Not-a-Bot (NAB): Improving Service Availability in the Face of Botnet Attacks

Ramakrishna (Ramki) Gummadi
MIT

Hari Balakrishnan (MIT), Petros Maniatis and Sylvia Ratnasamy (Intel Research)

The problem: Service unavailability

Subject: Mail delivery failed: returning message to sender
From: Mail Delivery System <Mailer-Daemon@csail.mit.edu>
Date: 4/21/2009 7:56 AM
To: ramki@csail.mit.edu

This message was created automatically by mail delivery software.

A message that you sent could not be delivered to one or more of its recipients. This is a permanent error. The following address(es) failed:

voelker@cs.ucsd.edu
SMTP error from remote mail server after end of data:
host cse-inbound.ucsd.edu [132.239.51.75]: 554 Service unavailable; Client host [outgoing.csail.mit.edu] blocked using Barracuda Reputation;
http://bbl.barracudacentral.com/q.cgi?ip=194.209.131.192

Bounced email

Crypto-gram

Misclassified email

Schneier's Crypto-Gram is getting flagged as spam by Razor. The reason is that some spam-detecting software will try to automatically detect spam *and then automatically report it*. So somebody's SpamAssassin mistakenly concludes that a copy of Crypto-Gram is spam and reports it to Razor, and this happens a few times over; now *everyone who uses Razor* will automatically be advised that Razor considers Crypto-Gram to be spam!

Botnets: Reduce service availability

- Email: 85% of spam from top six botnets
  - Over 95% of all inboxes affected
  - 120 billion messages/day: Overloaded mail servers
- DDoS [06]
- Click-fraud: ad fraud, search engine fraud
  - ~ 15% of all ad clicks
  - Compromise search results

Question: General way to distinguish bots from humans?

Existing solutions

CAPTCHAs

Drawback: Intrusive

User Account Control

Drawback: Default "yes" [Whitten, Tygar '99]

How to distinguish humans from bots automatically?

Strawman: Attesting human activity with Trusted Platform Modules

Google Suggest BETA

Web Images Groups News Froogle Local[New!] more »

tom cruise

tom cruise	6,670,000 results
tom clancy	1,990,000 results
tom cruise movies	2,670,000 results
tom cruise height	215,000 results
tom collins	5,780,000 results
tom cruse	145,000 results
tom clancy books	730,000 results
tom cruise filmography	481,000 results
tom chaplin	552,000 results
tom cochrane	347,000 results

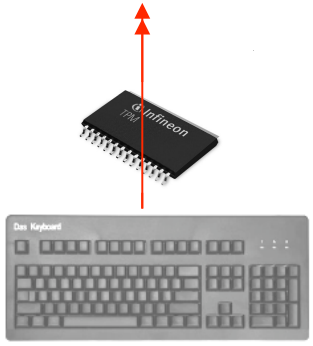

Attested Keystrokes

Problems with the strawman

Google™ Suggest BETA

Web Images Groups News Froogle Local[New!] more »

tom cruise

tom cruise	6,670,000 results
tom clancy	1,990,000 results
tom cruise movies	2,670,000 results
tom cruise height	215,000 results
tom collins	5,780,000 results
tom cruse	145,000 results
tom clancy books	730,000 results
tom cruise filmography	481,000 results
tom chaplin	552,000 results
tom cochrane	347,000 results

Attested Keystrokes

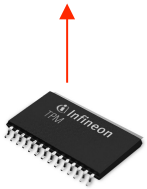

Slow

High-rate clicks

Browser

OS

Assumptions and Requirements

- Assumptions
  - Untrusted OS
  - Verifiable TPM bootup
  - Correct implementation of cryptographic primitives
- Requirements
  - Automatic
  - Fast (handle interactive traffic)
  - Small TCB (Trusted Computing Base)
  - Preserve privacy and anonymity

TPM Background

- Small, physically sealed chip
- Internal private key for measuring and reporting system integrity
- Two relevant protocols
  - Direct anonymous attestation
    - Group signatures using a key K_{priv}
  - Sealed storage
    - Secure location to store K_{priv} until system integrity verified

