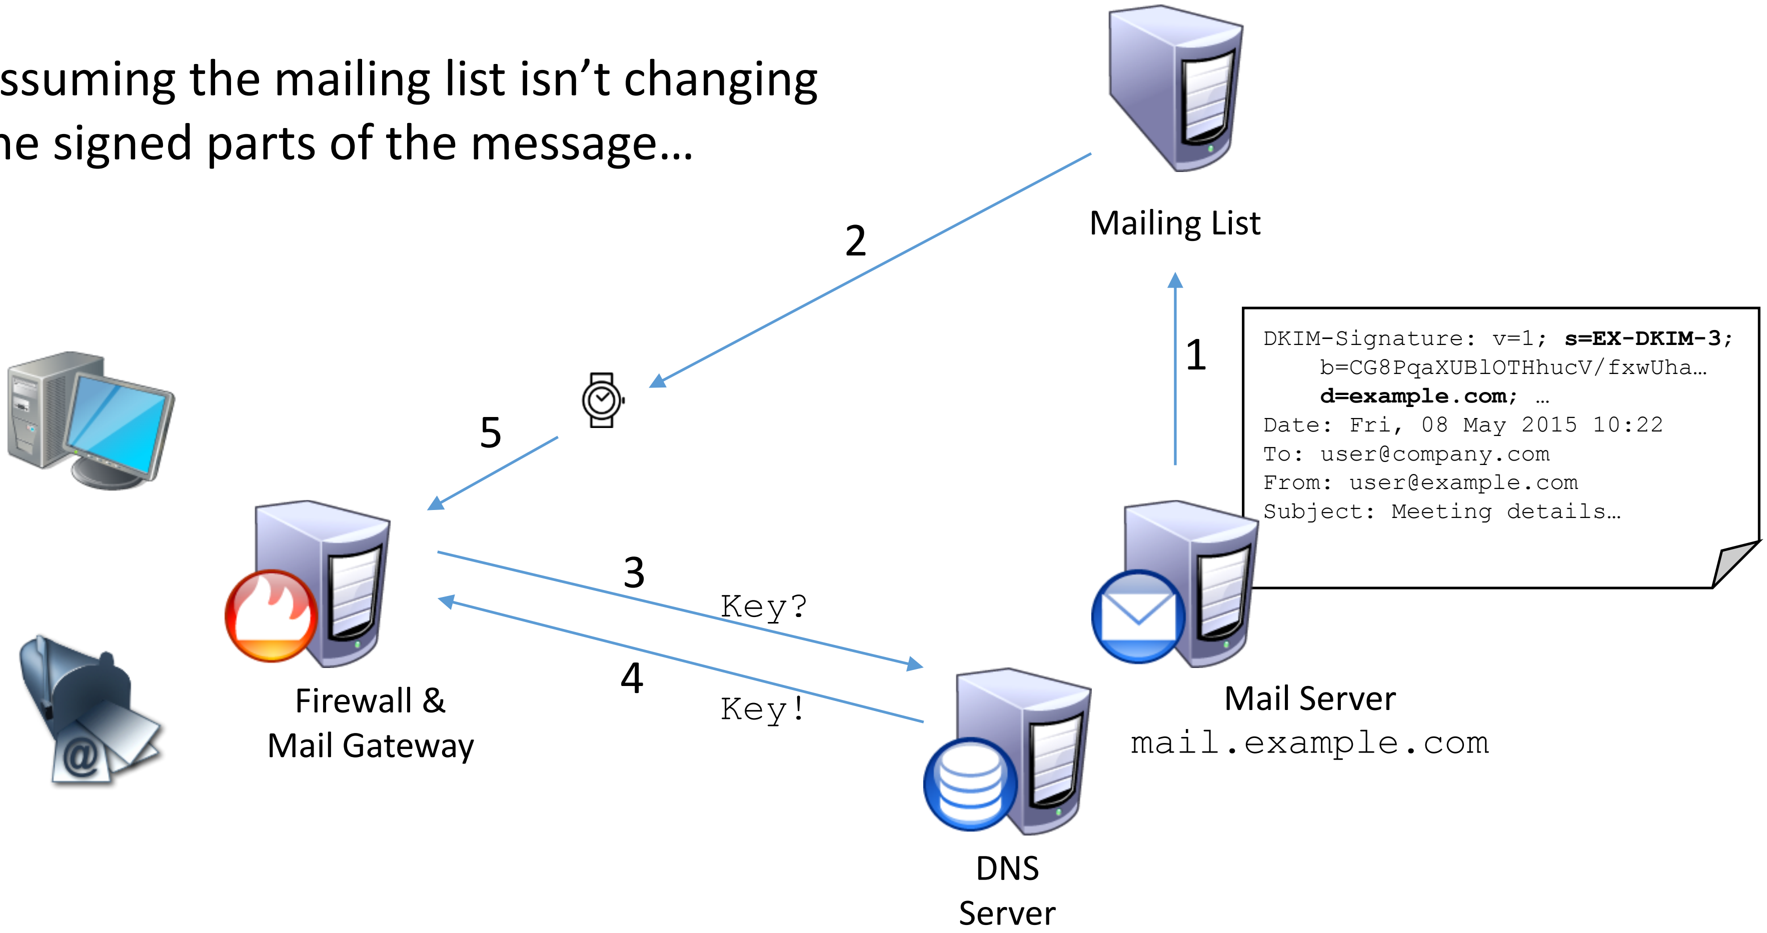

```
DKIM-Signature: v=1; s=EX-DKIM-3;
    b=CG8PqaXUBlOTHhucV/fxwUha…
    d=example.com; …
Date: Fri, 08 May 2015 10:22
To: user@company.com
From: user@example.com
Subject: Meeting details…
```

```
EX-DKIM-V3._domainkey.example.com   IN   TXT
"v=DKIM1; k=rsa; h=sha1; p=MIGfMA0GCSqGSIb3DQ…"
```

