IBM Research

Zombies suck the life out of the mail server

(“new developments” from LISA 2010 presentation)

Wietse Venema
IBM T. J. Watson Research Center
Hawthorne, NY, USA

Changing threats

- 2009: You built a mail system that has world-class email delivery performance.
  - *Problem*: your world-class performing mail system is now spending most of its resources *not delivering mail*.
    - *Solution*: work smarter.

Changing threats

92% Mail is spam, 95% spam is from botnets

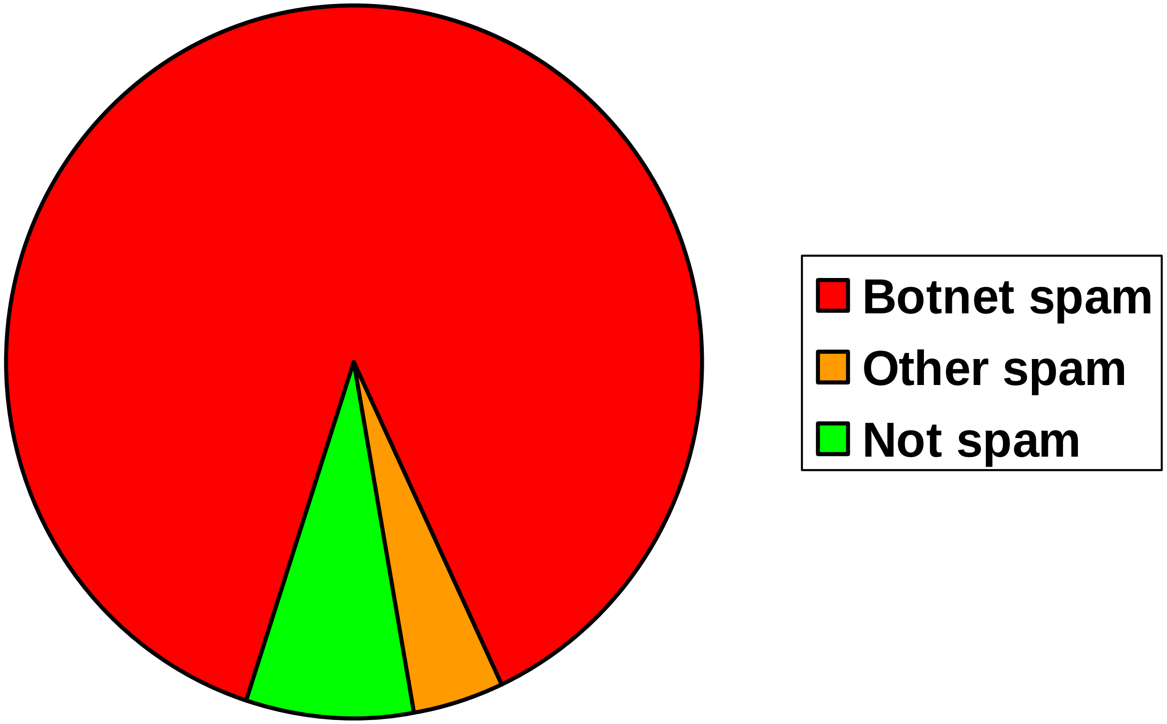

Source: MessageLabs Intelligence report, August 2010

Changing threats

Zombies keep mail server ports busy

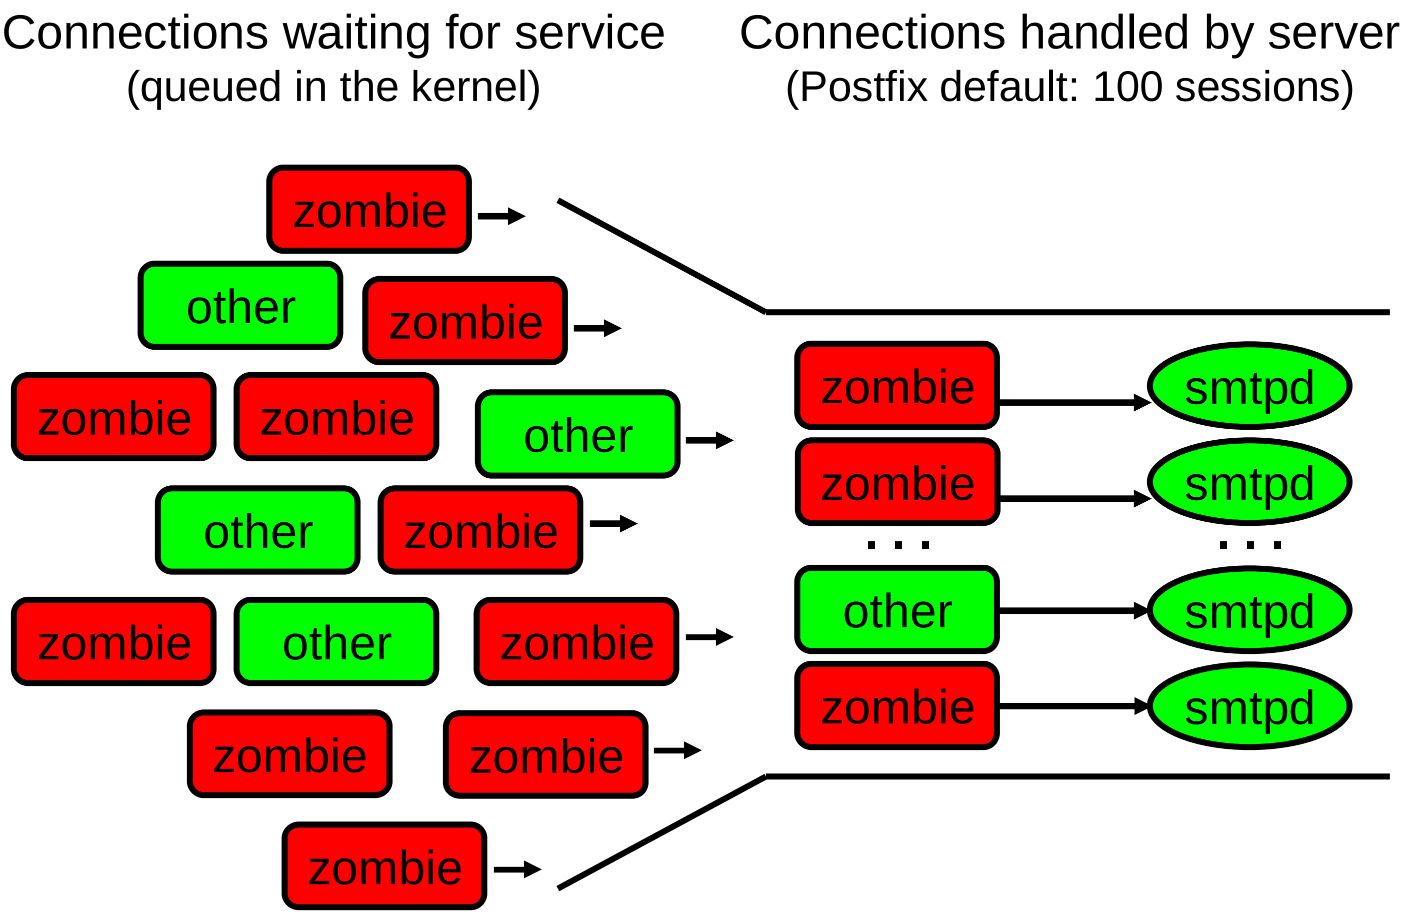

Changing threats

Zombies suck the life out of the mail server

- Worst-case example: Storm botnet.

```
13:01:36 postfix/smtpd: connect from [x.x.x.x]

13:01:37 postfix/smtpd: reject: RCPT from [x.x.x.x]:
          550 5.7.1 blah blah blah

13:06:37 postfix/smtpd: timeout after RCPT from [x.x.x.x]
```

  - RFC 5321 recommends 5-minute server-side timeout.
    - Postfix implements SMTP according to the standard.
      - Result: all SMTP server ports kept busy by Storm zombies.

Changing threats

Mail server overload strategies

Targeting small- and mid-size sites primarily

- Assumption: the zombie problem will get worse before things improve (if ever).
- Temporary overload:
  - Work faster: less time per SMTP client (load shedding).
- Persistent overload:
  - Work harder: handle more SMTP clients (forklift solution).
  - Work smarter: stop spambots up-stream (postscreen).

Temporary overload strategy

- Work faster: spend less time per SMTP client.
  - Reduce time limits, number of rejected commands, etc.
    - Automatic configuration switch in 21 lines of code (2007).
  - Will delay *some* legitimate email.
    - From sites with large network latency or packet loss.
    - From mailing lists with aggressive timeouts.
  - Better to receive *some* legitimate mail, than *no mail*.
    - OK as long as the overload condition is temporary.

Changing threats

Persistent overload strategies

- Work harder: configure more mail server processes.
  - The brute-force, fork-lift approach.
  - OK if you can afford network, memory, disk, and CPU.
- Work smarter: keep the zombies away from the server.
  - Before-server connection filter.
  - More SMTP processes stay available for legitimate email.

