

Understanding the Network-Level Behavior of Spammers

Anirudh Ramachandran Nick Feamster

College of Computing
Georgia Institute of Technology

Sept. 14, 2006



Why network-level properties?

Content-level filters do not always work



Why network-level properties?

Content-level filters do not always work

- Content-based filters: easy to evade and hard to maintain



Why network-level properties?

Content-level filters do not always work

- Content-based filters: easy to evade and hard to maintain
- Network-level properties of spam are usually less variable



Why network-level properties?

Content-level filters do not always work

- Content-based filters: easy to evade and hard to maintain
- Network-level properties of spam are usually less variable

This paper



Why network-level properties?

Content-level filters do not always work

- Content-based filters: easy to evade and hard to maintain
- Network-level properties of spam are usually less variable

This paper

- Characterizes various network-level properties of spam using a joint analysis of several datasets



Why network-level properties?

Content-level filters do not always work

- Content-based filters: easy to evade and hard to maintain
- Network-level properties of spam are usually less variable

This paper

- Characterizes various network-level properties of spam using a joint analysis of several datasets
- Identifies network-level properties which might be used to complement content-based filters



Major Findings

- *Where does spam come from?*



Major Findings

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space



Major Findings

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space

Implication: Insight about “spammier” prefixes could aid filter design



Major Findings

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space

Implication: Insight about “spammier” prefixes could aid filter design

- *Do spammers really hijack routes?*



Major Findings

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space

Implication: Insight about “spammier” prefixes could aid filter design

- *Do spammers really hijack routes?*

A: Yes, a small set of spammers *continually* do this



Major Findings

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space

Implication: Insight about “spammier” prefixes could aid filter design

- *Do spammers really hijack routes?*

A: Yes, a small set of spammers *continually* do this

I: Traceability of spam (or any malicious traffic) is not guaranteed



Major Findings

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space

Implication: Insight about “spammier” prefixes could aid filter design

- *Do spammers really hijack routes?*

A: Yes, a small set of spammers *continually* do this

I: Traceability of spam (or any malicious traffic) is not guaranteed

- *Who sends the most spam?*



Major Findings

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space

Implication: Insight about “spammier” prefixes could aid filter design

- *Do spammers really hijack routes?*

A: Yes, a small set of spammers *continually* do this

I: Traceability of spam (or any malicious traffic) is not guaranteed

- *Who sends the most spam?*

A: Most spam is sent by Windows hosts



Major Findings

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space

Implication: Insight about “spammier” prefixes could aid filter design

- *Do spammers really hijack routes?*

A: Yes, a small set of spammers *continually* do this

I: Traceability of spam (or any malicious traffic) is not guaranteed


- *Who sends the most spam?*

A: Most spam is sent by Windows hosts

I: Spam might be mostly from bots — bot-membership identification and in-network filters could help




Traces used in this study



- 1 Spam sinkhole; no real email addresses; spoofs **all** email received




Traces used in this study



- 1 Spam sinkhole; no real email addresses; spools **all** email received
 - Running *MailAvenger*, a highly configurable SMTP server




Traces used in this study



- 1 Spam sinkhole; no real email addresses; spoofs **all** email received
 - Running *MailAvenger*, a highly configurable SMTP server
- 2 BGP Monitor *in the same AS* as the sinkhole, with an iBGP session with the border router




Traces used in this study



- 1 Spam sinkhole; no real email addresses; spoofs **all** email received
 - Running *MailAvenger*, a highly configurable SMTP server
- 2 BGP Monitor *in the same AS* as the sinkhole, with an iBGP session with the border router
- 3 Traffic logs at a hijacked “command and control” (C&C) for a botnet (*Bobax*) used purely for spam



Traces used in this study



- 1 Spam sinkhole; no real email addresses; spoofs **all** email received
 - Running *MailAvenger*, a highly configurable SMTP server
- 2 BGP Monitor *in the same AS* as the sinkhole, with an iBGP session with the border router
- 3 Traffic logs at a hijacked “command and control” (C&C) for a botnet (*Bobax*) used purely for spam
- 4 Logs for accepted and rejected email (using a variety of spam filters) from a large email provider



