

Reducing Spam

APTIVATE
International IT Development

Hi, dear qb

Subject: Our prices don't know the competition Viagra - 1.79\$

4 Vigra pills FREE for every order!!! Special internet prices!

Your Co-operative Bank Card balance is currently overdue and we require an immediate payment of £43.97.

Chris Wilson

Be so kind to contact me at your earliest convenient for a possible business deal involving money transfer of about \$14.2 Million USD.

It may interest you to hear that I am a man of PEACE and INTERGRITY

You have therefore been approved for a total sum of £6,000.00 (Six Hundred Thousand Pounds) cash credited.....

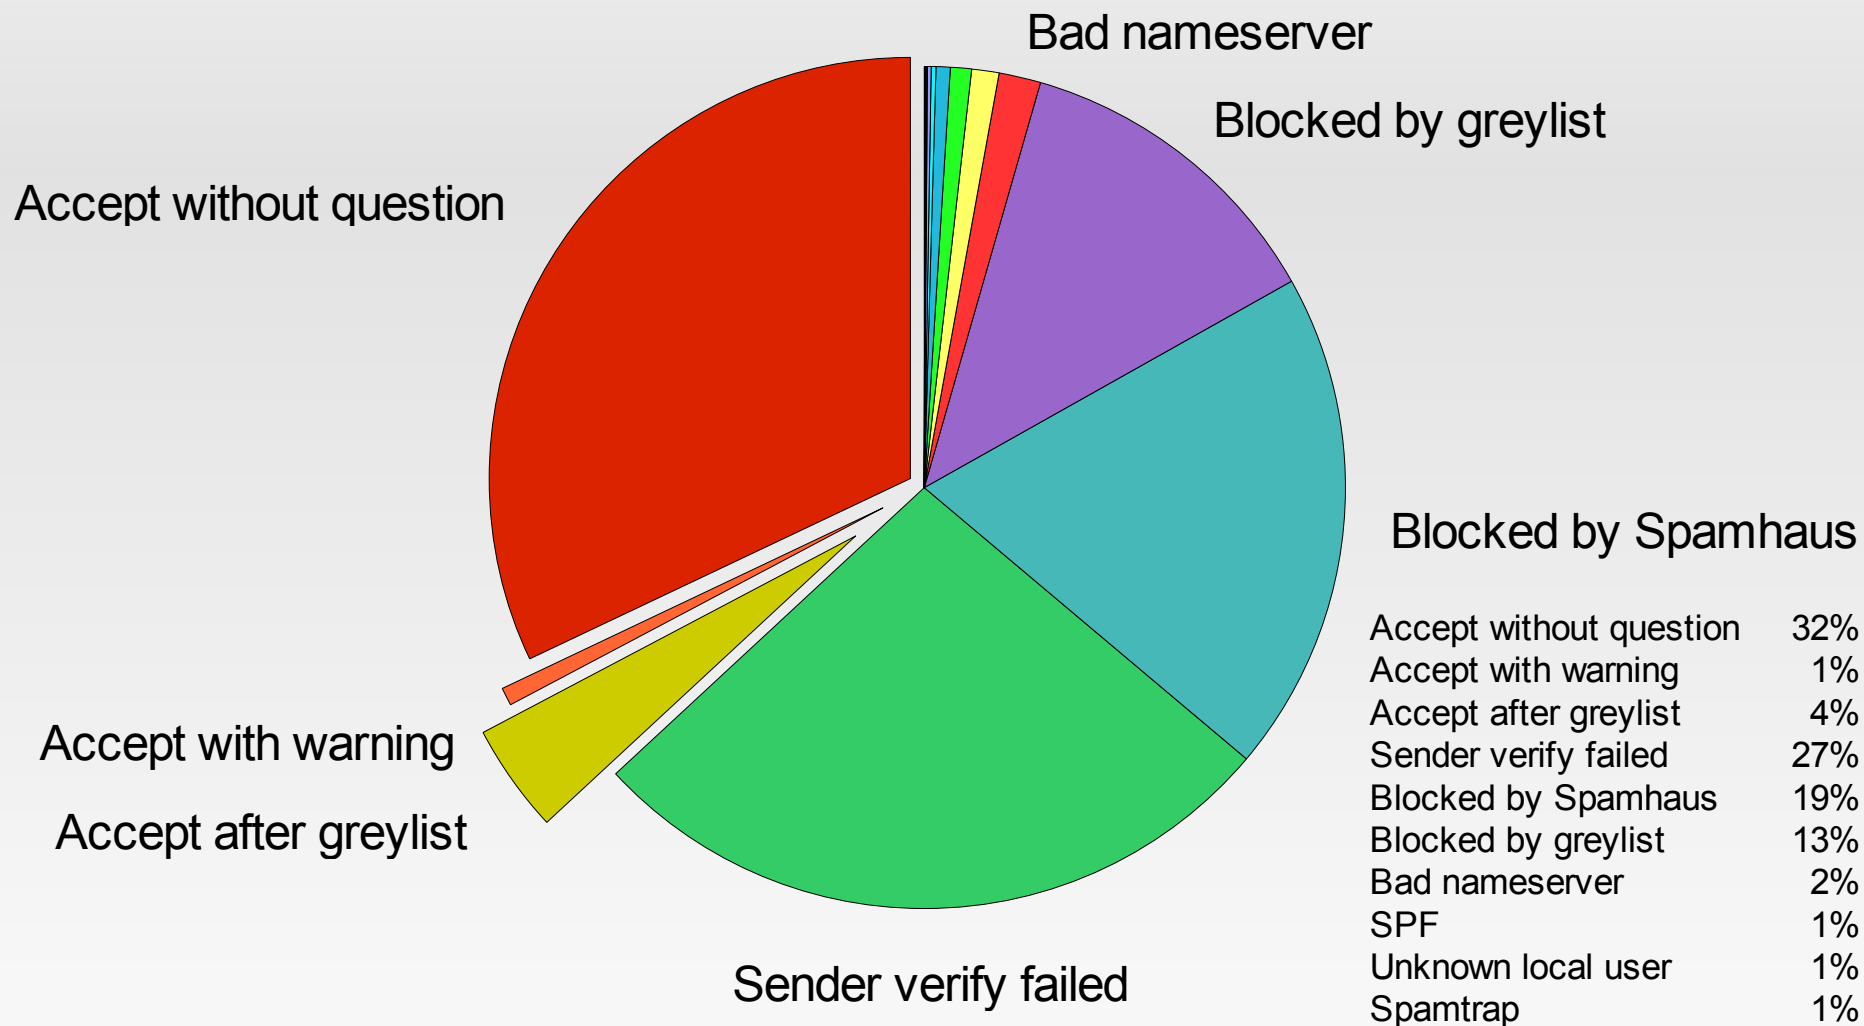


The deciphering of the Add Vinci Code discovered Jade Stewart as the descendent of the Davidic Dynasty. Her existence threatens the legitimacy of Christian orthodoxy

Structured Text

- Case Study of Aptivate Email
- Method comparison and implementation:
 - DNS Blacklists
 - Sender Verify
 - Greylisting
 - Local Policy (nameserver checks)
 - Spam Traps (Blacklisting)
 - Content scanning with DSPAM

What Happened to Our Mail?

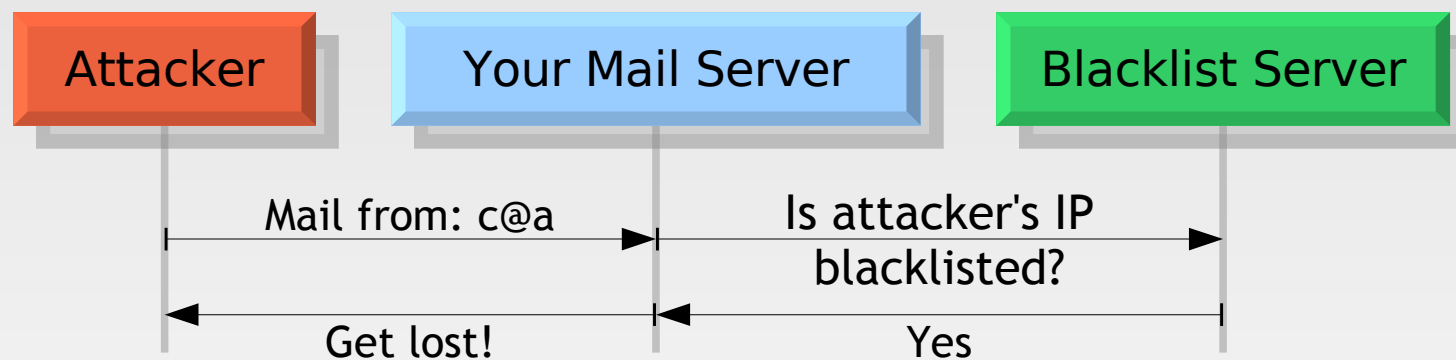