Changing threats

Persistent overload - before-smtpd connection filter

Prior work: OpenBSD spamd, MailChannels TrafficControl, M.Tokarev

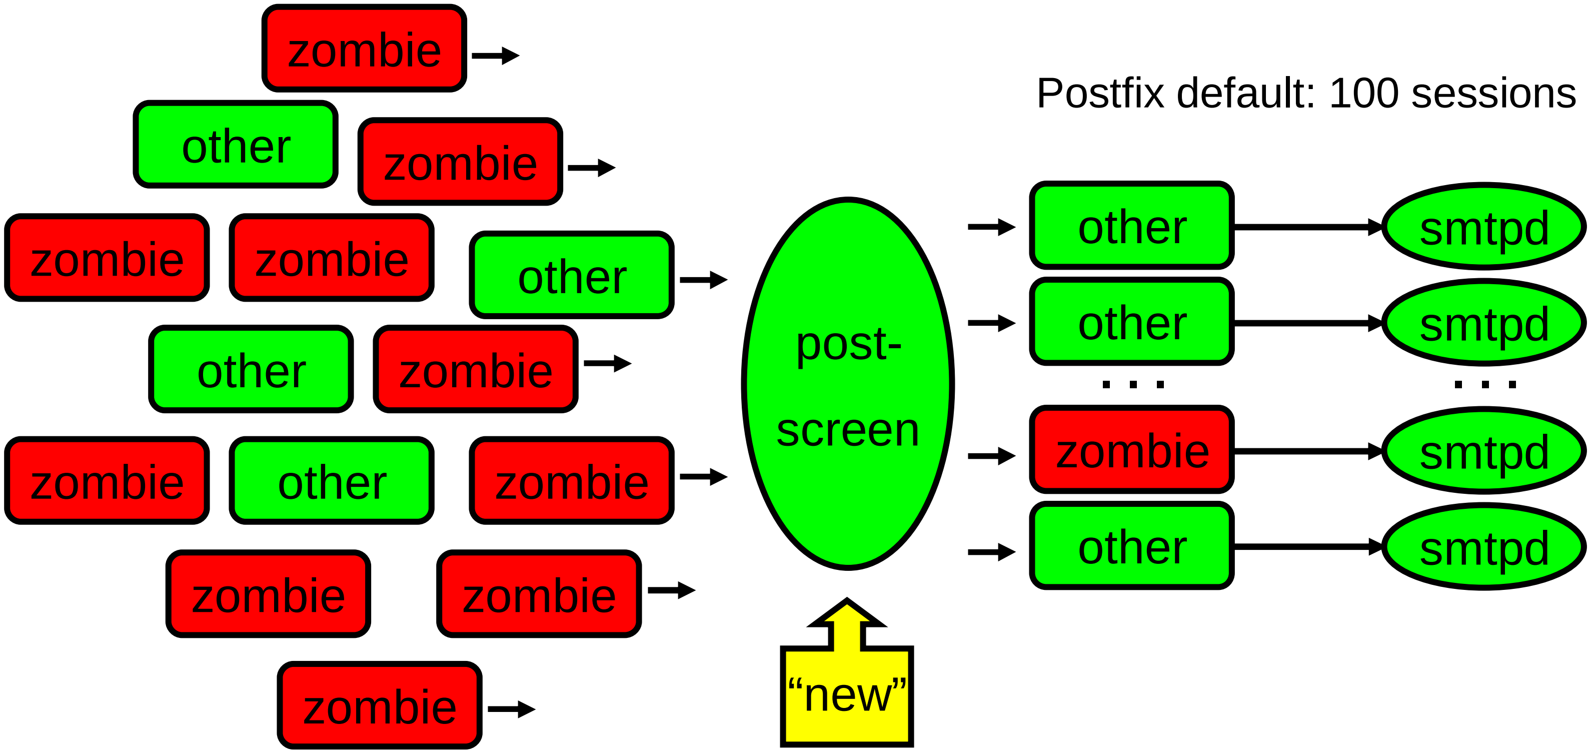

Changing threats

postscreen(8) challenges and opportunities

- Zombies are blacklisted within a few hours[1].
  - Opportunity: reject clients that are in a hurry to send mail.
    - Clients that talk too fast: pregreet, command pipelining.
    - Other blatant protocol violations.
    - Fake “temporary” error when stranger connects (greylisting).
- Zombies avoid spamming the same site repeatedly.
  - Challenge: decide “it’s a zombie” for single connections.
    - Use DNS white- and blacklists as shared intelligence source.

[1] Chris Kanich et al., Spamalytics: An Empirical Analysis of Spam Marketing Conversion, CCS 2008.

DNS white- and blacklists for email etc.

- Originally conceived by Paul Vixie of ISC.
  - The Internet Software Consortium provides reference implementations of DNS, DHCP and more.
  - To find out if address *1.2.3.4* is listed at *mail.abuse.org*, ask for the IP address of *4.3.2.1.mail.abuse.org*.
- Popular providers: spamhaus.org, spamcop.net, barracudacentral.org.
  - Spam traps and other sensors.
  - Some DNS[BW]L providers are free for small users.