DKIM and Mailing Lists

The list has changed a signed portion of the message, breaking the signature…

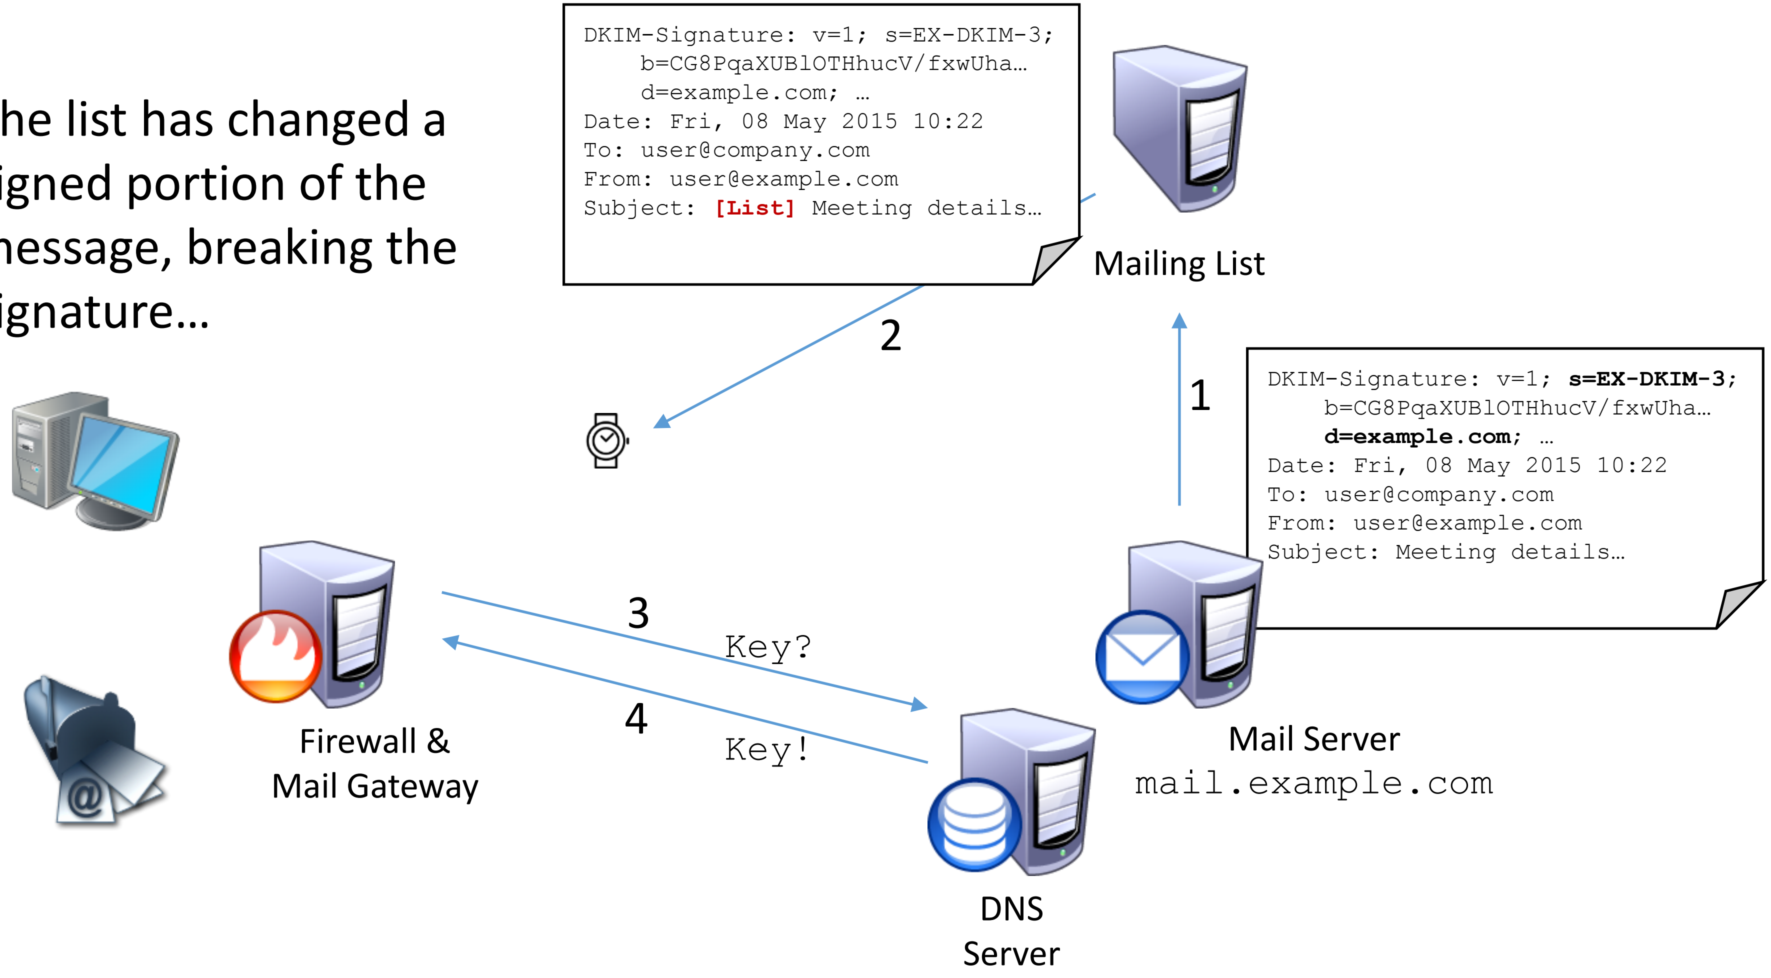

```
DKIM-Signature: v=1; s=EX-DKIM-3;
    b=CG8PqaXUBlOTHhucV/fxwUha…
    d=example.com; …
Date: Fri, 08 May 2015 10:22
To: user@company.com
From: user@example.com
Subject: [List] Meeting details…
```

```
DKIM-Signature: v=1; s=EX-DKIM-3;
    b=CG8PqaXUBlOTHhucV/fxwUha…
    d=example.com; …
Date: Fri, 08 May 2015 10:22
To: user@company.com
From: user@example.com
Subject: Meeting details…
```

```
EX-DKIM-V3._domainkey.example.com   IN   TXT
"v=DKIM1; k=rsa; h=sha1; p=MIGfMA0GCSqGSIb3DQ…"
```

DMARC Concepts

So Many Protocols, Why Add Another?