Spam Filtering Methods Tested

- DNS Black Lists (DNSBL)
- Sender Verify
- Greylisting
- Sender Policy Framework (SPF)
- Local Policy (nameserver checks)
- Spam Traps
- Content Scanning with DSPAM

DNS Blacklist Operation

“Do you have any friends?”



DNS Blacklist Implementation (1)

In the ACL controlling the response to the RCPT command, probably `acl_check_rcpt`:


```
deny
  # ! spf = pass
  ! hosts = +whitelist
  ! hosts = +relay_from_hosts
  ! authenticated = *
  dnslists = zen.spamhaus.org
  message = Your IP address ($sender_host_address) is blacklisted
by Spamhaus\n\
(http://www.spamhaus.org/query/bl?ip=$sender_host_address)\n\
$dnslist_text
```

DNS Blacklist Implementation (2)

```
warn
! spf = pass
! hosts = +whitelist
! hosts = +relay_from_hosts
! authenticated = *
dnslists = bl.csma.biz \
           : bl.spamcop.net \
           : cbl.abuseat.org ...
message = X-RBL-Warning: $sender_host_address is blacklisted by \
           $dnslist_domain ($dnslist_text)
log_message = $sender_host_address is blacklisted by \
           $dnslist_domain ($dnslist_text)
```

DNS Blacklist Comparison

“Is that your *final* answer?”