Outline

- *Where does spam come from?* ←←←

Answer: Most spam is received from very few regions of IP space

Implication: Insight about “spammier” prefixes

- *Do spammers really hijack routes?*

A: Yes, A small set of spammers *continually* do this

I: Traceability of spam (or any malicious traffic) cannot be guaranteed

- *Who sends the most spam?*

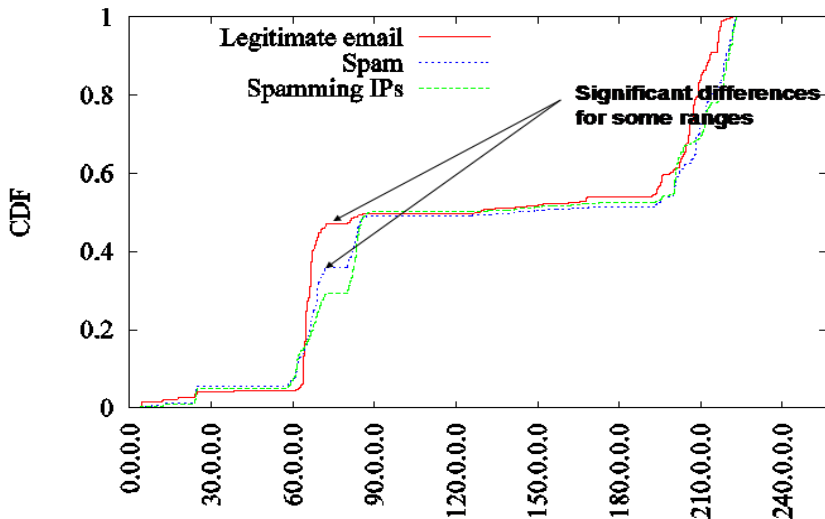
A: Most spam is received from Windows hosts

I: Windows hosts are likely to be bots — bot-membership identification and in-network filters might help



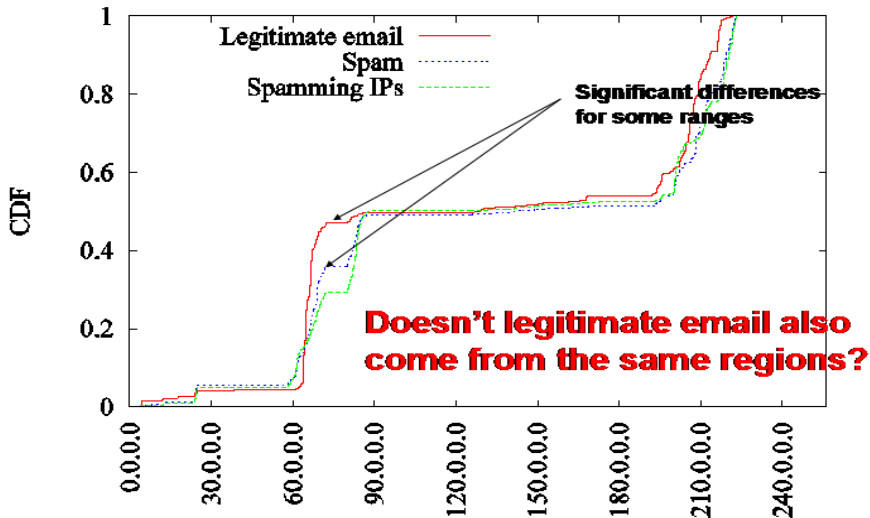
Where in IP space does spam come from?

Spam comes from few concentrated regions of IP space



Where in IP space does spam come from?

