The State of the Email Address

Mike Afergan
Rob Beverly
January 27, 2005

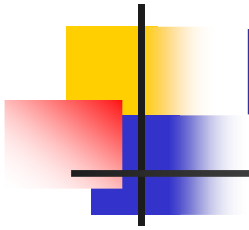

Emailtester - Outline

- Motivation/Goals/Background
- Methodology
- Results
- Questions

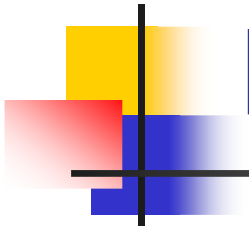

Motivation/Goals/Background

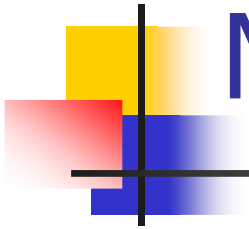

Motivation

- Electronic Mail is a widely-used, very important ($$) component of the Internet architecture
- But:
  - Simple Mail Transfer Protocol (SMTP) has been the standard protocol for over 20 years
  - In recent years, architecture has been strained by normal and unsolicited (i.e. spam) load
- General perception of Email: it "just works"
- Despite maturity and importance, surprisingly little data to substantiate this claim

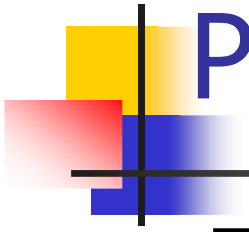

Project Goals

- To a large set of representative Internet SMTP servers, measure:
  - Loss
  - Latency
  - Paths
- *Without* administrative or user access to the servers
- Understand the results

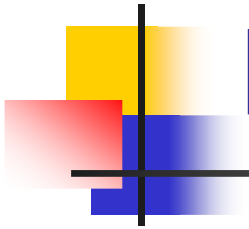

Methodology

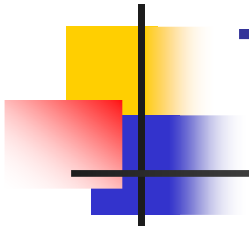

Testing Methodology

- Active measurement over several weeks
- Developed an email “traceroute” that relies on SMTP bounce-backs
- Traditionally servers inform users of errors: unknown user, undeliverable message, etc.

Methodology – High Level

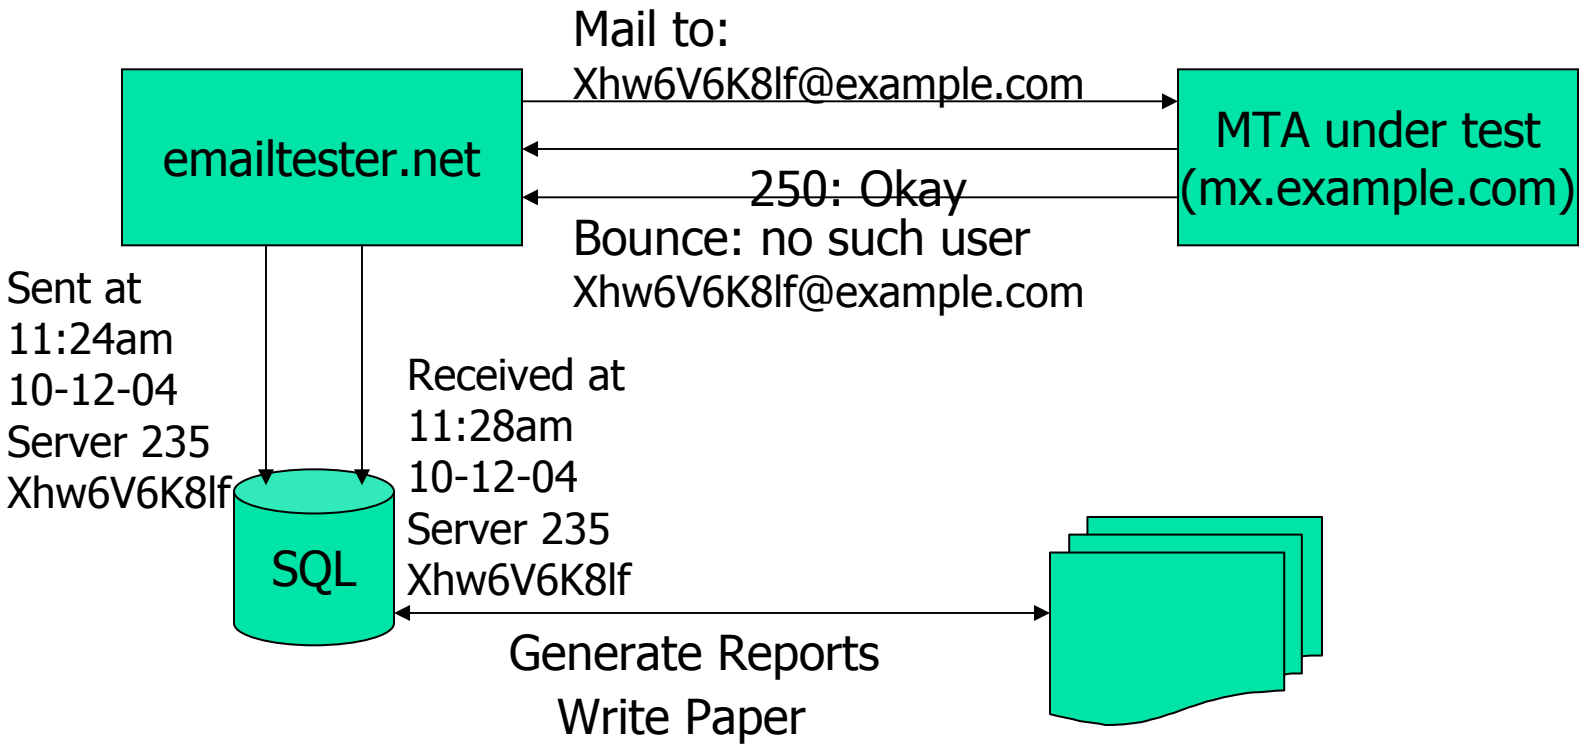

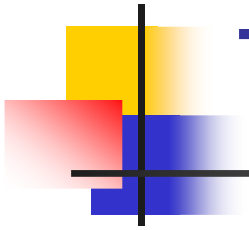

Testing Methodology

- Send emails addressed to unique, invalid recipients at each domain
- Record message ID (recipient), server ID and timestamp in database
- When and if the message bounces, disambiguate the message based on its message ID
- Record latency, statistics

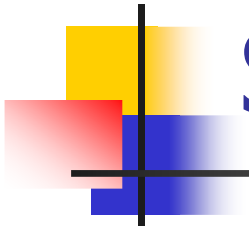

SMTP Bounce-backs

- Due to spam, only ~25% of the domains we survey respond with bounce-backs
- Despite the low return rate, our domain selection provides the most representative cross-section of SMTP servers possible
- How did we select domains?