postscreen(8) workflow

One daemon screens multiple connections simultaneously

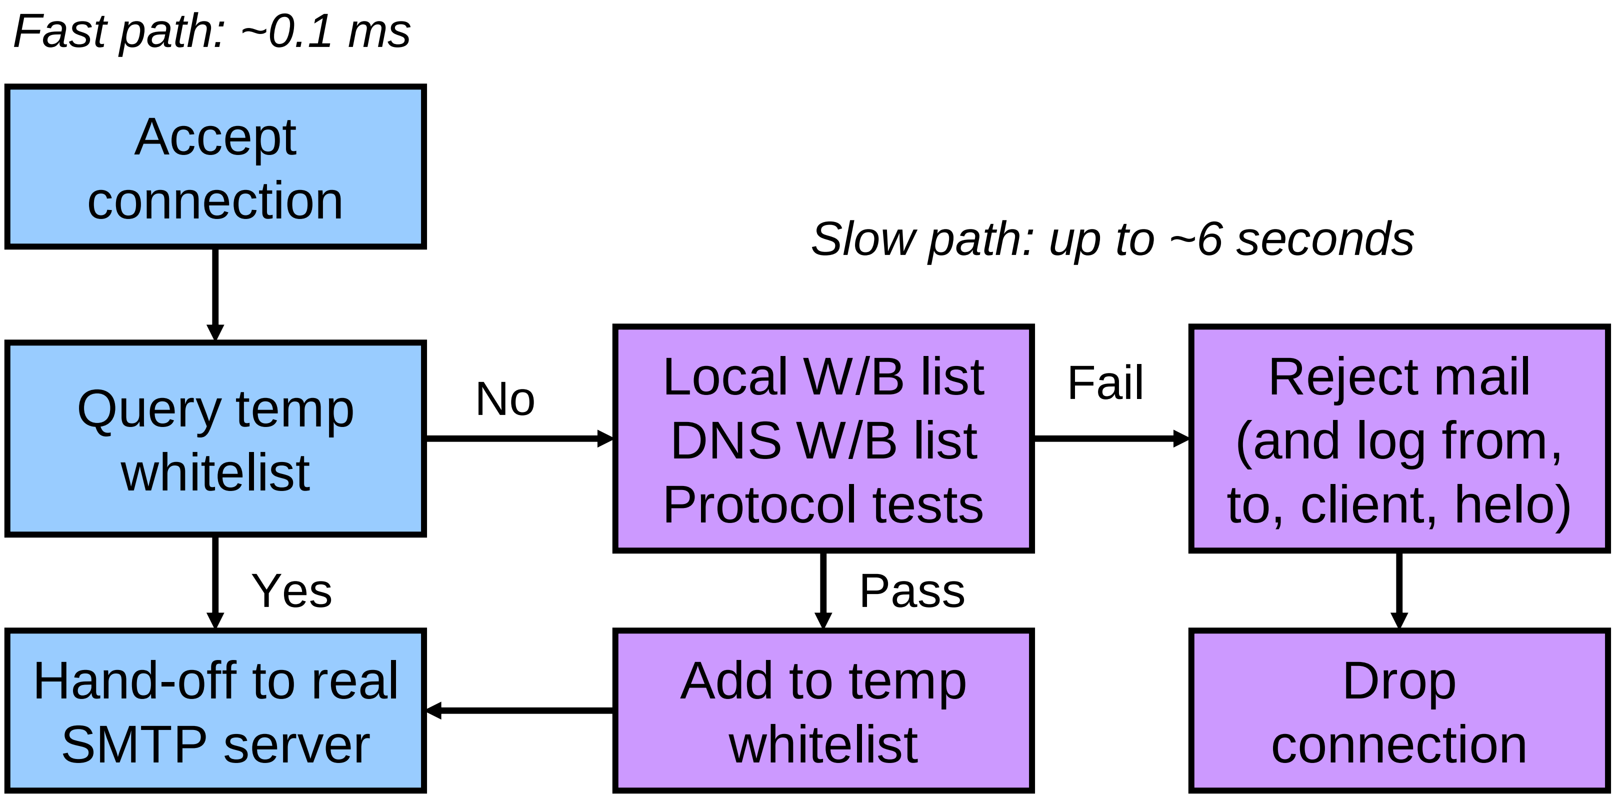

Changing threats

Detecting spambots that speak to early (pregreet)

- Good SMTP clients wait for the SMTP server greeting:

SMTP server: **220 server.example.com ESMTP Postfix<CR><LF>**

SMTP client: **EHLO client.example.org<CR><LF>**

- Sendmail *greet_pause* approach: wait several seconds before sending the 220 greeting.
  - Very few clients greet too early.
  - More clients just give up after a few seconds.
  - Manual whitelisting.

Question for dog catchers

- Q: How do I quickly find out if a house has a dog?
- A: Ring the doorbell, and the dog barks immediately.

- postscreen(8) uses a similar trick with botnet zombies.

Changing threats

Making zombies bark - multi-line greeting trap

- Good clients wait for the full multi-line server greeting:

```
mail server:  220–server.example.com ESMTP Postfix<CR><LF>
mail server:  220  server.example.com ESMTP Postfix<CR><LF>
good client:  HELO client.example.org<CR><LF>
```

- Many spambots talk immediately after the first line of the multi-line server greeting:

```
postscreen:   220–server.example.com ESMTP Postfix<CR><LF>
spambot:      HELO i-am-a-bot<CR><LF>
```