NAB (Not-A-Bot) Architecture

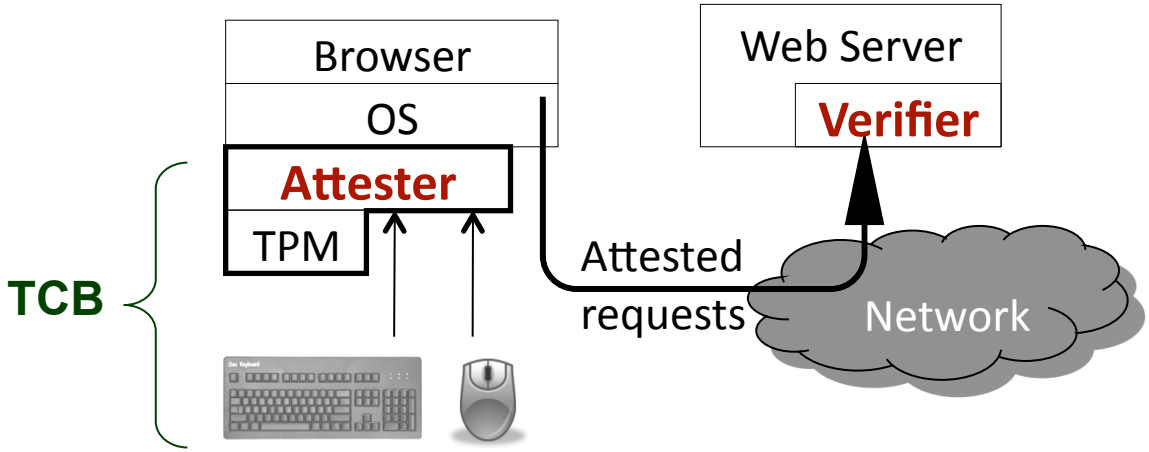

- Goal: Attest all human requests, reduce attested bot requests
  - No blacklisting: human requests from compromised hosts still receive service

Attestation security properties

- Non-transferable
  - Cannot generate at one host, use at another
- Bound to request content

  - No way to send spam from bots using one gmail account
- Single-use (verifier detects dupes)
- Limited valid time-window

When to attest?

- Simple, timing-based attestation
  - Requires human activity
- Allow attestation when request received within $\delta_{\{k,m\}}$ of last keyboard, mouse click
- Attester provides attestation only if $\delta_{\{k,m\}} < \Delta_{\{k,m\}}$ (= 1s for email)
  - Verifier checks $\delta_{\{k,m\}}$ in attestation for validity
- Reduces click harvesting

What to attest?

- Challenger-specific
  - Cannot be retargeted

To: bob@b.org
From: alice@a.org
Hi Bob,…

- Responder-specific
  - Cannot exploit manually configured whitelisting

- Content-specific
  - Cannot modify or piggyback on valid messages

What is in an attestation?

- Signed SHA-1 hash of message
- 160-bit signed nonce
  - Verifier stores nonces for application-defined period, checks duplicates
- Optional $\delta_{\{k,m\}}$ values (omitted for privacy)
- Certificate to verify K_{priv}

Attestation	$K_{priv}\{H(msg)\}$	Siged Nonce	$K_{priv}\{\delta_m, \delta_k\}$	certified K_{pub}