- No consistency in DKIM and SPF deployment
  - Incomplete coverage of mailstreams
  - Different usage for different senders using same domain
- Receivers could not rely on pass/fail results
  - "Sender wants me to block half their legitimate-looking email. Do they even know that's what they've asked for?"
  - No appetite for angry customers calling Receiver's support team
- Senders didn't know size of their coverage problem
  - Incomplete to zero visibility
  - No way to tell if things improve or worsen
- Nothing had broken this log jam in several years

DMARC – Domain-based Message Authentication, Reporting and Conformance

Developed from the experience of PayPal, Yahoo, GMail, and others in making DomainKeys effective for stopping abuse*

High-level principles or guidelines:

- Senders clearly opt-in by publishing DMARC policy
- Receivers provide feedback so Senders can close gaps
- Senders increase level of authenticated email
- Receivers can identify and block unauthenticated email
- Must work at Internet scale
- Succeed with this, then address more difficult threats

DMARC - Overview

- Layers on top of DKIM and SPF
  - Signatures (DKIM) may survive when a path (SPF) doesn’t
  - SPF may work even if the sender screws up DKIM temporarily
- Allows policy assertions to quarantine or block messages that do not authenticate
- Lots of data collection and reporting
- For DMARC to pass, either DKIM or SPF must pass ***BUT***:
- Additional requirements on DKIM and SPF results
  - A DKIM/SPF pass is not always a DMARC pass

DMARC - Limitations

- More complicated if you want to assert an active policy
  - Address alignment requirements
- Policy enforcement requires good execution, operational control
- Inherits some limitations of DKIM and SPF
  - Does not work well with indirect mailflows that modify messages
  - *SPF pass with changed RFC5321.MailFrom not a DMARC pass*
- Not simple for small-scale senders
  - Most adoption today is B2C with large or at-risk organizations
  - Can be difficult for ESPs to support, depending on infrastructure or customer capabilities

DMARC – Design Decisions

- DMARC operates on the RFC5322.From address
  - The one author field most email clients display (at least partially)
  - The one bad actors actively exploit
  - This domain drives all DMARC policy lookups
- Why use the RFC5322.From address?
  - Owner of this domain has a clear interest in the message
  - Field should be present in all email messages
  - Shown to end-user by almost every mail client programs (MUA)
- Why not use the RFC5322.Sender address?
  - Not shown to end-user by most mail client programs
    - If it passes they will see the RFC5322.From instead of Sender
  - How would you arbitrate divergent policies between the two?

DMARC – A Few New Concepts

- Organizational Domain
- Identifier Alignment
- Reporting
  - Aggregate Reports
  - Failure Reports
- Terminology
  - Domain Owner
  - Mail Receiver
  - Report Receiver

DMARC – Organizational Domain

- Addresses in email may use one , none, or several levels of sub-domain
  - `example.com`, `mail.example.com`, `a.b.c.d.example.com`
- Organizational Domain would be the smallest name that is not a Top-Level Domain (TLD) according to Internet Assigned Name Authority (IANA)
  - `.com` is a TLD
  - `example.com` is not a TLD
- Work is underway to standardize how to detect these domain boundaries; for now, there are heuristics in RFC7489

