How spammers break email for the rest of us

Andrew Richards, May 2009.

www.acrconsulting.co.uk

Outline

- About SMTP
- SMTP functionality used and abused
- 'Improving' SMTP with DKIM, SPF, BATV etc
- SMTP-time measures
- Graph showing these measures in action
- Conclusions, questions

SMTP – What it is, what's wrong with it

- SMTP: Simple Mail Transport Protocol – RFC 5321
- Relatively old and simple protocol
- Predates spam – so it's quite naïve.
- Messages get put in 'envelopes' (= instructions to mail server)
- Various later extensions to enhance functionality.

SMTP: Message & Envelope

```
Date: Thu, 28 May 2009 11:23:05 +0100
From: Major Tom <tom@outer.space>
Subject: Space
To: Ground Control <gc@planet.earth>

Dear Ground Control,

The stars look very different today.
```

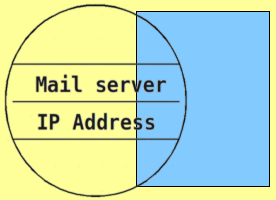


```
MAIL FROM:<tom@outer.space>
RCPT TO:<gc@planet.earth>
```

SMTP Sample Session

```
(Connect to mail server for the domain planet.earth on port 25)
(Establish TCP connection)
220 mx.planet.earth ESMTP
HELO sputnik.outer.space
250 mx.planet.earth
MAIL FROM:<tom@outer.space>
250 ok
RCPT TO:<gc@planet.earth>
250 ok
DATA
354 go ahead
Date: Thu, 28 May 2009 11:23:05 +0100
From: Major Tom <tom@outer.space>
Subject: Space
To: David <david@planet.earth>

Dear Ground Control,

The stars look very different today.
.
250 ok 1064691857 qp 3569
quit
221 mx.planet.earth
(Connection closes)
```

Receiving system knows the sending IP address,

the stated machine name,

the stated sender,

the destination mailbox,

and the actual message.

Open relays

Just accept email from/to anywhere. A gift for spammers.

Easy to close this loophole. Also use RBL services to reject messages from open relays.

Downsides:

- RBLs mean you no longer accept mail from anywhere.
- Need to authenticate senders: Roaming users can no longer trivially email.
- More work for users and sysadmins.
- Open relay detection may generate false positives.

VRFY & dictionary attacks

One old trick for a spammer get their list of addresses to spam:

```
VRFY <smith@some.domain>
250 John Smith <smith@some.domain>
```

So hobble VRFY,

```
VRFY peter
252 send some mail, i'll try my best
```

Result: SMTP degraded. Sender can't determine if message is deliverable with VRFY.

Dictionary attacks now unusual since more effective methods exist.

Issues with bounces

- A valid email address like ***yours*** may make a spammer's message look more authentic to spam filters.
- You get bounces for messages you didn't send.
- Malware can be spread by such bounces if opened.
- Result: Some mail systems choose to destroy or quarantine incoming bounces – degrading reporting of unsuccessful deliveries.

'Improving' SMTP with DKIM, SPF, BATV etc

To address some of these issues:

- DKIM: Server-to-server signing (end user unaffected)
- SPF: Check sending machine. It makes assumptions about mail flows, so breaks forwarding despite workarounds.
- BATV: Ensures bounces are valid by adding a cookie to the envelope sender. Breaks some valid mailing list setups.
- Recipient verification: Check destination mailbox(es) exists before accepting message. Dictionary attacks possible.

Result: Mail servers more complex; new problems may be introduced.

SMTP-Time: Lightweight measures

Consider how spam is sent: Esp. botnets. Try not to accept messages!

Measures that are lightweight to implement (don't need the message itself):

- Block if no reverse DNS address
- Block early talkers (greetdelay)
- Recipient verification
- Greylisting ('Try again later') (but whitelist large 'Cloud' senders, online signups)
- Nolisting (implement carefully - ensure NO response not temp fail)

Valid senders get a rejection for false positives.