Attester Interfaces

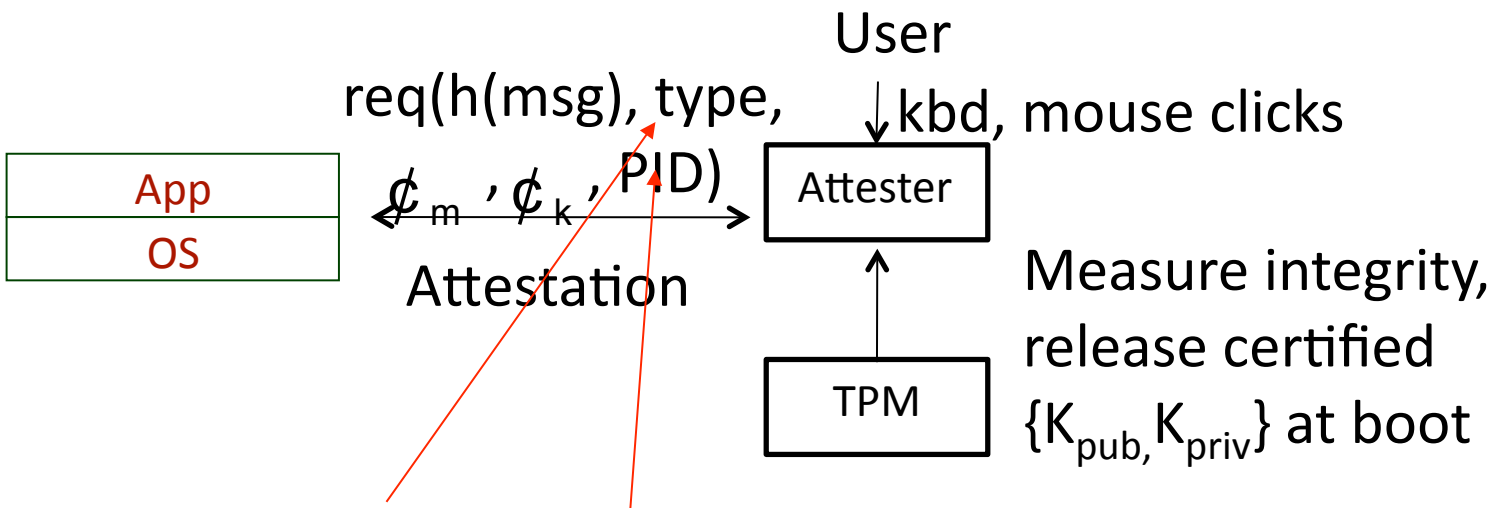

Type: Anonymous or non-anonymous

PID: Delayed attestation release for a process

Attester Operation

- Installation: Set to use TPM register# 18:

  PCRExtend(18, Hash(Attester Code))

- Sealing private key K_{priv} to host:

  S=*Seal*(18,K_{priv})

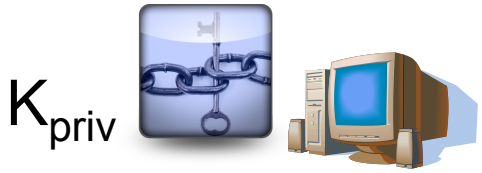

- Booting: Release K_{priv} to attester:

  K_{priv} =*Unseal*(S,(18,PCR_{18}))

Recomputed attester's hash

Verifier Operation

- Checks validity of K_{priv}, attestation, nonce
- Uses application-specific policies
- Email:

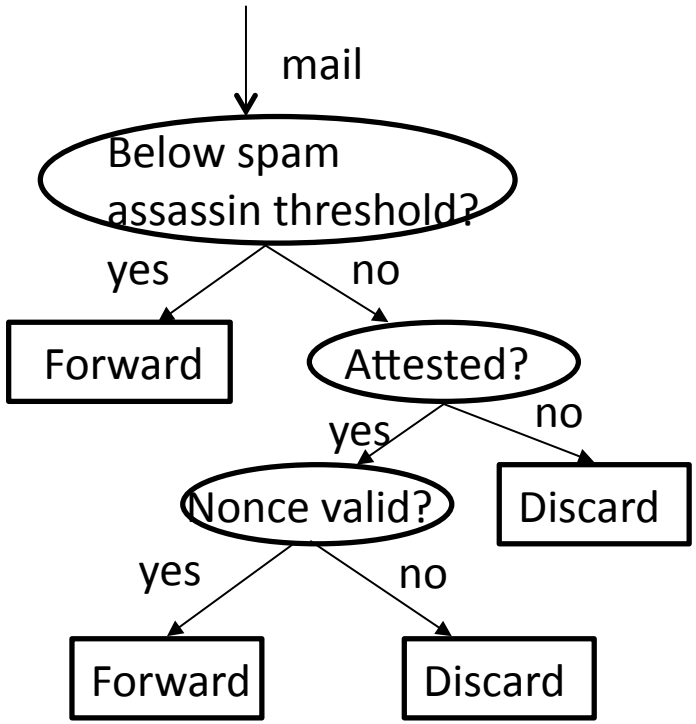

Email: Usage scenarios and incentives

- Mailing lists
  - Verifier checks subscription to mailing list name in “To:” field
- Offline mode
  - Attestation requested when user hits “send”
- Sender incentive
  - Better email reliability
- Recipient incentive
  - Reduced mail server load, better reliability