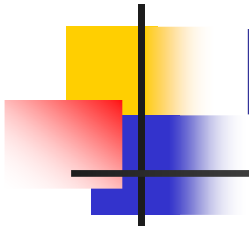

Domains

- Want large and diverse set of representative SMTP servers (large heterogeneity on Internet)
- Many Mail Transfer Agents (MTAs) in the wild:
  - qmail
  - exchange
  - postfix, etc…
- Each MTA may be uniquely configured
- Different servers may have vastly different:
  - Load (legitimate and spam)
  - Internet connectivity

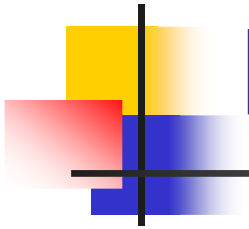

Domains

- Fortune 500:
  - Domains corresponding to Fortune 500 list
  - Likely more robust and fault-tolerant systems
- Topbits:
  - Most popular servers from an ISP web cache
- Random:
  - Pick a random 32-bit number
  - Use BGP table to determine if IP address is routable
  - Perform inverse DNS lookup on IP address
  - Positive DNS responses are truncated to the TLD

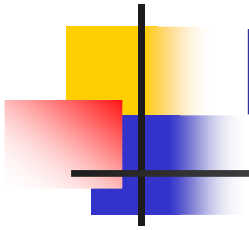

DNS and MX Records

- Need to remove non-determinism due to DNS caching and load-balancing
- Mail Exchanger (MX) records:
  - Map domain names to a set of mail servers
  - Each MX record has a preference value
  - If the most preferred server is unavailable, the remaining servers are tried in order of preference
  - More than one server may have the same preference value in order to load-balance

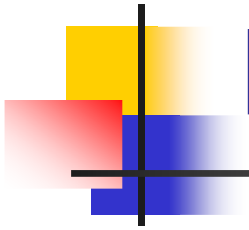

DNS and A Records

- Address (A) records:
  - Each server named in an MX record has one or more A records
  - May be a single IP or multiple addresses, again for load balancing or multi-homed hosts

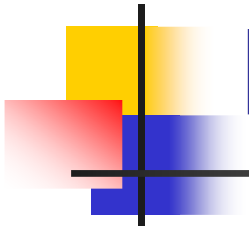

Pre-processing Step

- To remove non-determinism, each domain is resolved into the full set of MX servers
- Each MX record is further resolved into the set of corresponding IP addresses
- The atomic unit of testing is the IP address of a server supporting a domain
- Categorize all servers into:
  - Primary: most preferred
  - Secondary: all others

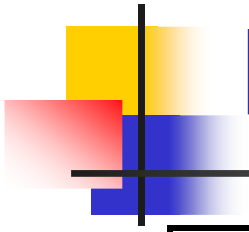

Domains

Category	Domains	Primary Servers	Total Servers
Fortune 500	282	486	735
Random	216	309	436
TopBits	73	212	297

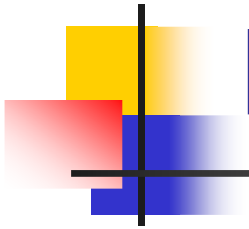

Methodology – More Details

- Important: Send emails to servers (not domains)!
- Randomize server order before each run
- Email body the same for every message:
  - Designed to be innocuous and pass through any spam filters
  - Provides information on the study and an opt-out link
  - Six domains opted out over the course of the study
  - Implies that some administrators monitor sources of invalid email closely!

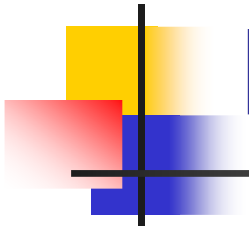

Results

- Overview
- Latency
- Loss
- Errors (not today)

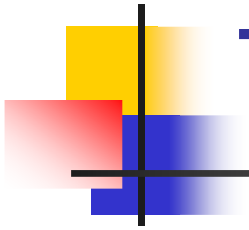

Testing

- Send to every server every 15 minutes
- 2880 15 minute "rounds" in Sept.
- Expected: uninteresting results (no loss, low latency)
- Unexpectedly we found:
  - Non-trivial loss rates
  - Bursty loss
  - Latencies longer than days
  - Non-deterministic server behavior

Most of our analysis and conversations were designed to explain away these strange observations. We failed.