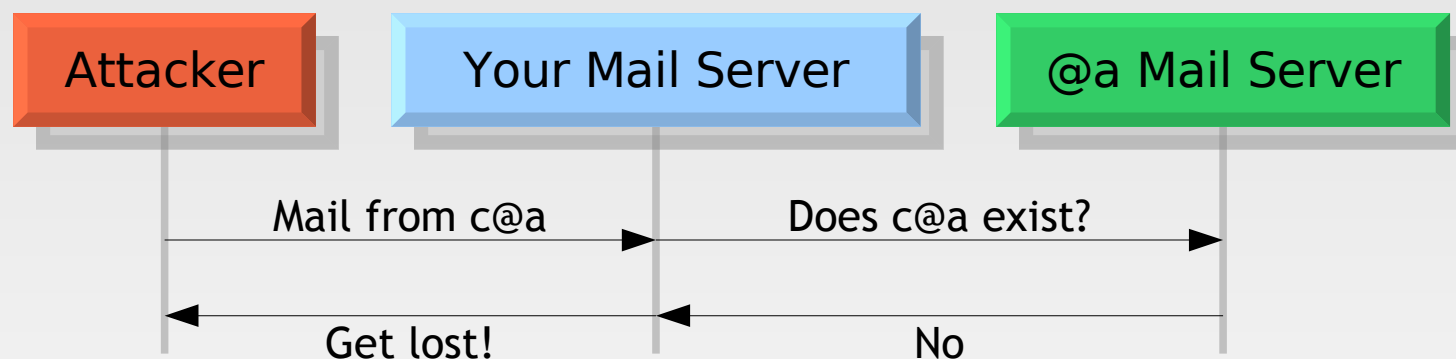


DNS Blacklist Conclusion

- Spamhaus is excellent
- PSBL is very good, but some false positives
- Spamcop is accurate, but didn't help us much
- The others may help content scanners, but don't block on them

Sender Verify Operation

“Yes, Mr. Mouse, we have your reservation.”



Sender Verify Implementation

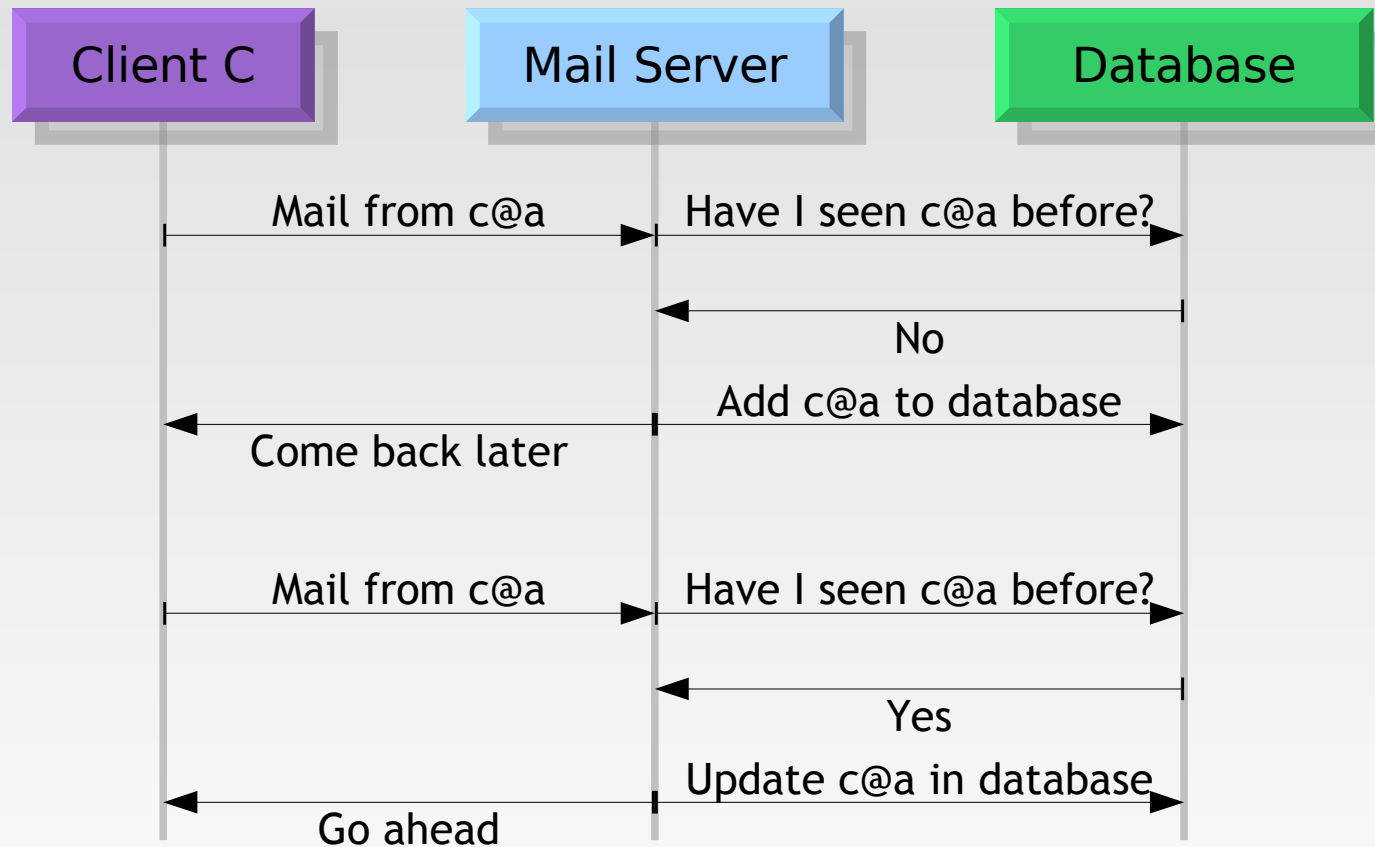
In the ACL controlling the response to the RCPT command, probably `acl_check_rcpt`:

```
# Deny unless the sender address can be verified.
```

```
require verify = sender/callout=120s  
    message = Sender verify failed
```