Over 60% of bots pregreet at mail.charite.de

8% Not on DNS blacklists. Berlin, Aug 26 – Sep 29, 2010

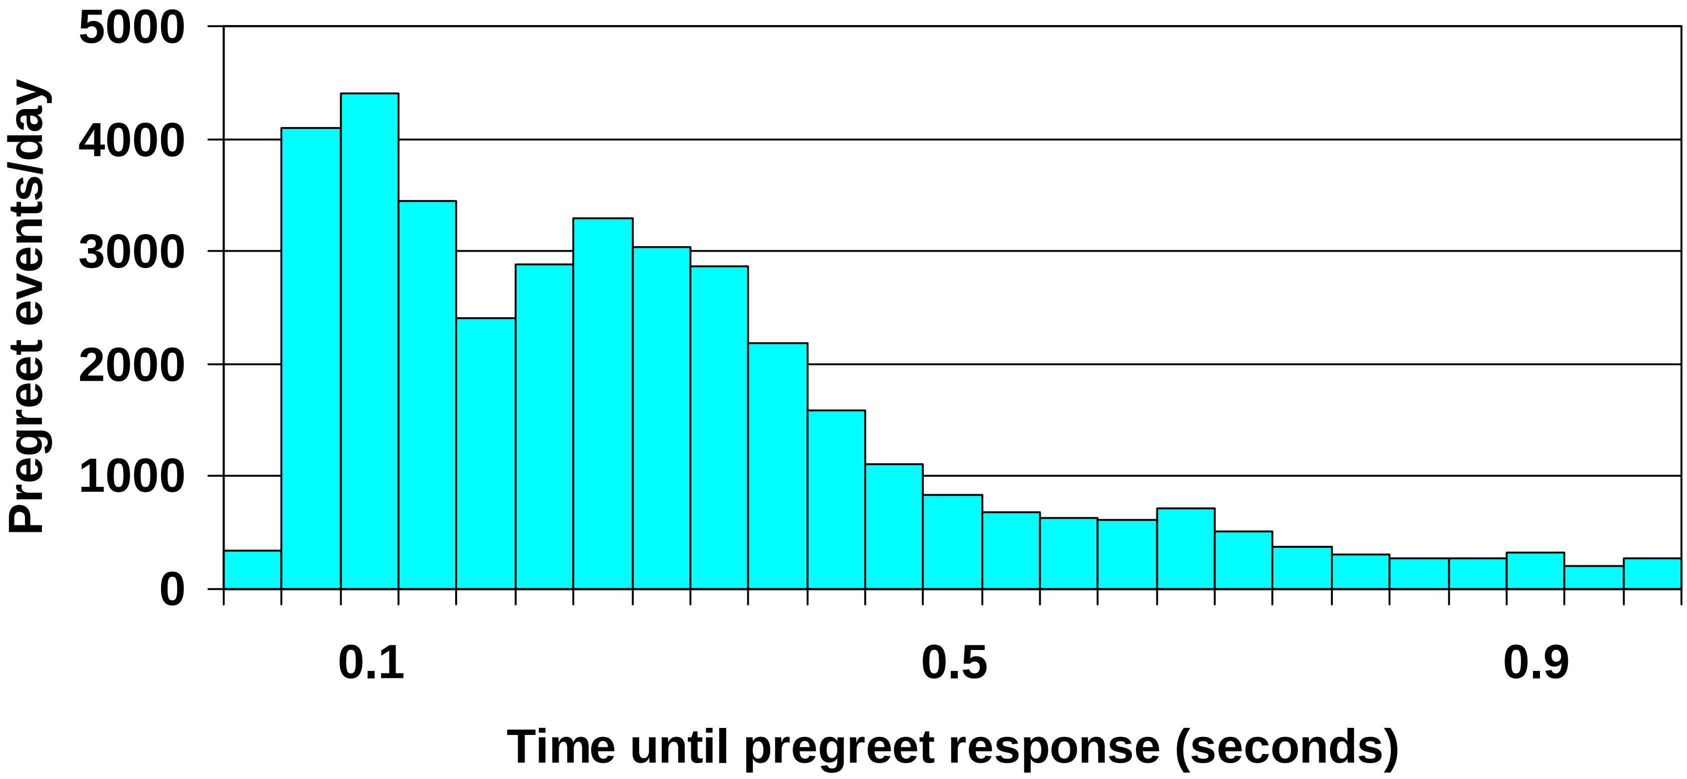

Over 60% of bots pregreet at mail.charite.de