Spam comes from few concentrated regions of IP space



Distribution across Autonomous Systems



Distribution across Autonomous Systems

<i>AS Name</i>	<i>Country</i>
Korean IX	Korea
China Telecom	China
Sprint	US
China Net. Comm.	China
Hanaro Telecom	Japan
JKS Media, LLC	US
Polish Telecom	Poland
AT&T WorldNet	US
Verizon Global	US
Microsoft	US

Top 10 ASes by spam email count



Distribution across Autonomous Systems

<i>AS Name</i>	<i>Country</i>
Korean IX	Korea
China Telecom	China
Sprint	US
China Net. Comm.	China
Hanaro Telecom	Japan
JKS Media, LLC	US
Polish Telecom	Poland
AT&T WorldNet	US
Verizon Global	US
Microsoft	US

Top 10 ASes by spam email count

<i>AS Name</i>	<i>Country</i>
Google Inc.	US
AT&T WorldNet	US
Yahoo	US
Savvis	US
Earthlink, Inc	US
Schlund Partner AG	Germany
Microsoft Corp	US
Inktomi Corporation	US
GTE.net LLC	US
Inktomi Corporation	US

Top 10 ASes by legit email count



Distribution across Autonomous Systems

<i>AS Name</i>	<i>Country</i>
Korean IX	Korea
China Telecom	China
Sprint	US
China Net. Comm.	China
Hanaro Telecom	Japan
JKS Media, LLC	US
Polish Telecom	Poland
AT&T WorldNet	US
Verizon Global	US
Microsoft	US

<i>AS Name</i>	<i>Country</i>
Google Inc.	US
AT&T WorldNet	US
Yahoo	US
Savvis	US
Earthlink, Inc	US
Schlund Partner AG	Germany
Microsoft Corp	US
Inktomi Corporation	US
GTE.net LLC	US
Inktomi Corporation	US

Top 10 ASes by spam email count

Top 10 ASes by legit email count

Points to note

Distribution across Autonomous Systems

<i>AS Name</i>	<i>Country</i>
Korean IX	Korea
China Telecom	China
Sprint	US
China Net. Comm.	China
Hanaro Telecom	Japan
KS Media, LLC	US
Polish Telecom	Poland
AT&T WorldNet	US
Verizon Global	US
Microsoft	US

<i>AS Name</i>	<i>Country</i>
Google Inc.	US
AT&T WorldNet	US
Yahoo	US
Savvis	US
Earthlink, Inc	US
Schlund Partner AG	Germany
Microsoft Corp	US
Inktomi Corporation	US
GTE.net LLC	US
Inktomi Corporation	US

Top 10 ASes by spam email count

Top 10 ASes by legit email count

Points to note

- The two top spamming ASes make up over 10% of received spam

Distribution across Autonomous Systems

<i>AS Name</i>	<i>Country</i>
Korean IX	Korea
China Telecom	China
Sprint	US
China Net. Comm.	China
Hanaro Telecom	Japan
JKS Media, LLC	US
Polish Telecom	Poland
AT&T WorldNet	US
Verizon Global	US
Microsoft	US

<i>AS Name</i>	<i>Country</i>
Google Inc.	US
AT&T WorldNet	US
Yahoo	US
Savvis	US
Earthlink, Inc	US
Schlund Partner AG	Germany
Microsoft Corp	US
Inktomi Corporation	US
GTE.net LLC	US
Inktomi Corporation	US

Top 10 ASes by spam email count

Top 10 ASes by legit email count

Points to note

- The two top spamming ASes make up over 10% of received spam
- ASes based in the US, however, account for most spam overall

Distribution across Autonomous Systems

<i>AS Name</i>	<i>Country</i>
Korean IX	Korea
China Telecom	China
Sprint	US
China Net. Comm.	China
Hanaro Telecom	Japan
JKS Media, LLC	US
Polish Telecom	Poland
AT&T WorldNet	US
Verizon Global	US
Microsoft	US

<i>AS Name</i>	<i>Country</i>
Google Inc.	US
AT&T WorldNet	US
Yahoo	US
Savvis	US
Earthlink, Inc	US
Schlund Partner AG	Germany
Microsoft Corp	US
Inktomi Corporation	US
GTE.net LLC	US
Inktomi Corporation	US