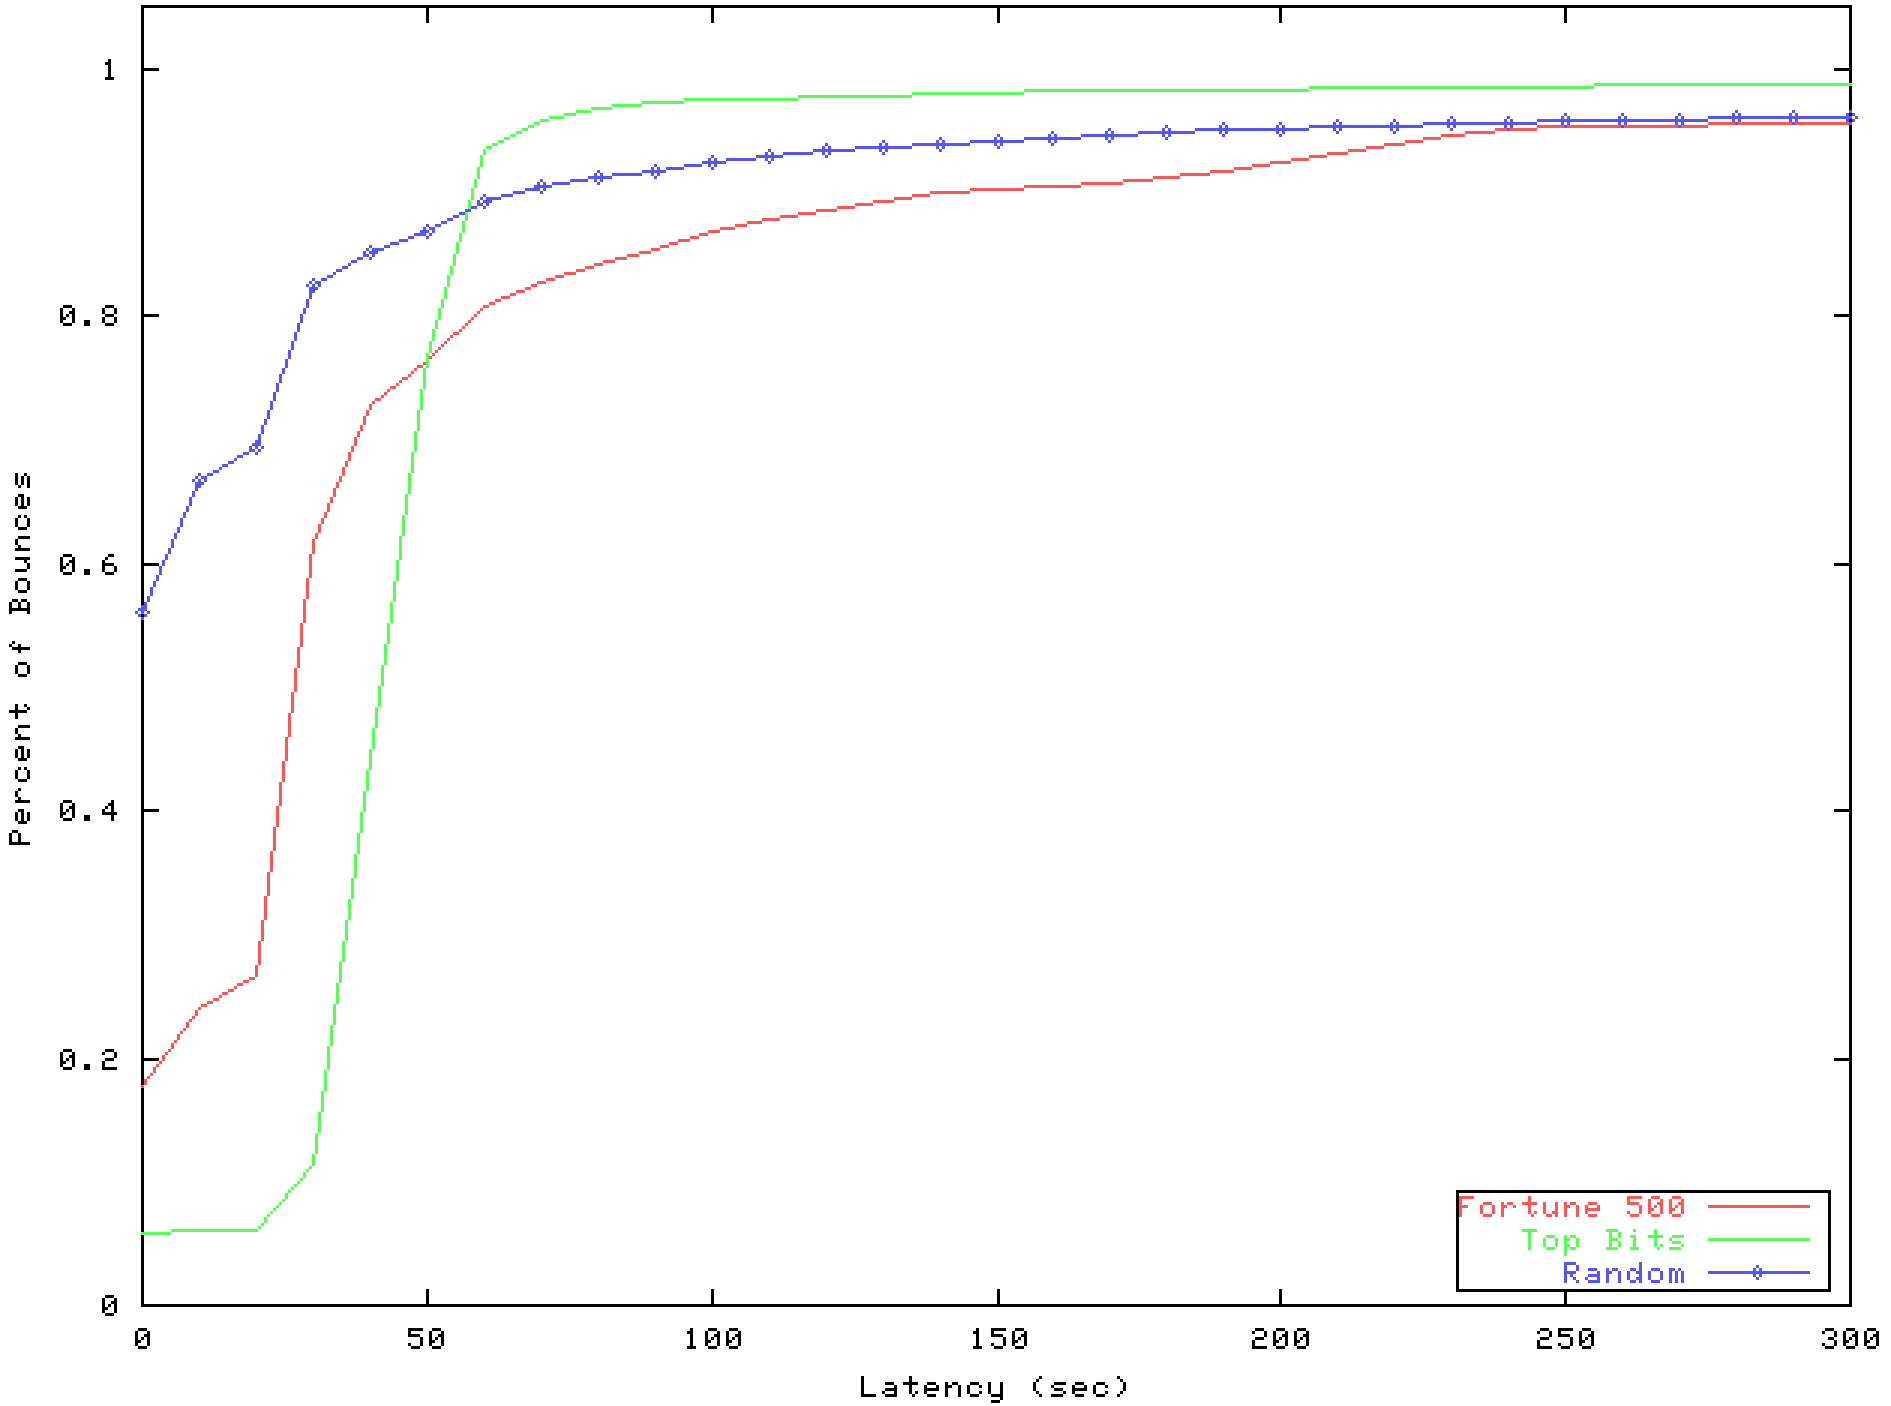
Percent of Bounces
Latency (sec)
0
0.2
0.4
0.6
0.8
1
0
50
100
150
200
250
300
Fortune 500
Top Bits
Random