8% Not on DNS blacklists. Berlin, Aug 26 – Sep 29, 2010

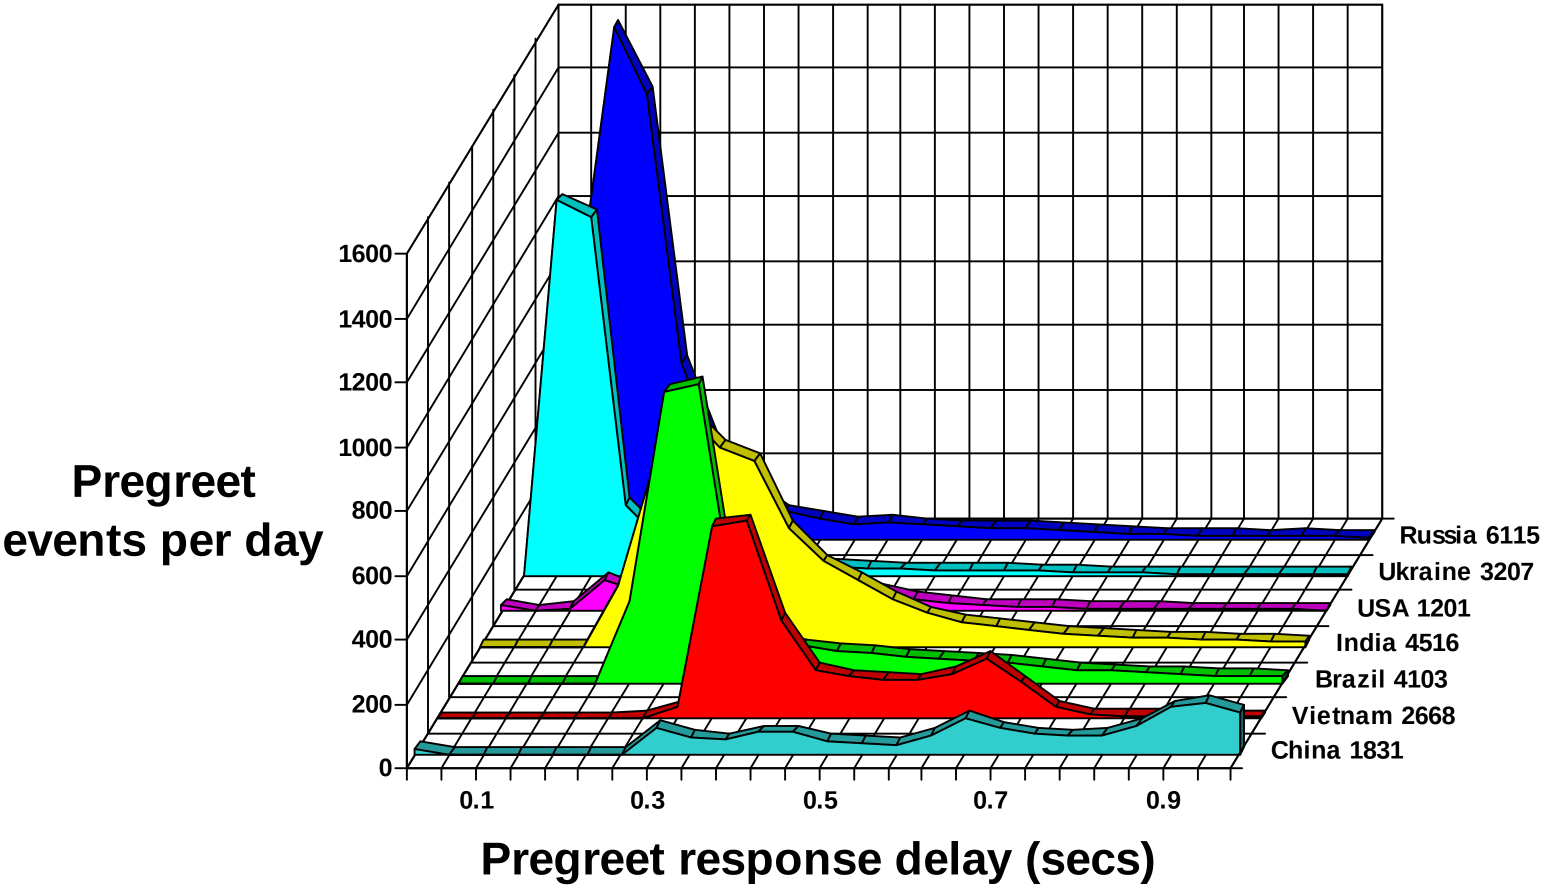

Over 70% of bots pregreet at mail.python.org
1% Not on DNS blacklists. Amsterdam, Sep 16 – 29, 2010