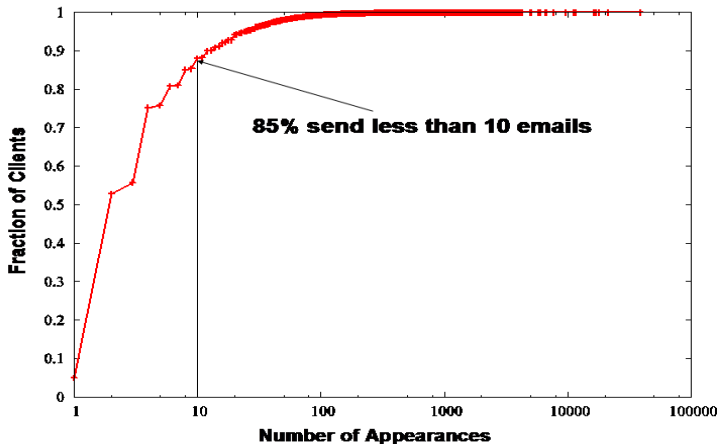
Top 10 ASes by spam email count

Top 10 ASes by legit email count

Points to note

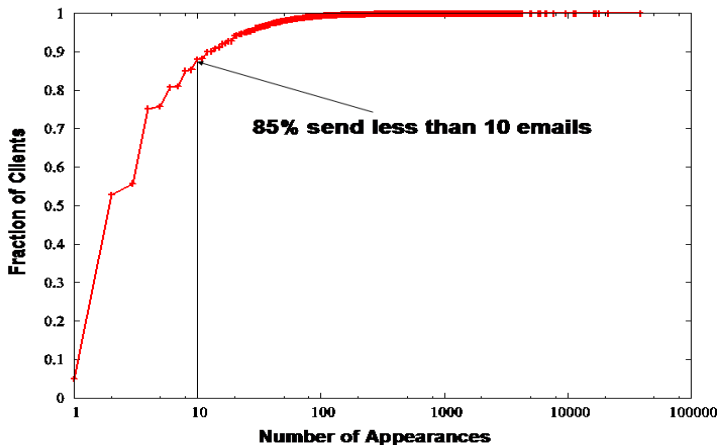
- The two top spamming ASes make up over 10% of received spam
- ASes based in the US, however, account for most spam overall
- Top ASes by legitimate email count substantially different

How prolific is each spamming host?



Network-level Characteristics: Hypotheses

How prolific is each spamming host?



Network-level Characteristics: Hypotheses

- Spam volume by IP **ranges** far more persistent

Outline

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space

Implication: Insight about “spammier” prefixes

- *Do spammers really hijack routes?* ←←

Answer: Yes, a small set of spammers *continually* do this

Implication: Traceability of spam (or any malicious traffic) cannot be guaranteed

- *Who sends the most spam?*

Answer: Most spam is received from Windows hosts

Implication: Windows hosts are likely to be bots — bot-membership identification and in-network filters could help



Spam from hijacked routes

Basic Idea:



Spam from hijacked routes

Basic Idea: 1) Gain access to a router,



Spam from hijacked routes

Basic Idea: 1) Gain access to a router, 2) Announce a “hijacked” prefix,



Spam from hijacked routes

Basic Idea: 1) Gain access to a router, 2) Announce a “hijacked” prefix,
3) Send some spam,




Spam from hijacked routes

Basic Idea: 1) Gain access to a router, 2) Announce a “hijacked” prefix,
3) Send some spam, 4) Withdraw the route



Spam from hijacked routes

- Basic Idea:* 1) Gain access to a router, 2) Announce a “hijacked” prefix, 3) Send some spam, 4) Withdraw the route



Spam from 61/8 for a 10 minute period



Spam from hijacked routes: “Spectrum Agility”

- Why such large address ranges?
 - Short prefixes (larger number of addresses) are *less likely to be filtered*
 - Does *not* disrupt legitimate, more specific routes
 - Allows the spammer to *hop between a large “spectrum” of IP addresses*
- How much spam from short-lived BGP routes?
 - Less than 10% of spam (for routes announced for less than 1 day)
 - A small *group of persistent ASes* seem to be using this technique to send spam (*i.e.*, intentionally)



Outline

- *Where does spam come from?*

Answer: Most spam is received from very few regions of IP space

Implication: Certain prefixes are “spammier” than others

- *Do spammers really hijack routes?*

Answer: A small set of spammers *continually* use short-lived route announcements to remain untraceable