Greylisting Operation

Do I know you?



Greylisting Tables

```
CREATE DATABASE exim_db;
```

```
USE exim_db;
```

```
GRANT ALL ON exim_db.* TO exim@localhost IDENTIFIED BY  
'sUp3rs3cr3t';
```

```
CREATE TABLE exim_greylist  
(  
  id int(11) NOT NULL auto_increment PRIMARY KEY,  
  relay_ip varchar(21),  
  from_domain varchar(85),  
  block_expires datetime NOT NULL,  
  record_expires datetime NOT NULL,  
  origin_type enum('MANUAL','AUTO') NOT NULL DEFAULT 'MANUAL',  
  create_time datetime NOT NULL,  
  KEY exim_lookup (relay_ip,from_domain)  
);
```

Greylisting Macros (1)

Based on “*Greylisting with MySQL and Exim*” available from <http://theinternetco.net/projects/exim/greylist>

```
hide mysql_servers = localhost/exim_db/exim_user/sUp3rs3cr3t
```

```
GREYLIST_TEST = SELECT IF(NOW() > block_expires, 2, 1) \  
FROM exim_greylist \  
WHERE relay_ip = '${quote_mysql:$sender_host_address}' \  
AND from_domain = '${quote_mysql:$sender_address_domain}' \  
AND record_expires > NOW()
```

Greylisting Macros (2)

```
GREYLIST_ADD = \  
  INSERT INTO exim_greylist \  
  SET relay_ip      = '${quote_mysql:$sender_host_address}', \  
  from_domain      = '${quote_mysql:$sender_address_domain}', \  
  block_expires    = DATE_ADD(NOW(), INTERVAL 10 MINUTE), \  
  record_expires   = DATE_ADD(NOW(), INTERVAL 28 DAY), \  
  origin_type      = 'AUTO', \  
  create_time      = NOW()
```

```
GREYLIST_UPDATE = \  
  UPDATE exim_greylist \  
  SET record_expires = DATE_ADD(now(), INTERVAL 28 DAY) \  
  WHERE relay_ip      = '${quote_mysql:$sender_host_address}' \  
  AND from_domain     = '${quote_mysql:$sender_address_domain}' \  
  AND record_expires > NOW()
```

Greylisting ACL (1)

In the ACL controlling the response to the RCPT command, probably `acl_check_rcpt`:

```
warn set acl_m2 = ${lookup mysql{GREYLIST_TEST}{$value}{0}}

defer
! spf = pass
! hosts = +whitelist
! hosts = +relay_from_hosts
! authenticated = *
condition = ${if eq{$acl_m2}{0}{yes}}
condition = ${lookup mysql{GREYLIST_ADD}{yes}{no}}
message = Now greylisted - please try again in five minutes.
```

Greylisting ACL (2)

```
defer
! spf = pass
! hosts = +whitelist
! hosts = +relay_from_hosts
! authenticated = *
condition = ${if eq{$acl_m2}{1}{yes}}
message = Still greylisted - please try again in five minutes.
```

```
defer
! spf = pass
! hosts = +whitelist
! hosts = +relay_from_hosts
! authenticated = *
condition = ${lookup mysql{GREYLIST_UPDATE}{no}{no}}
message = Greylist update failed
```