DMARC – Identifier Alignment

- DMARC operates on the RFC5322.From address
  - This domain drives all DMARC policy lookups
- Identifier Alignment concept requires that:
  - DKIM: `d=` domain must match RFC5322.From domain
  - SPF: `smtp.mfrom` domain must match RFC5322.From domain
- For a DKIM or SPF "pass" to generate a DMARC "pass," the identifiers must meet this alignment requirement
- Two modes, `strict` or `relaxed`
  - `strict` requires an exact match between the two domains
  - `relaxed` requires that the two Organizational Domains match

DMARC – Identifier Alignment

- What does Identifier Alignment look like? Assume that:

RFC5321.MailFrom	`bounces@mail.example.net`
DKIM d= domain	`@example.com`
RFC5322.From	`all-hands@mail.example.com`

- Under strict alignment:
  - SPF does not have Identifier Alignment – no exact match
  - DKIM does not have Identifier Alignment – no exact match
- Under relaxed alignment:
  - SPF does not have Identifier Alignment – Org Domains don't match
  - DKIM does have Identifier Alignment – Org Domains do match

DMARC - Reporting

- Aggregate Reports
    - Report from a Mail Receiver of all email traffic using a given domain in the RFC5322.From
        - Doesn’t matter what source it came from, you’ll see it
    - Message counts broken out by
        - Sending IP address
        - Authentication results
        - Disposition
    - Generally sent daily, or up to several times a day, depending on the Mail Receiver
    - XML format

DMARC – Reporting – Aggregate XML

```xml
<?xml version="1.0" encoding="UTF-8" ?>
<feedback>
  <report_metadata>
    <org_name>google.com</org_name>
    <email>noreply-dmarc-support@google.com</email>
    <extra_contact_info>http://support.google.com/a/bin/answer.py?answer=2466580</extra_contact_info>
    <report_id>14093921091532388656</report_id>
    <date_range>
      <begin>1432598400</begin>
      <end>1432684799</end>
    </date_range>
  </report_metadata>
  <policy_published>
    <domain>dmarctest.org</domain>
    <adkim>r</adkim>
    <aspf>r</aspf>
    <p>none</p>
    <sp>none</sp>
    <pct>100</pct>
  </policy_published>
```

Policy this domain published during this reporting period

DMARC – Reporting – Aggregate XML

```
<record>
   <row>
     <source_ip>2607:f8b0:400e:c03::232</source_ip>
     <count>1</count>
     <policy_evaluated>
       <disposition>none</disposition>
       <dkim>pass</dkim>
       <spf>fail</spf>
     </policy_evaluated>
   </row>
   <identifiers>
     <header_from>dmarctest.org</header_from>
   </identifiers>
   <auth_results>
     <spf>
       <domain>dmarctest.org</domain>
       <result>softfail</result>
     </spf>
   </auth_results>
 </record>
```

```
<record>
   <row>
     <source_ip>72.52.75.16</source_ip>
     <count>2</count>
     <policy_evaluated>
       <disposition>none</disposition>
       <dkim>pass</dkim>
       <spf>pass</spf>
     </policy_evaluated>
   </row>
   <identifiers>
     <header_from>dmarctest.org</header_from>
   </identifiers>
   <auth_results>
     <spf>
       <domain>dmarctest.org</domain>
       <result>pass</result>
     </spf>
   </auth_results>
 </record>
```

DMARC - Reporting

Failure Reports

- Report from a Mail Receiver documenting a specific message that failed to authenticate
- Not all Mail Receivers will generate these, and each may redact different elements
- Sent when the authentication failure occurs
- Leverages ARF/AFRF per RFC6591
- Generally includes header information needed to debug authentication failures
- May include URLs and other data for investigation

- **Caution**: Abuse activity could generate millions/day!