Implication: Traceability of spam (or any malicious traffic) is no longer guaranteed

- *Who sends the most spam?* ←←

Answer: Most spam is received from Windows hosts, each of which send a small volume of spam *to our domain*

Implication: Windows hosts are likely to be bots, and volume-based or single-domain blacklists are unlikely to be effective



Characteristics of spamming bots

● **Distribution across IP space for bots**

- Similar to IP space distribution for all spam
- Lower bot activity in ranges where spam also comes from hijacked routes

● **Operating Systems of Spamming Hosts**

Spam:

- Windows-based: 58%
- Unix-based: 8%
- Unidentified/No Fingerprint: 34%

Clients:

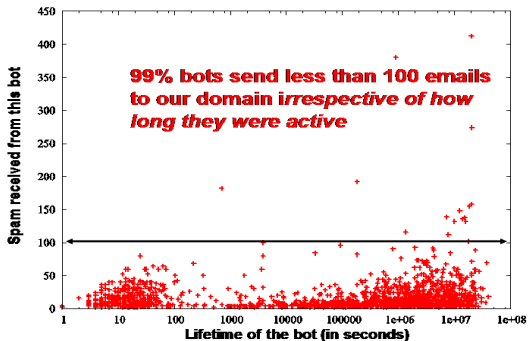
- Windows-based: 70%
- Unix-based: 4%
- Unidentified/No Fingerprint: 27%



How much spam does each bot send?




How much spam does each bot send?



Amount of spam from a bot vs. its lifetime



How much spam does each bot send?



Amount of spam from a bot vs. its lifetime

- 99% of bots send less than 100 pieces of spam *to our domain*
- 65% of known bots send spam only *once*; 75% for less than two minutes \Rightarrow collaborative blacklists are the way to go
- “Single-shot” bots send only 25% of all bot spam, but could become heavy spammers over time



Where do DNS Blacklists stand?



Where do DNS Blacklists stand?

- 80% of spamming hosts are listed in *at least one* blacklist
 - *Catch*: Requires checking **8** blacklists



Where do DNS Blacklists stand?

- 80% of spamming hosts are listed in *at least one* blacklist
 - *Catch*: Requires checking **8** blacklists
 - Even the most aggressive blacklist has only **40%** IPs listed



Where do DNS Blacklists stand?

- 80% of spamming hosts are listed in *at least one* blacklist
 - *Catch*: Requires checking **8** blacklists
 - Even the most aggressive blacklist has only **40%** IPs listed
- A large fraction (**50%**) of spam hosts using short-lived routes are not listed in *any* blacklist



Where do DNS Blacklists stand?

- 80% of spamming hosts are listed in *at least one* blacklist
 - *Catch*: Requires checking **8** blacklists
 - Even the most aggressive blacklist has only **40%** IPs listed
- A large fraction (**50%**) of spam hosts using short-lived routes are not listed in *any* blacklist
- However, a much *higher* fraction of Bobax bots are blacklisted



Where do DNS Blacklists stand?

- 80% of spamming hosts are listed in *at least one* blacklist
 - *Catch*: Requires checking **8** blacklists
 - Even the most aggressive blacklist has only **40%** IPs listed
- A large fraction (**50%**) of spam hosts using short-lived routes are not listed in *any* blacklist
- However, a much *higher* fraction of Bobax bots are blacklisted
 - Bots send email to *many* domains



Takeaway lessons

- *Network-level spam filtering*
 - Network-level properties are less **malleable**, and are observable **closer to the source** of spam
 - **Aggregate properties** (e.g., IP prefix, ASN, route used etc.) are more suitable to detect spam than, say, an IP address
 - Some **network-level properties can be easily incorporated into spam-filters**; could be used as a first-pass filter.