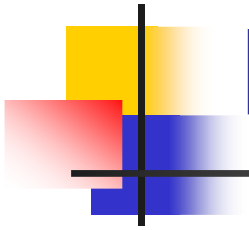

Pathological Latency Data

- 295 (0.035%) of bounces arrived more than 24 hours later
- One bounce came 30 days later!
- Examine the latency via the headers
  - No smoking guns
  - Clear evidence of delay within corporate and ISP email infrastructure

Loss Summary (Overall)

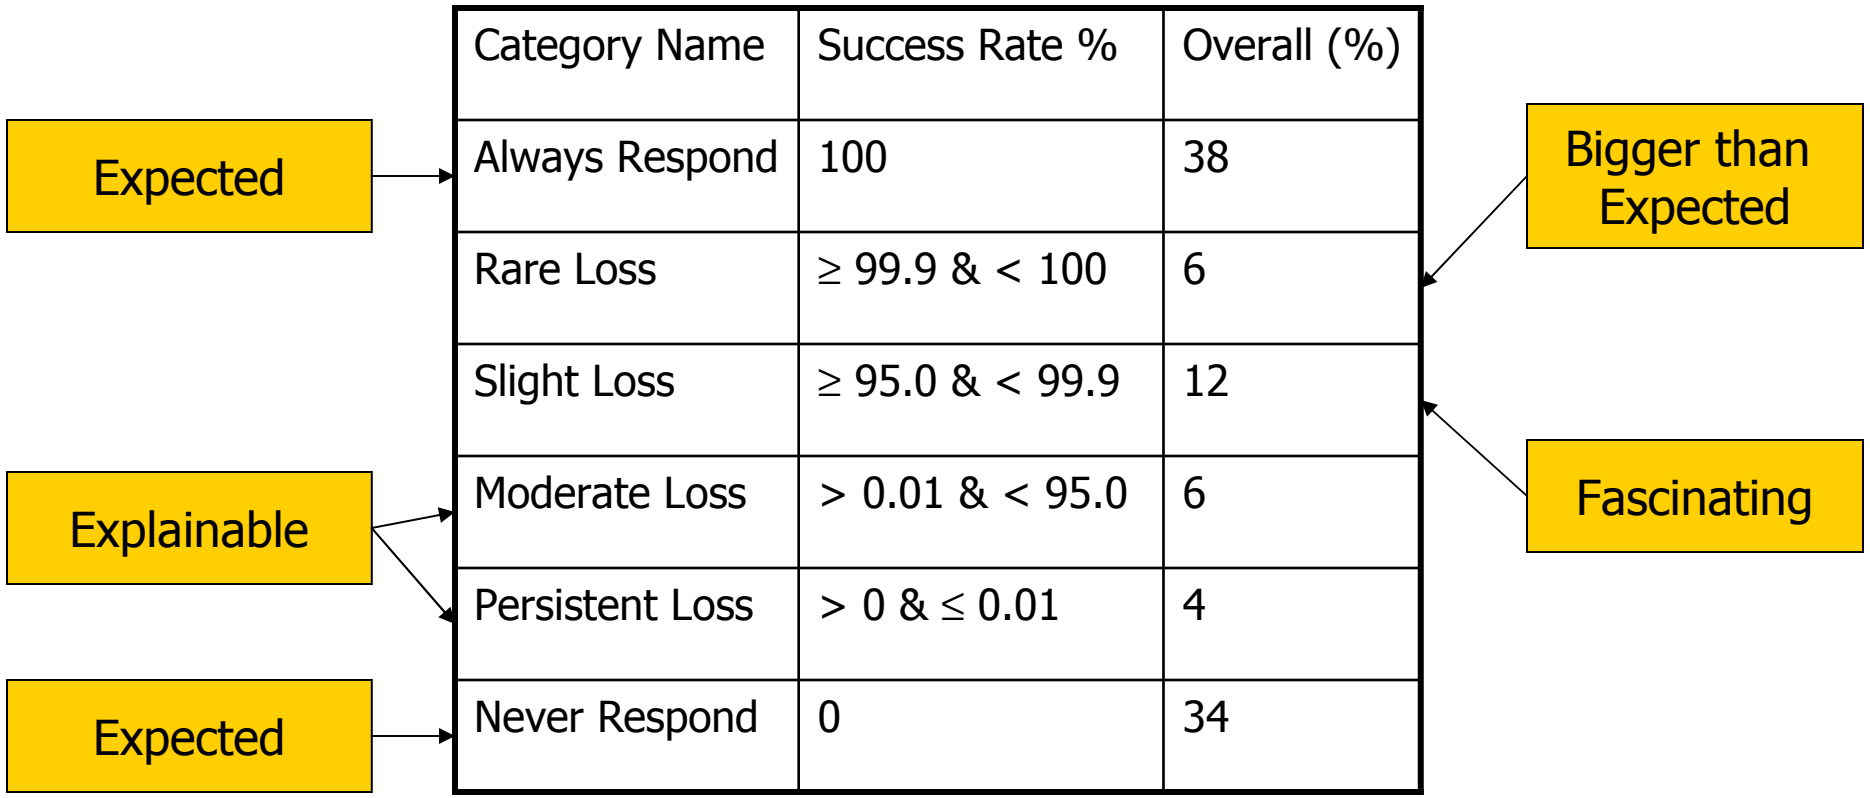

Category Name	Success Rate %	Overall (%)
Always Respond	100	38
Rare Loss	≥ 99.9 & < 100	6
Slight Loss	≥ 95.0 & < 99.9	12
Moderate Loss	> 0.01 & < 95.0	6
Persistent Loss	> 0 & ≤ 0.01	4
Never Respond	0	34

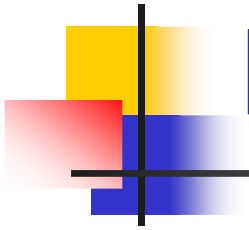

Possible Explanations

- Explanation #1: Bounces may not be representative of normal email behavior
  - Servers might implement different queues for bounces
  - Different policy under different load conditions
- Explanation #2: The IPs we see are virtual
- Can't find any evidence of this:
  - Loss not correlated with peak traffic hours
  - Greeting banners different for same IP address, but not correlated with loss