SMTP-Time: Resource-intensive measures

Scan before accepting message (CPU/bandwidth intensive, need the message):

- Spam
- Viruses

Valid senders get a rejection for false positives rather than their message disappearing into quarantine.

By scanning at SMTP-time more messages can be rejected, shrinks quarantine folders.

Putting it all together

Here's some data from a small mail system showing how much can be rejected before scanning:

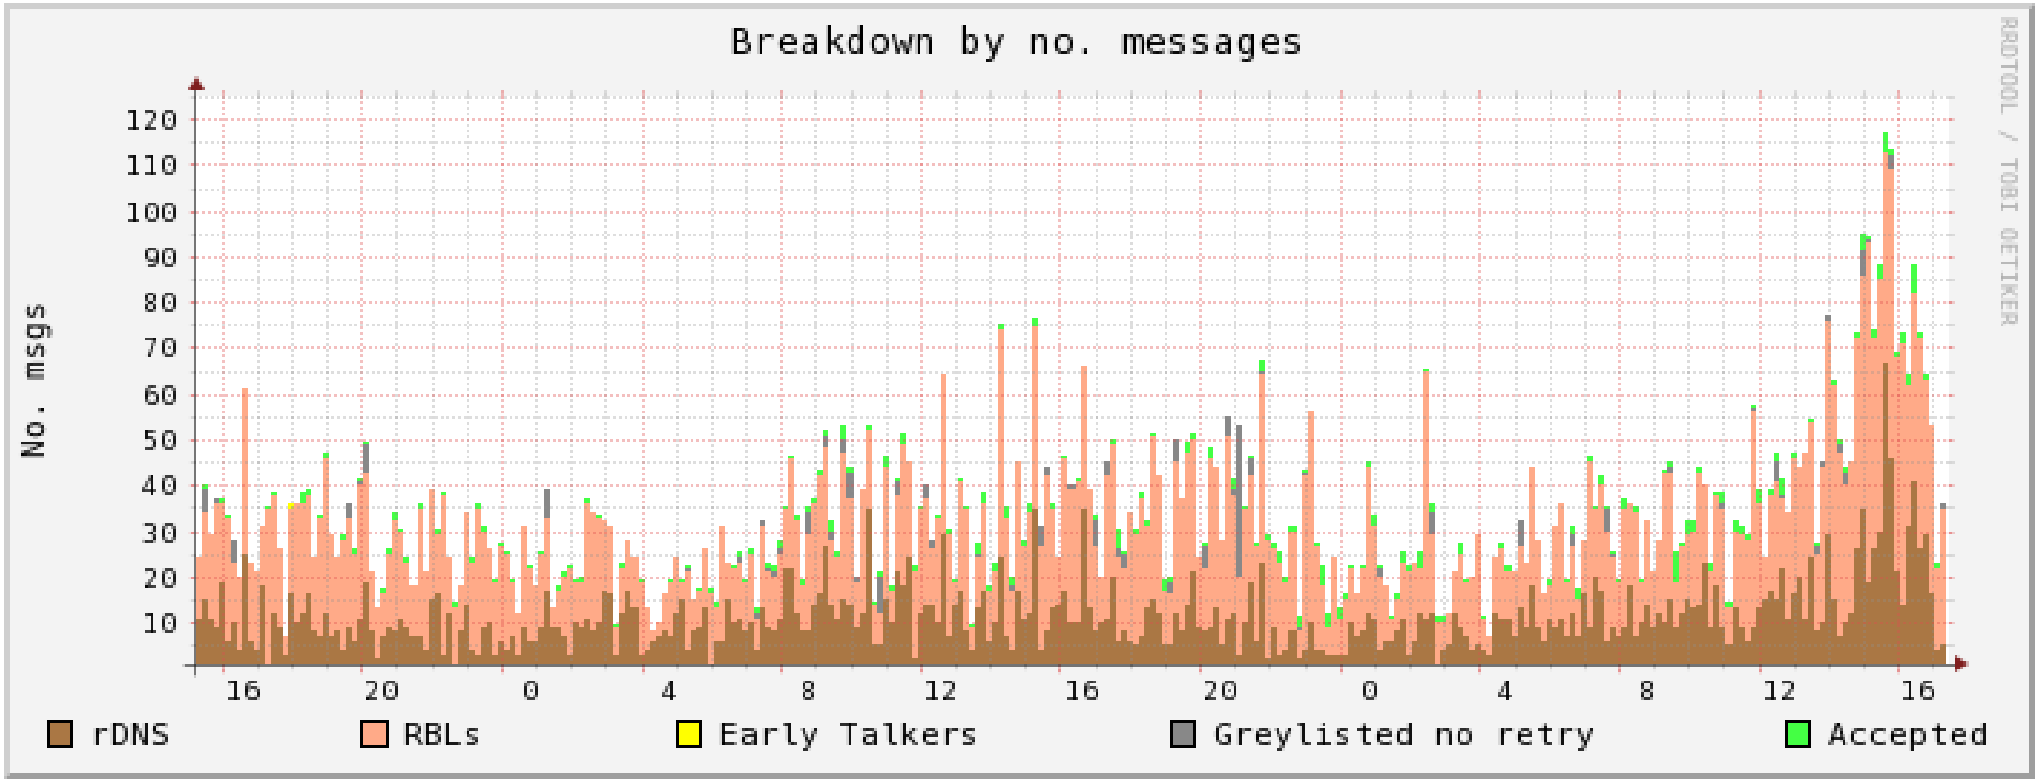

www.acrconsulting.co.uk

Putting it all together: As percentages

Expressed in percentage terms:

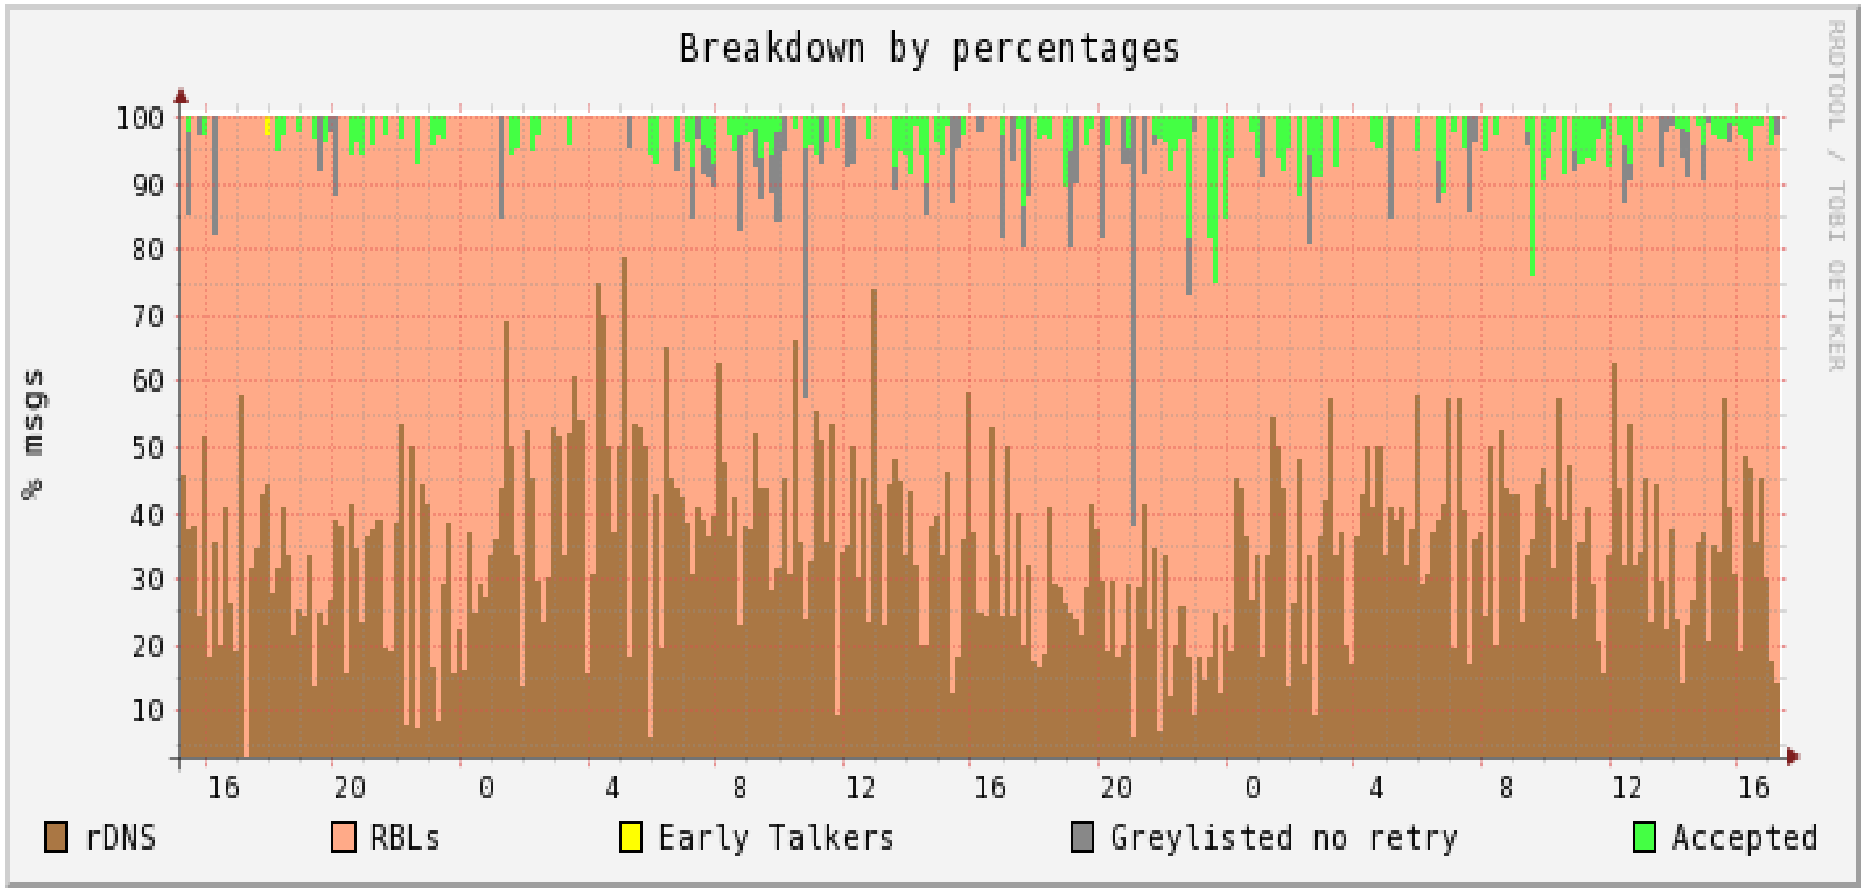

Putting it all together: In numbers

	No. messages/conns.	% messages/conns.
No reverse DNS	3424	34.11
Listed in chosen RBLs	6210	61.86
Early talkers	1	0.01
Greylisted: Didn't resend	197	1.96
Accepted	206	2.05

How SMTP has been degraded

Some thoughts on how email has changed:

- Harder for users and sysadmins to setup/configure mail systems – harder to send messages!
- Broken spam counter-measures result in disappearing messages, or valid mail flows being disrupted
- Quarantining hides valid messages from recipients
- Email address harvesting means techies are less likely to contribute to mailing lists

Conclusions, questions

- Degradation of SMTP ==> Slower, less reliable email
- Various retrofitting possibilities, some useful, some troublesome
- Optimise SMTP-time measures to provide better information to valid senders/recipients and reduce quarantining/discards.
- Lightweight SMTP-level measures scale well.
- Easy to inadvertently block valid email

Questions