- *Redefining end-host identifiers*
 - Spam filtering requires a **better notion of end-host identity**
 - **Securing the Internet routing infrastructure** is key to traceability



Summary

- Joint analysis of many data sets to characterize the network-level behavior of Spammers



Summary

- Joint analysis of many data sets to characterize the network-level behavior of Spammers
- Network-level properties offer insight into new techniques of spamming



Summary

- Joint analysis of many data sets to characterize the network-level behavior of Spammers
- Network-level properties offer insight into new techniques of spamming
- Filters based on network-level properties could be incorporated easily into spam filters and could flag spam where content-based methods fail



Summary

- Joint analysis of many data sets to characterize the network-level behavior of Spammers
- Network-level properties offer insight into new techniques of spamming
- Filters based on network-level properties could be incorporated easily into spam filters and could flag spam where content-based methods fail

- **Current and Future Work**



Summary

- Joint analysis of many data sets to characterize the network-level behavior of Spammers
- Network-level properties offer insight into new techniques of spamming
- Filters based on network-level properties could be incorporated easily into spam filters and could flag spam where content-based methods fail
- **Current and Future Work**
 - Incorporation of network-level properties of spam into spam filters/MTAs



Summary

- Joint analysis of many data sets to characterize the network-level behavior of Spammers
- Network-level properties offer insight into new techniques of spamming
- Filters based on network-level properties could be incorporated easily into spam filters and could flag spam where content-based methods fail
- **Current and Future Work**
 - Incorporation of network-level properties of spam into spam filters/MTAs
 - Implementation of a massively collaborative system for filtering spam closer to its source




Questions?




IP distribution for Bobax drones




IP distribution for Bobax drones



IP distribution for Bobax drones




IP distribution for Bobax drones



- Bobax drones have a similar distribution across IP space as all spamming hosts \Rightarrow much of received spam may be due to botnets like Bobax



IP distribution for Bobax drones



- Bobax drones have a similar distribution across IP space as all spamming hosts \Rightarrow much of received spam may be due to botnets like Bobax
- 60.*–67.* has significant spamming hosts, but less spam from Bobax bots — indication that other techniques might be in use (BGP Spectrum Agility)?



Operating Systems of Spamming Hosts



Operating Systems of Spamming Hosts

<i>OS</i>	<i>Clients</i>	<i>Total Spam</i>
Windows	854404 (70%)	5863112 (58%)
Linux	28132 (2.3%)	557377 (5.5%)
FreeBSD	6584 (0.5%)	152456 (1.5%)
MacOS	2944 (0.2%)	46151 (0.4%)
Solaris	1275 (< 0.1%)	18084 (0.2%)
Unidentified	128580 (10.4%)	1212722 (12%)
No Fingerprint	204802 (16.7%)	2225410 (22%)
Total	1228403	10103837

The operating system of each unique sender of received spam, as determined by passive OS fingerprinting.



Operating Systems of Spamming Hosts

<i>OS</i>	<i>Clients</i>	<i>Total Spam</i>
Windows	854404 (70%)	5863112 (58%)
Linux	28132 (2.3%)	557377 (5.5%)
FreeBSD	6584 (0.5%)	152456 (1.5%)
MacOS	2944 (0.2%)	46151 (0.4%)
Solaris	1275 (< 0.1%)	18084 (0.2%)
Unidentified	128580 (10.4%)	1212722 (12%)
No Fingerprint	204802 (16.7%)	2225410 (22%)
Total	1228403	10103837

The operating system of each unique sender of received spam, as determined by passive OS fingerprinting.



Operating Systems of Spamming Hosts

OS	Clients	Total Spam
Windows	854404 (70%)	5863112 (58%)
Linux	28132 (2.3%)	557377 (5.5%)
FreeBSD	6584 (0.5%)	152456 (1.5%)
MacOS	2944 (0.2%)	46151 (0.4%)
Solaris	1275 (< 0.1%)	18084 (0.2%)
Unidentified	128580 (10.4%)	1212722 (12%)
No Fingerprint	204802 (16.7%)	2225410 (22%)
Total	1228403	10103837

- 70% of hosts that we receive spam (and 95% that we have a fingerprint for) run a version of Windows (*cf.* most legitimate mail servers are Unix-based)

The operating system of each unique sender of received spam, as determined by passive OS fingerprinting.