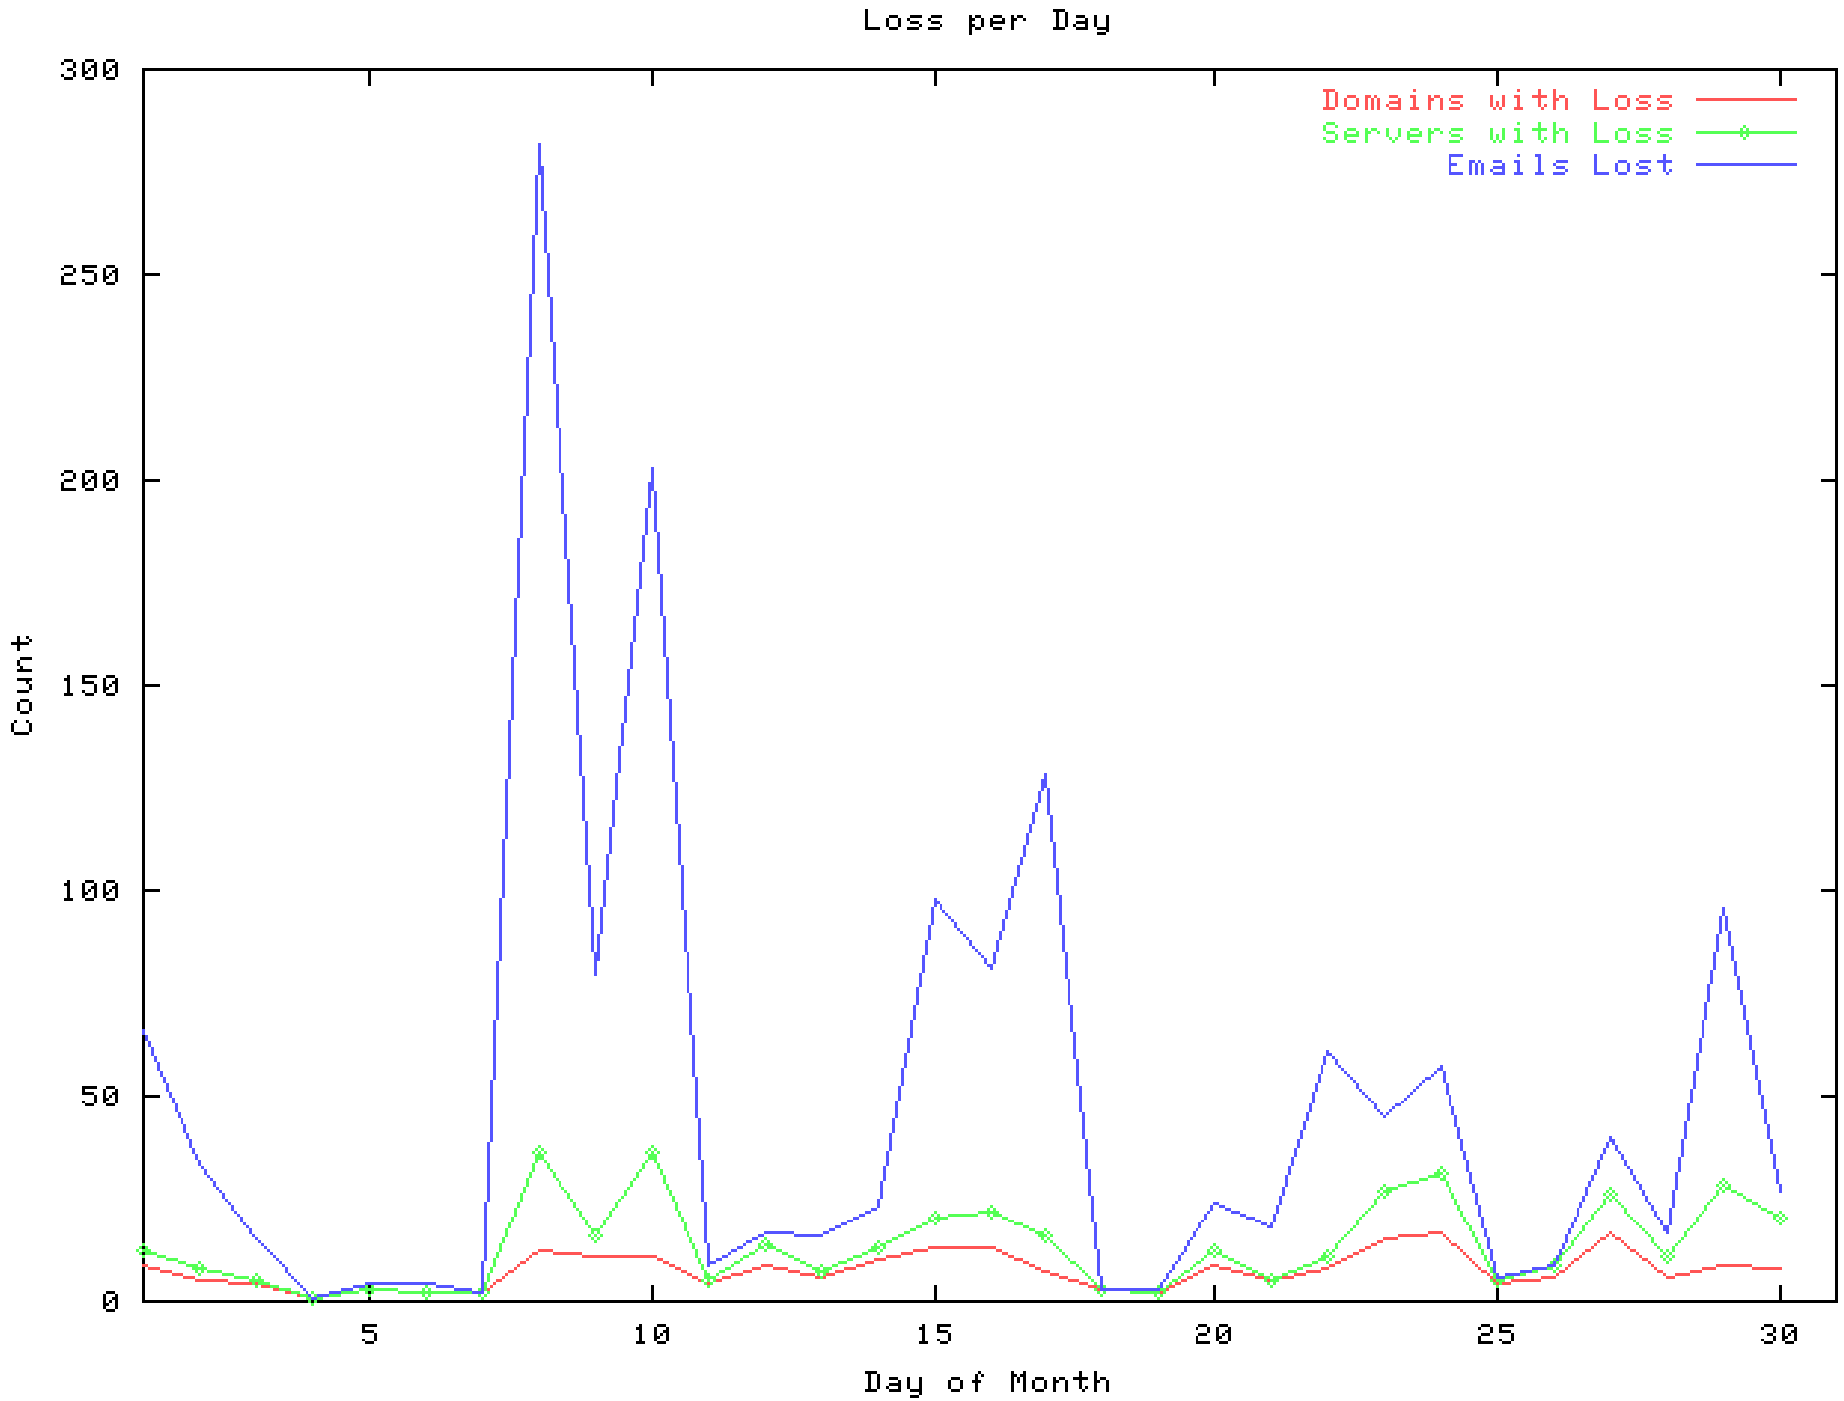
Loss per Day
Domains with Loss
Servers with Loss
Emails Lost
Count
Day of Month
0
50
100
150
200
250
300
5
10
15
20
25
30

Loss Summary

Success Rate %	Fortune 500 (%)	TopBits (%)	Random (%)	Overall (%)
100	36	16	53	38
≥ 99.9 & < 100	7	3	5	6
≥ 95.0 & < 99.9	13	26	5	12
> 0.01 & < 95.0	8	0	3	6
> 0 & ≤ 0.01	5	0	2	4
0	31	55	32	34

Observation: Corporate servers perform worse than random servers.

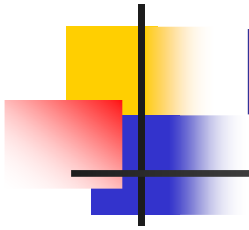

Loss Summary

- We observe much more loss than expected
- Most loss comes from bursts of loss, often from the same server and/or domain.
- Clearly demonstrate atypical and/or non-deterministic behavior

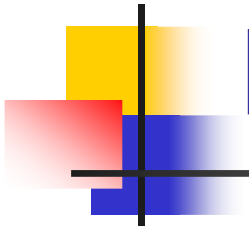

Errors

- Our system also records errors
- Errors before sending the email address are irrefutable.
- Interesting results
  - Many primary MXes are unreachable
  - Occasional odd error messages
  - Fortune 500 servers were more well-behaved

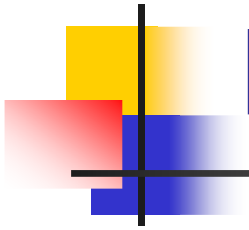

Potential Future Research

- Investigate non-deterministic behavior
- Analyze paths, path stability from SMTP headers
- Consider alternative testing techniques
- Measure loss on Hotmail, GMail, Yahoo, etc.
- Protocols for e2e reliable email/store-and-forward systems

Questions?