Tracking Down Some Spammers

```
chris@fen-ndiyo2(~)$ whois wailmail.com
```

```
Domain servers in listed order:
```

```
PARK25.SECURESERVER.NET
```

```
PARK26.SECURESERVER.NET
```

```
chris@fen-ndiyo2(~)$ whois zehnerfamily.com
```

```
Domain servers in listed order:
```

```
PARK17.SECURESERVER.NET
```

```
PARK18.SECURESERVER.NET
```

```
chris@fen-ndiyo2(~)$ whois windowmessages.com
```

```
Domain servers in listed order:
```

```
CPNS01.SECURESERVER.NET
```

```
CPNS02.SECURESERVER.NET
```

What do they have in common? *Nameservers.*

Into the Den of Thieves

```
chris@fen-ndiyo2(~)$ google secureserver.net
```

“URIBL lists domains that appear in spam...”

Top 50 name server hosters with URIBL Listed Domains (<http://rss.uribl.com/ns/>)

- name-services.com (#1 with 460 domains)
- secureserver.net (#8 with 199 domains)
- ip4dns.com (#13 with 100 domains)

Local Policy

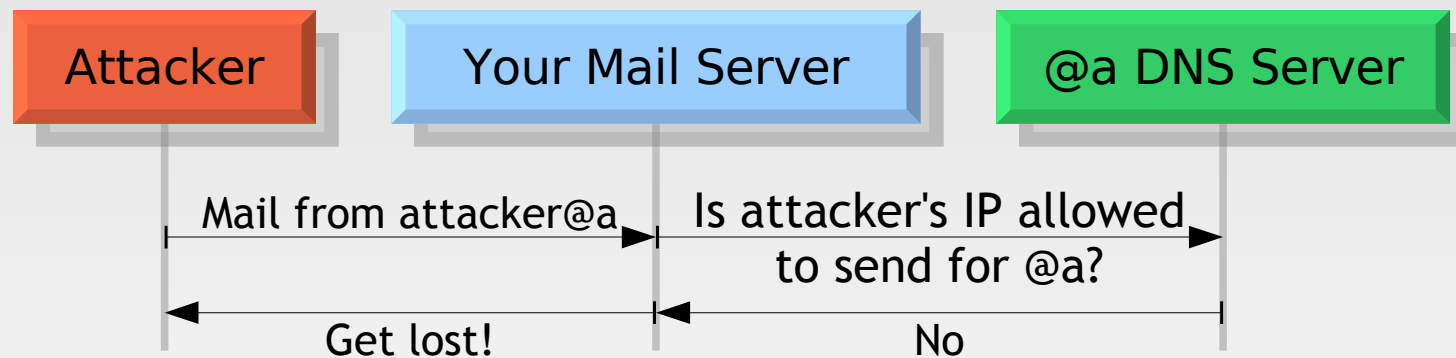
Do you really want to receive mail from companies that support spammers?

Do you care about blocking mail from GoDaddy customers?

```
deny
  condition = ${if match \
    ${lookup dnsdb {zns=${sender_address_domain}}}} \
    {.*\.seureserver\.net}}
  message = "You look like a spammer to me (ns=seureserver.net)"
```