DMARC – Additional Terminology

- Domain Owner
  - The entity that owns or has registered a DNS domain
- Mail Receiver
  - The ultimate destination for an email message
  - Mailbox Provider sometimes used
  - Usually a Report Generator too
- Report Generator
  - Entity creating and sending DMARC reports
- Report Receiver
  - Entity receiving DMARC reports sent by a Report Generator
  - Usually denotes they are receiving reports for third parties
  - Report Processor also used

DMARC Mechanics

DMARC – DNS Record Structure

```
v=DMARC1; p=none; sp=quarantine; pct=100; ri=46,200;
rua=mailto:reports@dmarc.org; ruf=mailto:reports@dmarc.org
```

Field	Meaning	Default
v	Protocol version	DMARC1
p	Policy for the domain	none
sp	Policy for any subdomains	p= value
pct	% of messages to apply policy	100
adkim	DKIM alignment mode	r
aspf	SPF alignment mode	r
rua	Aggregate reporting URI(s)	
ruf	Failure reporting URI(s)	
rf	Failure report format	afrf
ri	Aggregate reporting interval	86400
fo	Failure reporting options	0

DMARC – Policy Options

- Three policies can be requested for unauthenticated email:
  - None – Take no action ("monitor mode")
  - Quarantine – Deliver to quarantine or spam folder
  - Reject – Don't deliver the message at all
- Receivers will apply these to unauthenticated message
  - However each has exceptions for "known forwarders," mailing lists, and other things lumped under "local policy" or "receiver policy"
- The `pct=` tag intended for gradual rollout
  - `pct=50` – Mail Receiver applies requested policy to half of unauthenticated messages

DMARC – Reporting Options

- `rua` and `ruf` tags (`mailto:reports@example.com!10M`)
  - URI indicates a request for each report type, where to send it, optional size limit
  - Only URI type currently supported is `mailto:`
  - Originally included an HTTP POST method, and it could include something similar again…
  - Reports are supposed to be sent with authentication
  - Simple case, these are addresses in the same domain the DMARC record was retrieved from
    - Mail from `example.com`, `rua=mailto:reports@example.com`

DMARC – Reporting Options

- Using another domain in `rua`/`ruf` could be a DDoS
- A domain can signal that it will accept reports generated for another domain
  - External Reporting Addresses, RFC7489 Sections 7.1 and 12.5

```
example.com TXT … rua=mailto:rua@example.net
```

- The Report Generator will check for the following record:

```
example.com._report._dmarc.example.net TXT … v=DMARC1
```

- Wildcards are commonly used by report processors

DMARC – Reporting Options

- `fo` tag – failure reporting options
  - `0`: Generate report if all authentication methods fail to produce an aligned result
  - `1`: Generate report if any authentication method fails to produce aligned result
  - `d`: Generate DKIM-specific report if message has a signature that failed to verify for any reason
  - `s`: Generate an SPF-specific report if message failed SPF, aligned or not

DMARC – Retrieving DMARC Records

- Mail Receiver accepts a message
- Locates RFC5322.From header (`user@mail.example.com`)
- Extracts domain from address (`mail.example.com`)
- Prepend `_dmarc.` to the domain (`_dmarc.mail.example.com`)
- Looks up TXT record in DNS using that name
- No record found? Repeat with the Organizational Domain instead of extracted domain (`_dmarc.example.com`)