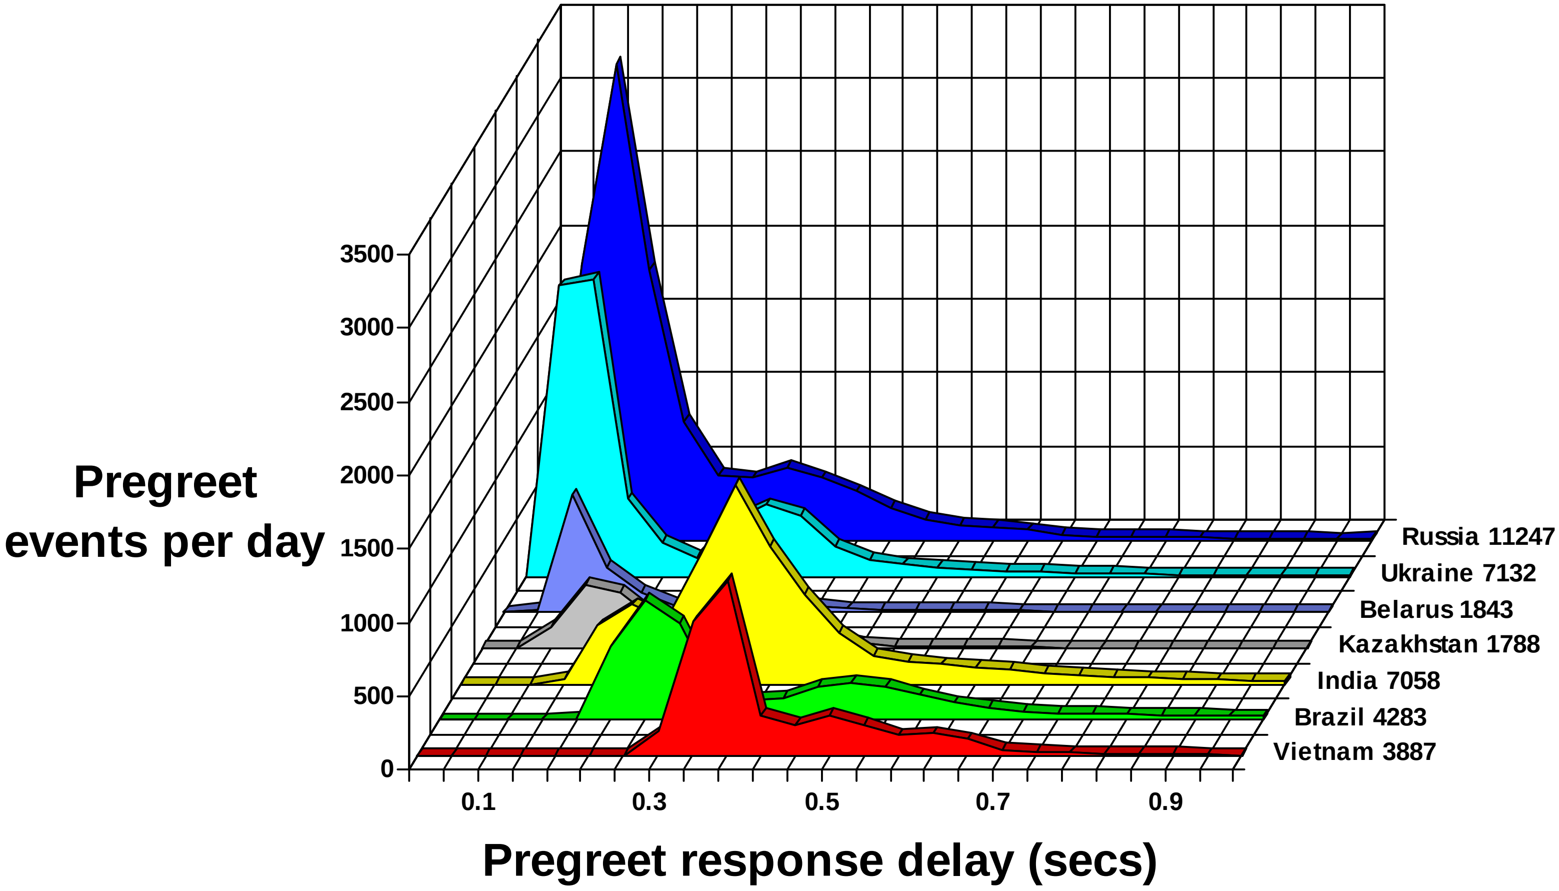

SPAM load varies by receiver and time of day

- SPAM load at different receivers:
  - A handful countries sends most of today's spam, but different receivers see different sender volumes.
- SPAM load at different times of day:
  - SPAM is a 24-hour operation, but spambots are not.
    - SPAM tends to be sent later in the day than HAM[1].

[1] S. Hao et al., Detecting Spammers with SNARE: Spatio-temporal Network-level Automatic Reputation Engine. Usenix Security 2009.