Operating Systems of Spamming Hosts

OS	Clients	Total Spam
Windows	854404 (70%)	5863112 (58%)
Linux	28132 (2.3%)	557377 (5.5%)
FreeBSD	6584 (0.5%)	152456 (1.5%)
MacOS	2944 (0.2%)	46151 (0.4%)
Solaris	1275 (< 0.1%)	18084 (0.2%)
Unidentified	128580 (10.4%)	1212722 (12%)
No Fingerprint	204802 (16.7%)	2225410 (22%)
Total	1228403	10103837

The operating system of each unique sender of received spam, as determined by passive OS fingerprinting.

- 70% of hosts that we receive spam (and 95% that we have a fingerprint for) run a version of Windows (*cf.* most legitimate mail servers are Unix-based)
- 4% Unix hosts send 8% email \Rightarrow a fraction of the spamming infrastructure is still unix-based




But what about DNS Blacklists?

Each spamming host IP is queried in 8 blacklists **at the time the spam is received**



But what about DNS Blacklists?

Each spamming host IP is queried in 8 blacklists **at the time the spam is received**




Minimum Number of DNSBLs listing an IP
for each type of received spam



But what about DNS Blacklists?

Each spamming host IP is queried in 8 blacklists **at the time the spam is received**




Minimum Number of DNSBLs listing an IP
for each type of received spam



But what about DNS Blacklists?

Each spamming host IP is queried in 8 blacklists **at the time the spam is received**




- 80% of spamming hosts are listed in *at least one* blacklist

Minimum Number of DNSBLs listing an IP
for each type of received spam



But what about DNS Blacklists?

Each spamming host IP is queried in 8 blacklists **at the time the spam is received**




- 80% of spamming hosts are listed in *at least one* blacklist
- A much *higher* fraction of Bobax drones are blacklisted

Minimum Number of DNSBLs listing an IP
for each type of received spam



But what about DNS Blacklists?

Each spamming host IP is queried in 8 blacklists **at the time the spam is received**



Minimum Number of DNSBLs listing an IP
for each type of received spam


- 80% of spamming hosts are listed in *at least one* blacklist
- A much *higher* fraction of Bobax drones are blacklisted
- A large fraction (50%) of spam hosts using short-lived routes are not listed in *any* blacklist



Effectiveness of DNS Blacklists: The Full Story




Effectiveness of DNS Blacklists: The Full Story



Cumulative fraction of emails listed in each blacklist (blacklists sorted from most aggressive to least aggressive)




Effectiveness of DNS Blacklists: The Full Story



Cumulative fraction of emails listed in each blacklist (blacklists sorted from most aggressive to least aggressive)



Effectiveness of DNS Blacklists: The Full Story




- 80% spammers blacklisted— isn't that pretty good?

Cumulative fraction of emails listed in each blacklist (blacklists sorted from most aggressive to least aggressive)



Effectiveness of DNS Blacklists: The Full Story




- 80% spammers blacklisted— isn't that pretty good?
 - **Only if you check 8 blacklists**

Cumulative fraction of emails listed in each blacklist (blacklists sorted from most aggressive to least aggressive)



Effectiveness of DNS Blacklists: The Full Story




- 80% spammers blacklisted— isn't that pretty good?
 - **Only if you check 8 blacklists**
- 95% bots blacklisted

Cumulative fraction of emails listed in each blacklist (blacklists sorted from most aggressive to least aggressive)



Effectiveness of DNS Blacklists: The Full Story




- 80% spammers blacklisted— isn't that pretty good?
 - **Only if you check 8 blacklists**
- 95% bots blacklisted
 - Each bot sends low volumes of email to **many** domains.

Cumulative fraction of emails listed in each blacklist (blacklists sorted from most aggressive to least aggressive)



Effectiveness of DNS Blacklists: The Full Story



Cumulative fraction of emails listed in each blacklist (blacklists sorted from most aggressive to least aggressive)

- 80% spammers blacklisted— isn't that pretty good?
 - **Only if you check 8 blacklists**
- 95% bots blacklisted
 - Each bot sends low volumes of email to **many** domains.
 - Indicates how massively collaborative spam filters could be useful