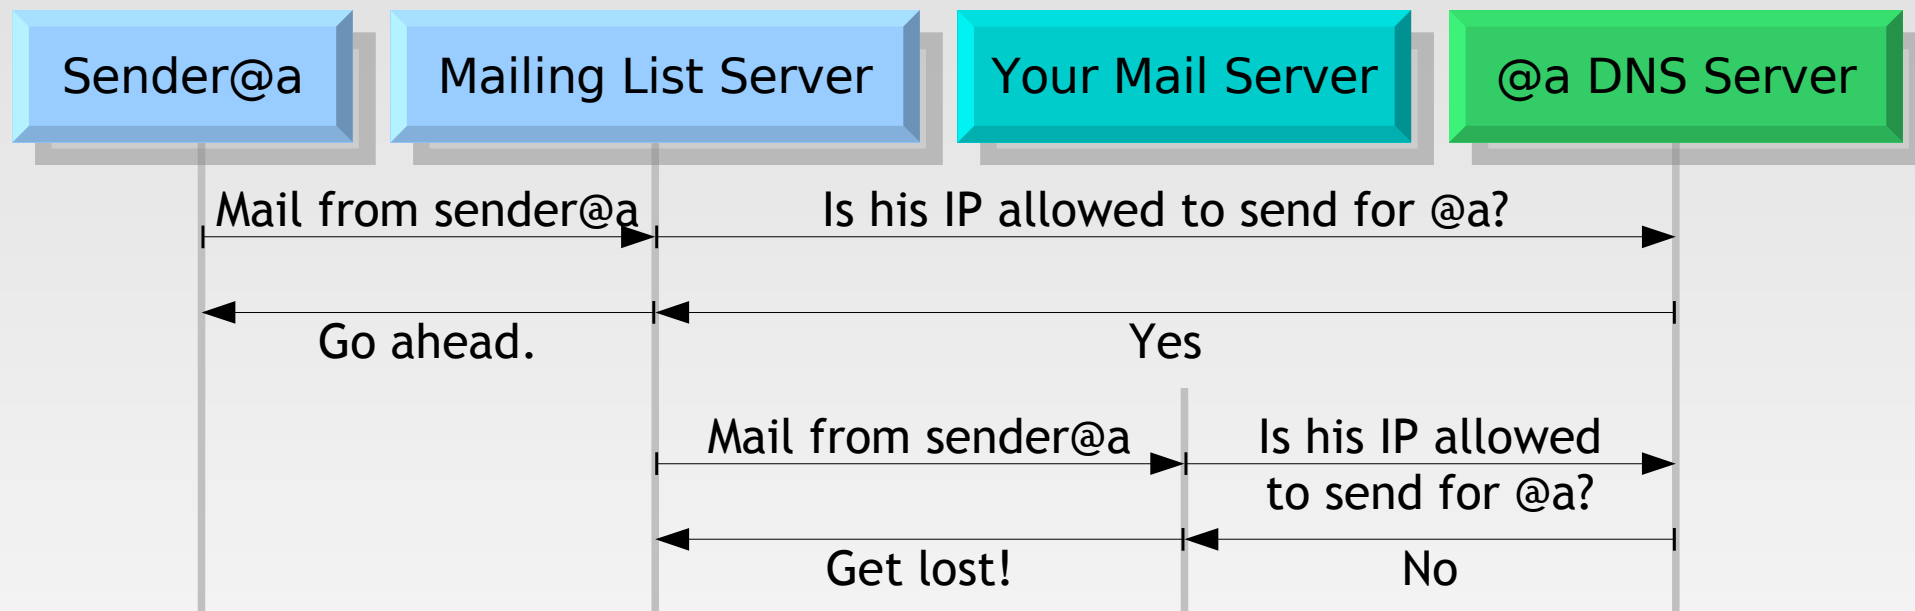
SPF In Theory

I'll need form X32, signed in triplicate.



SPF In Practice

Q: How does this work with mailing lists?



A: It doesn't.

SPF Implementation

SPF is still useful, because it allows domain owners to set policy if they don't mind the mailing list problem (e.g. they don't use them).

deny

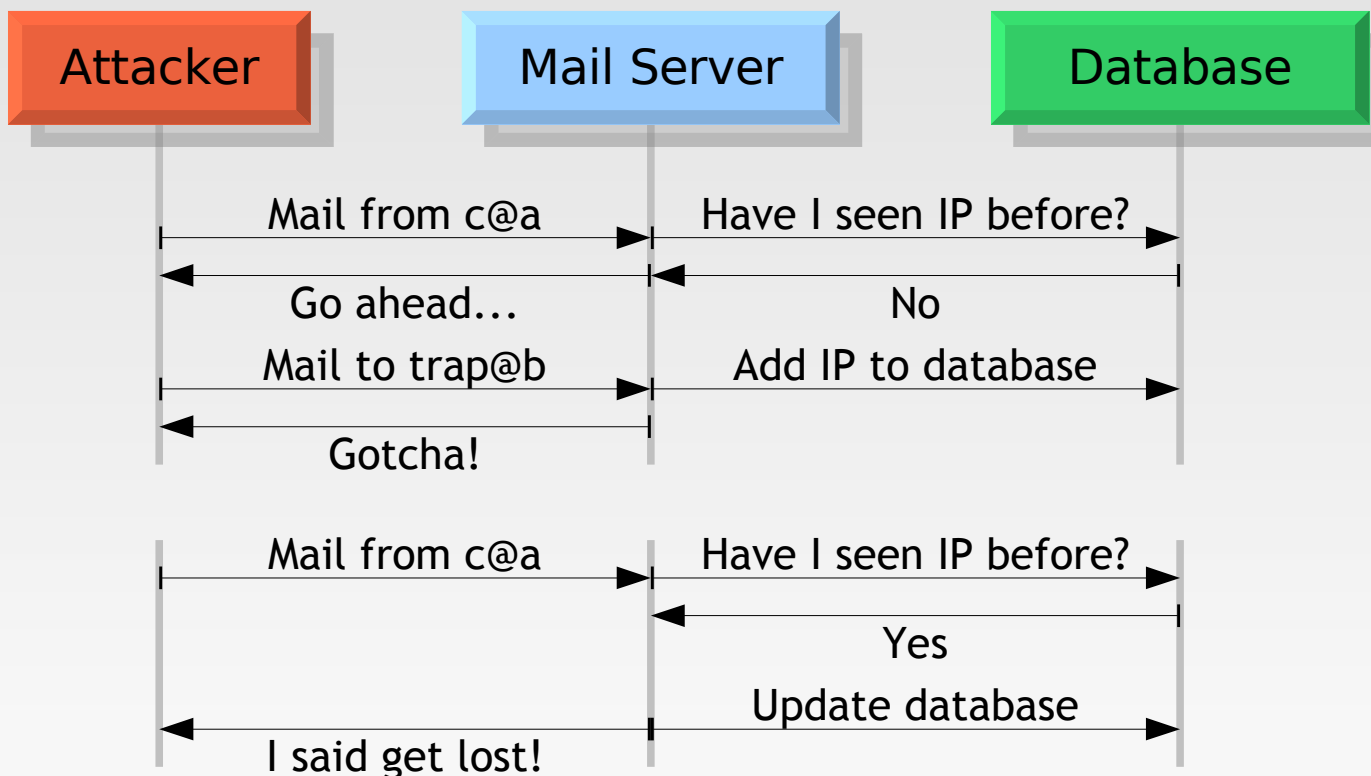
```
! hosts = +whitelist
! hosts = +relay_from_hosts
! authenticated = *
```

```
spf = fail
```

```
message = Your IP address ($sender_host_address) is not allowed
to send mail from $sender_address_domain\n\
(SPF check failed: see http://www.openspf.org/)\n\
$spf_smtp_comment
```