MX Record Example

```
;; ANSWER SECTION:
yahoo.com.      7200    IN      MX      1 mx2.mail.yahoo.com.
yahoo.com.      7200    IN      MX      1 mx3.mail.yahoo.com.
yahoo.com.      7200    IN      MX      5 mx4.mail.yahoo.com.
yahoo.com.      7200    IN      MX      1 mx1.mail.yahoo.com.

;; ADDITIONAL SECTION:
mx1.mail.yahoo.com.     1800    IN      A       64.157.4.78
mx1.mail.yahoo.com.     1800    IN      A       67.28.113.10
mx1.mail.yahoo.com.     1800    IN      A       67.28.113.11
mx2.mail.yahoo.com.     1800    IN      A       67.28.114.36
mx2.mail.yahoo.com.     1800    IN      A       64.156.215.8
mx2.mail.yahoo.com.     1800    IN      A       67.28.114.35
mx3.mail.yahoo.com.     1800    IN      A       64.156.215.5
mx3.mail.yahoo.com.     1800    IN      A       64.156.215.6
mx3.mail.yahoo.com.     1800    IN      A       64.156.215.7
mx3.mail.yahoo.com.     1800    IN      A       64.156.215.18
. . .
```

Methodology

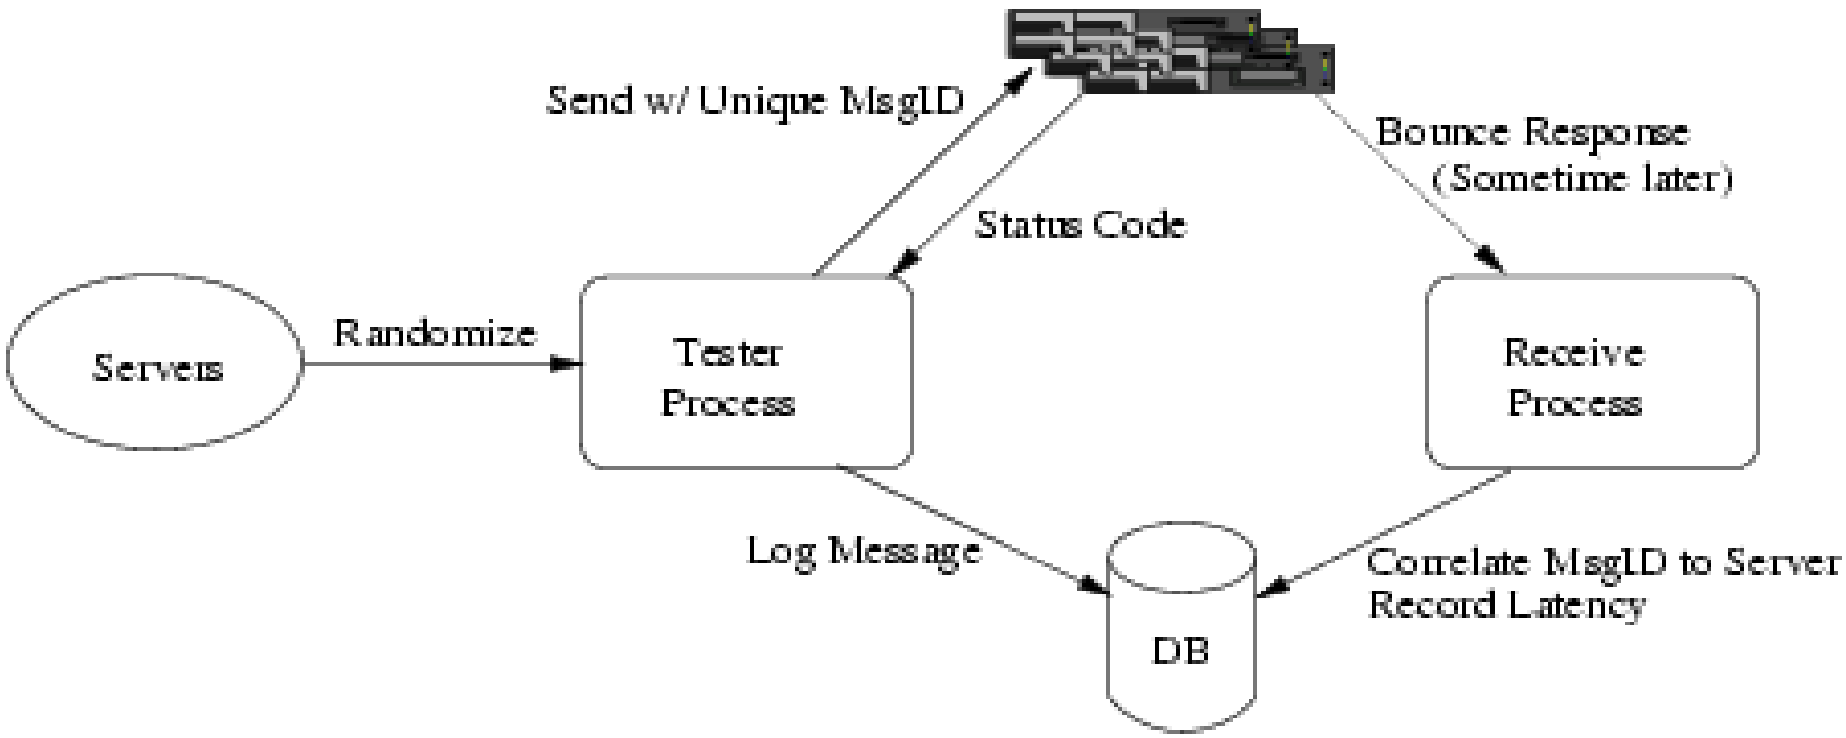

Loss

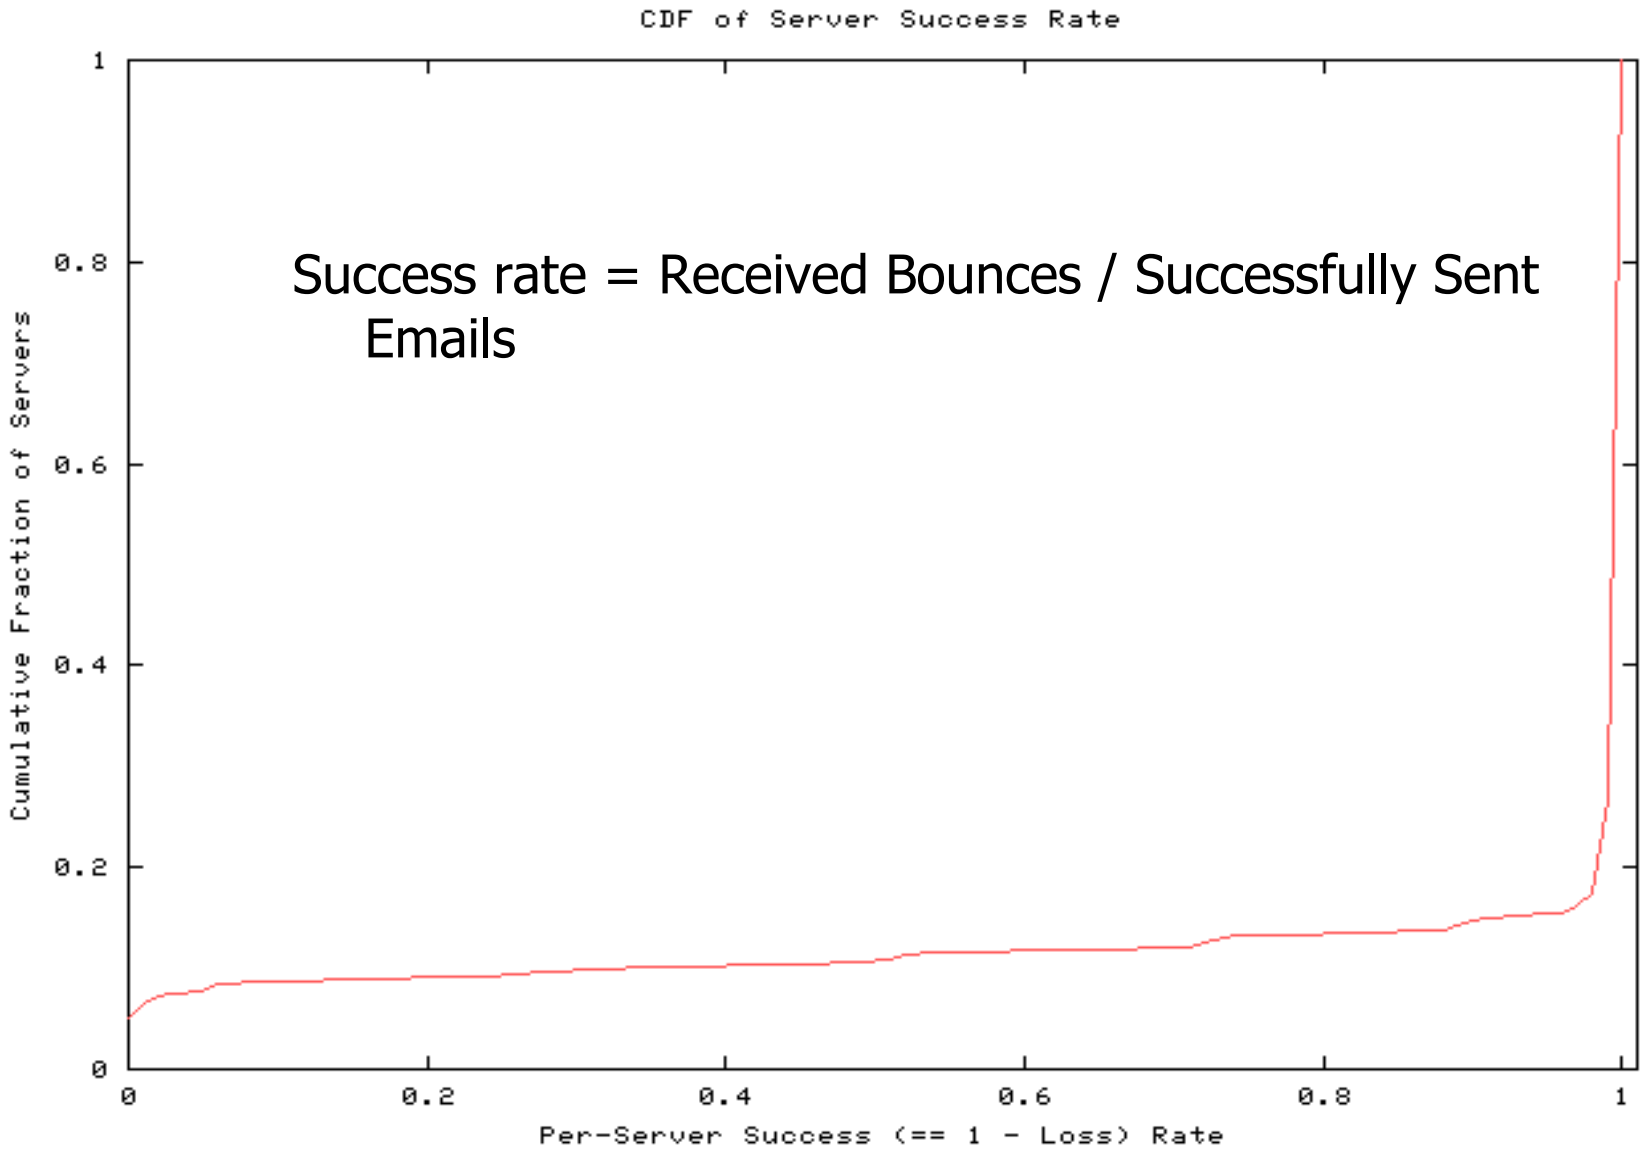

Success rate = Received Bounces / Successfully Sent Emails