DMARC – Retrieving DMARC Records

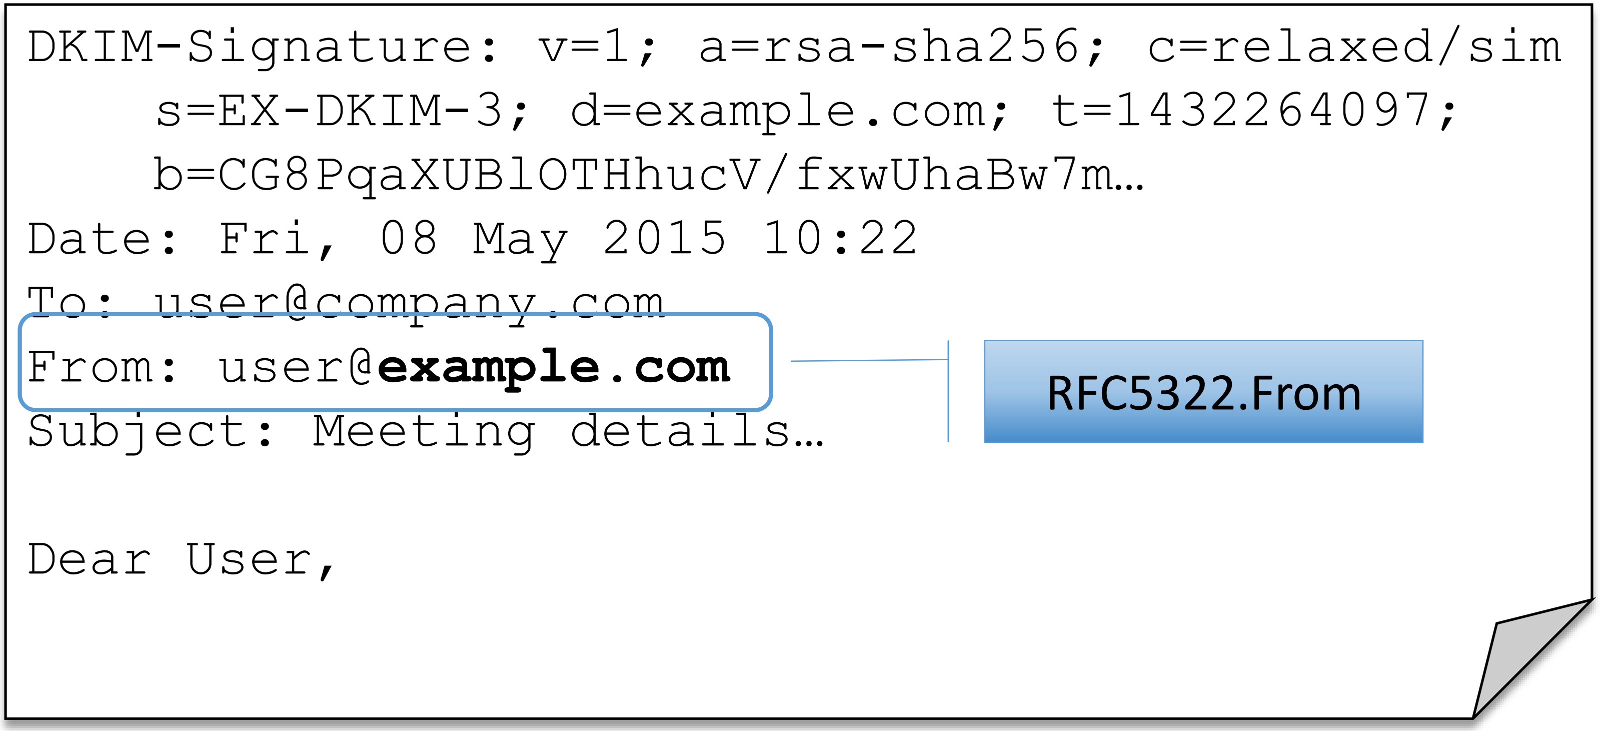

```
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/sim
    s=EX-DKIM-3; d=example.com; t=1432264097;
    b=CG8PqaXUBlOTHhucV/fxwUhaBw7m…
Date: Fri, 08 May 2015 10:22
To: user@company.com
From: user@example.com
Subject: Meeting details…

Dear User,
```

RFC5322.From

```
% dig +short txt _dmarc.example.com
"v=DMARC1\; p=none\; rua=mailto:reports@example.com; …"
```

DMARC Records and Sub-domains

- A record applies to subdomains by default
  - Publish for `example.com`, all subdomains will "inherit" when they fail to find a specific match and lookup the Organizational Domain
  - Use the `sp=` field to specify different default for all subdomains
  - Publish different policies for specific subdomains as needed

`example.com`	`p=none, sp=quarantine`
`mail.example.com`	inherits from `sp=quarantine`
`web.example.com`	publishes record with `p=reject`

- `sp=` useful to prevent fraudsters from making up subdomains
- Allows staged rollout across complex hierarchies

Questions?