Spam connections/day at small European sites

Spam according to zen.spamhaus.org, Sep 3 – 23, 2010

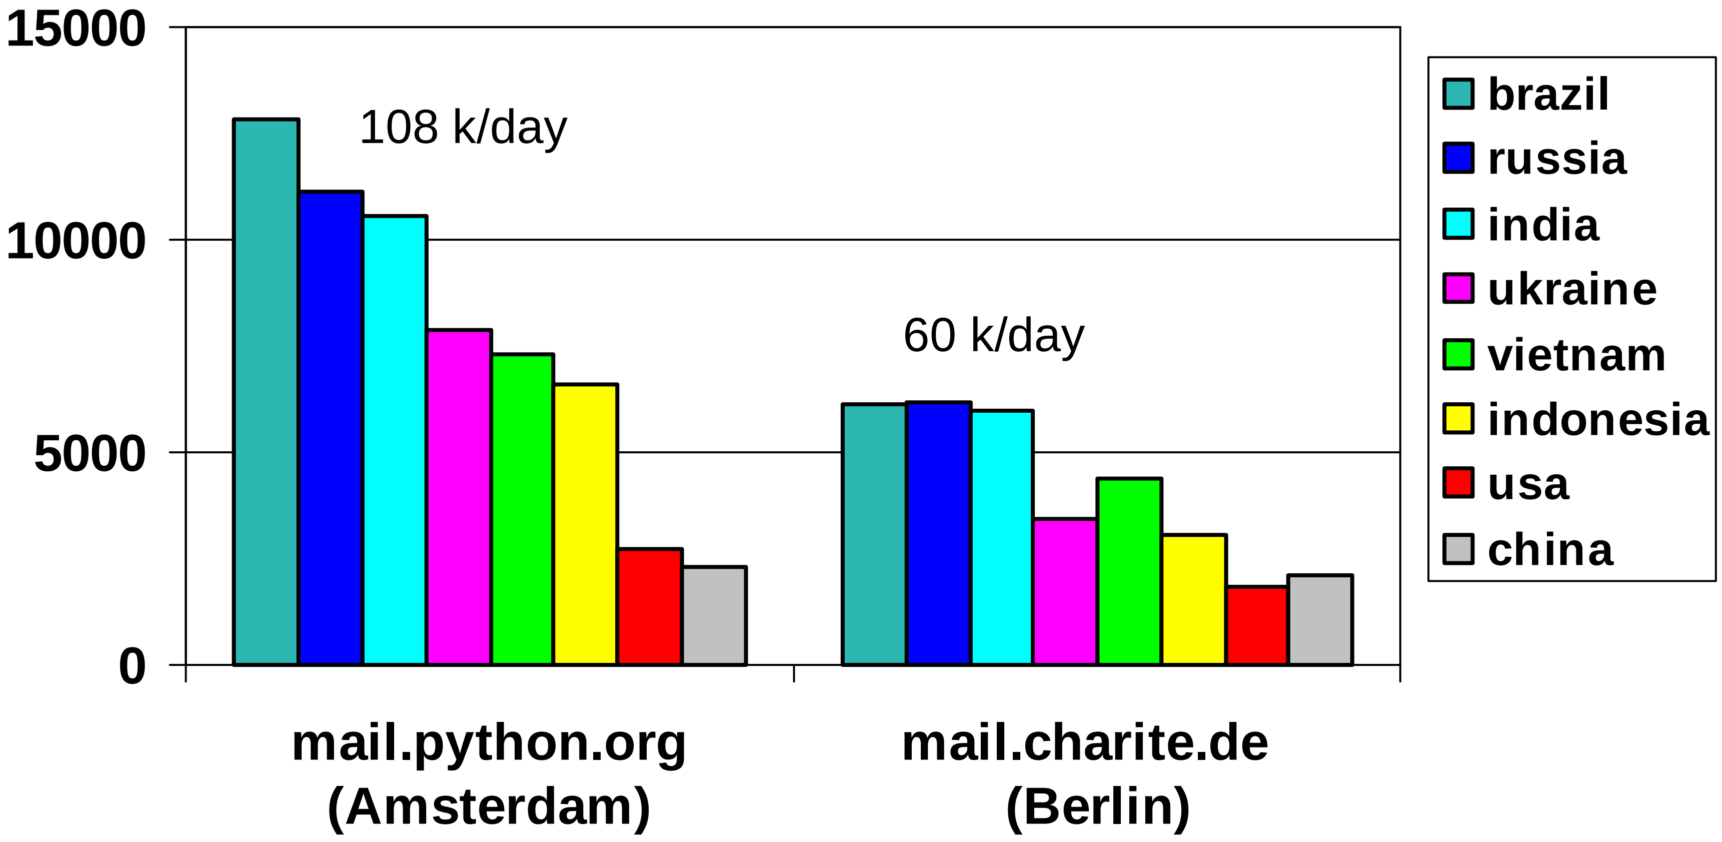

Changing threats

Spam volume by source country and hour at mail.charite.de UTC+2
Spam according to zen.spamhaus.org, Aug 26 – Sep 29, 2010

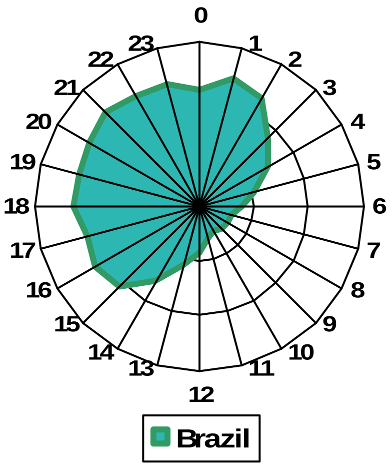

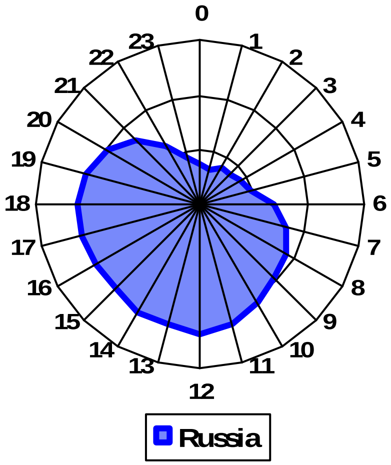

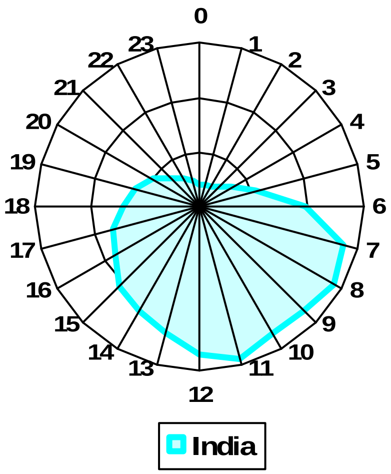

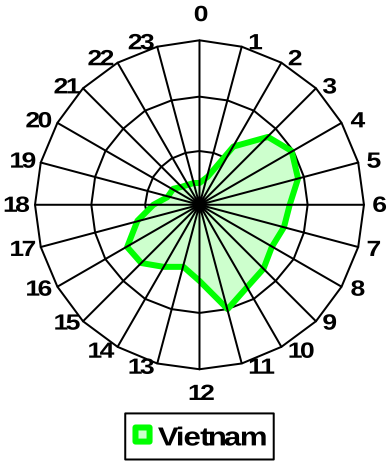

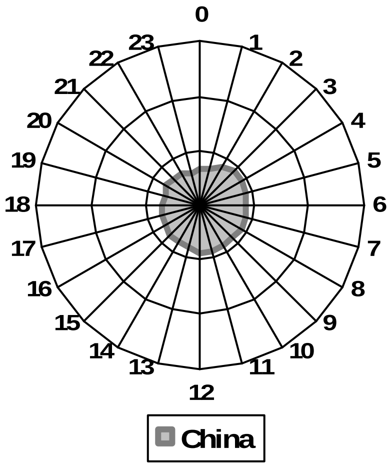

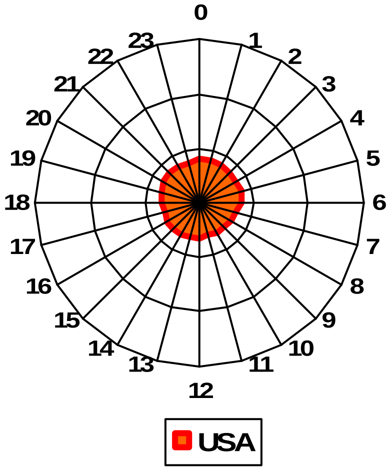

Changing threats

postscreen(8) results and status

- Parallel, weighted, DNS white/blacklist lookup.
- Static white/blacklist, dynamic “fast path” cache.
- Pilot results (small sites, up to 200k connections/day):
  - Pregreet (talk too early): up to ~10% not on DNS blacklist.
  - Pipelining (multiple commands): ~1% of spambots.
  - Hanging zombies (read timeout): ~1% of spambots.
- Other protocol tests to be added as botnets evolve.
- Start planning for extension interfaces.
- Expected release with Postfix 2.8, early 2011.

Changing threats