Spam Trap Overview

‘Diplomacy is the art of saying “Nice doggie” until you can find a rock.’



Spam Trap Macros

```
BLACKLIST_TEST = SELECT 1 \
    FROM    exim_blacklist \
    WHERE   relay_ip = '${quote_mysql:$sender_host_address}' \
    AND     NOW() < expires
```

```
BLACKLIST_ADD = INSERT INTO exim_blacklist \
    SET relay_ip = '${quote_mysql:$sender_host_address}', \
        expires = DATE_ADD(NOW(), INTERVAL 1 DAY), \
        created = NOW(), \
        sender = '${quote_mysql:$sender_address}', \
        recipient =
    '${quote_mysql:$original_local_part@$original_domain}'
```

```
BLACKLIST_UPDATE = UPDATE exim_blacklist \
    SET expires = DATE_ADD(NOW(), INTERVAL 1 WEEK) \
    WHERE relay_ip = '${quote_mysql:$sender_host_address}'
```

Spam Trap ACL

```
warn set acl_m3 = ${lookup mysql{BLACKLIST_TEST}{$value}{0}}
```

```
deny
```

```
! hosts = +whitelist_hosts
```

```
! senders = +whitelist_users
```

```
! authenticated = *
```

```
condition = ${if eq{$acl_m3}{1}{yes}}
```

```
condition = ${lookup mysql{BLACKLIST_UPDATE}{yes}{yes}}
```

```
message = You are still blacklisted for hitting a spam trap
```

```
deny
```

```
! hosts = +whitelist_hosts
```

```
! senders = +whitelist_users
```

```
! authenticated = *
```

```
recipients = dominc@aidworld.org
```

```
condition = ${lookup mysql{BLACKLIST_ADD}{yes}{yes}}
```

```
message = You are now blacklisted for hitting a spam trap (1)
```

What Else Can We Do

- We filtered out a lot of spam so far, but:
- I still get a lot of spam!
- 3000 messages in 9 months
- About 20% of my mail is spam
- Need to look at message contents

DSPAM Content Scanner



- DSPAM is “a scalable and open-source content-based spam filter”
- Designed for multi-user enterprise systems
- Supports many different MTAs
- Many new algorithms and approaches
- Command line and web-based interfaces

DSPAM Integration

You can integrate DSPAM directly into Exim, but we haven't tested this. Not all of our users use DSPAM.

We use Procmail anyway, so we call DSPAM from Procmail.

Procmail Implementation

```
# This router runs procmail for users who have a .procmailrc file
```

```
procmail:
```

```
  driver = accept
  check_local_user
  transport = procmail_pipe
  require_files =
  ${local_part}:${home}:${home}/.procmailrc:/usr/bin/procmail
  no_verify
```

```
# This transport is used for procmail
```

```
procmail_pipe:
```

```
  driver = pipe
  command = "/usr/bin/procmail"
  return_path_add
  delivery_date_add
  envelope_to_add
```

DSPAM in Procmail

In ~/.procmailrc for each user who runs DSPAM:

```
SPAMCAN = /home/chris/mail/spam
```

```
:0wc: dspam.lock
```

```
* for chris\+spam@
```

```
  | /usr/bin/dspam --source=error --class=spam
```

```
:0wc: dspam.lock
```

```
* for chris\+ham@
```

```
  | /usr/bin/dspam --source=error --class=innocent
```

```
:0fw: dspam.lock
```

```
  | /usr/bin/dspam --user chris --mode=teft --stdout --  
  deliver=spam,innocent
```

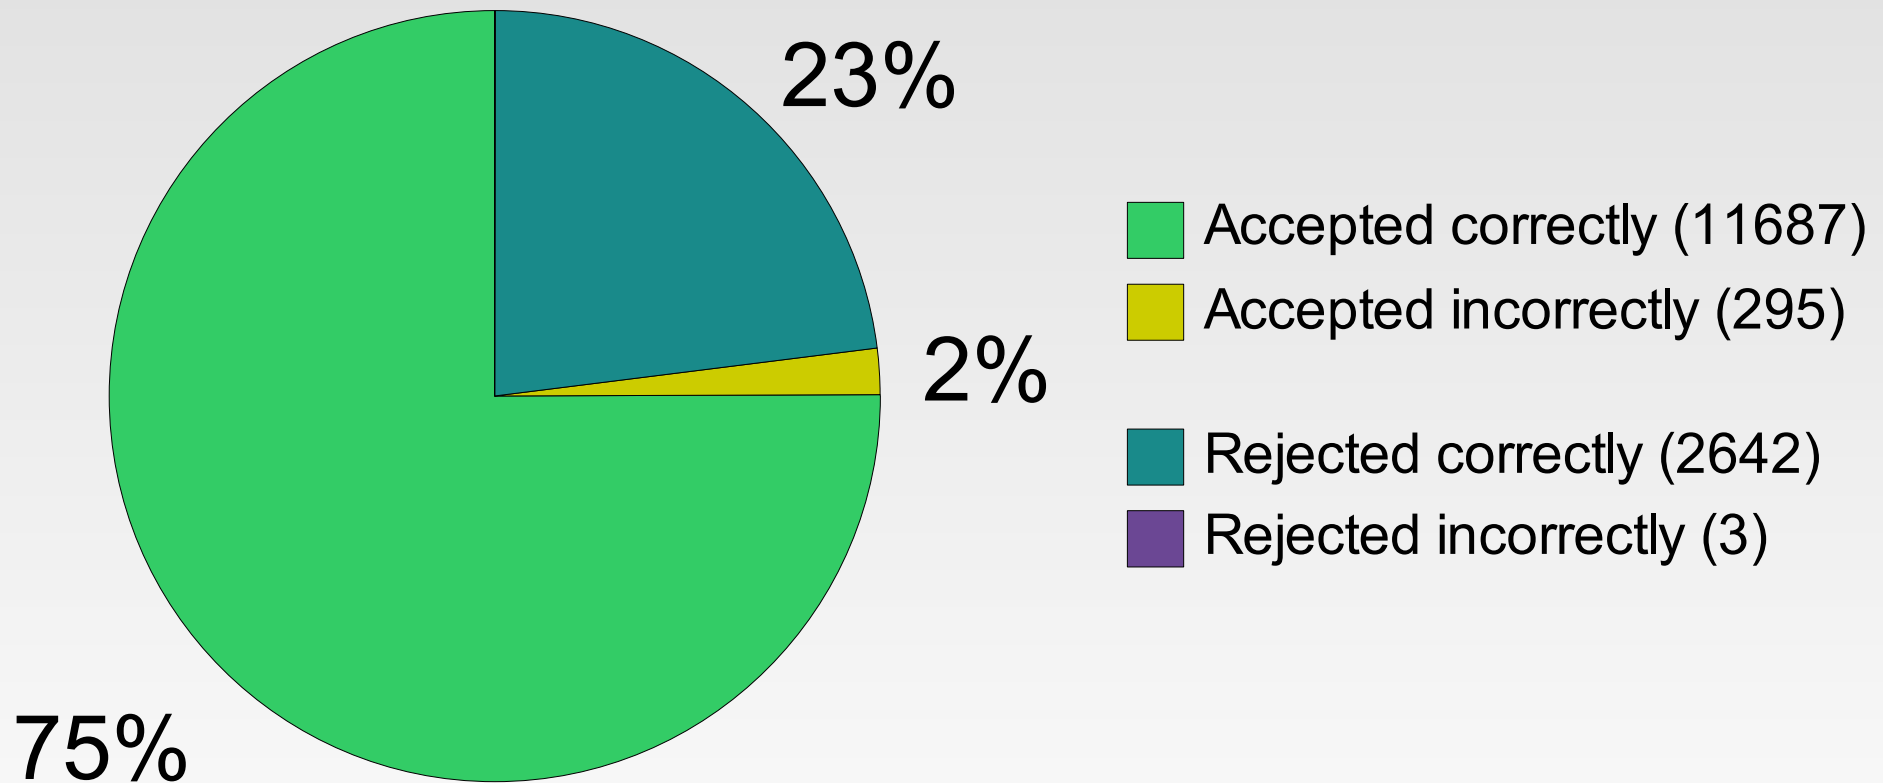
```
:0:
```

```
* X-DSPAM-Result: Spam
```

```
$SPAMCAN
```

DSPAM Results

How well does it work?



Conclusion

How well did we do?

- 63% of mail rejected as spam
- 28% good mail delivered
- 9% filed in spam folder by DSPAM
- 2% false negatives (failed to mark spam)
- 0.1% false positives (wrongly marked as spam)

Success!

Health Warning

All of our stunts were carried out by certifiable professionals.

NO WARRANTY - Not even that you won't lose your job by following my advice!

A few sub-humans were harmed in the making of this presentation.