Request processing at verifier

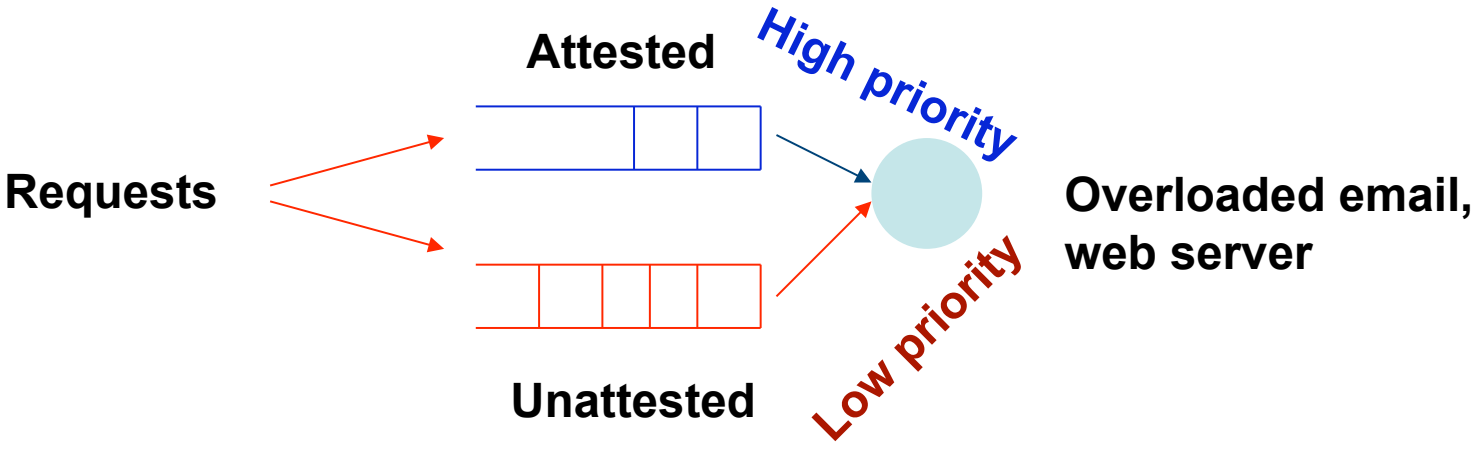

Prioritize attested requests

DDoS, Click-fraud: Usage and incentives

- Browser gets attestation when requesting document root (“http://foo.com/”)
  - Verifier stores attestation, accepts same attestation in future for all embedded links
  - 10 minutes expiry
    - Browser forced to use new attestation for next fetch
- Incentive: Attester distributed in search engine toolbars

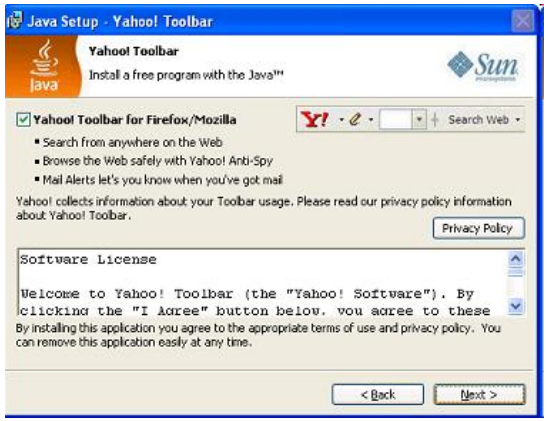

Evaluation

- Implemented attester with Xen VMM
  - Uses domain disaggregation [Murray et al.,'08]
  - Attester within a paravirtualized Xen domain built with miniOS, isolated from untrusted OS
- Trace-driven verifier evaluation
  - Click traces of 328 users in one month [Giroire et al.,'08]
  - Publicly available spam, DDoS and click-fraud traces
  - Worst-case scenario with adaptive bots

Attester evaluation

- CPU cost: At most 10 ms on 2 GHz CPU
  - RSA signatures, 1024-bit modulus
- Complexity metric: lines of code
  - Attester kernel module: 500 lines
  - miniOS: 30,000 lines
- Applications: NET::SMTP (Email), cURL (Web)
  - 250 lines of code modified
  - Attestations as extended protocol objects

Verifier evaluation

- Methodology: 328 click traces at 1s intervals
  - Adaptive bot: steals as many clicks as possible
  - Generates traffic using all stolen clicks
  - Compare against status quo (normal bot without NAB) within the same time
  - 328 data points, one for each user's trace
- Other metrics
  - Nonce storage cost (< 600 GB for one-month nonces with million clients)
  - Throughput: 10,000 attestations/s

Spam mitigation

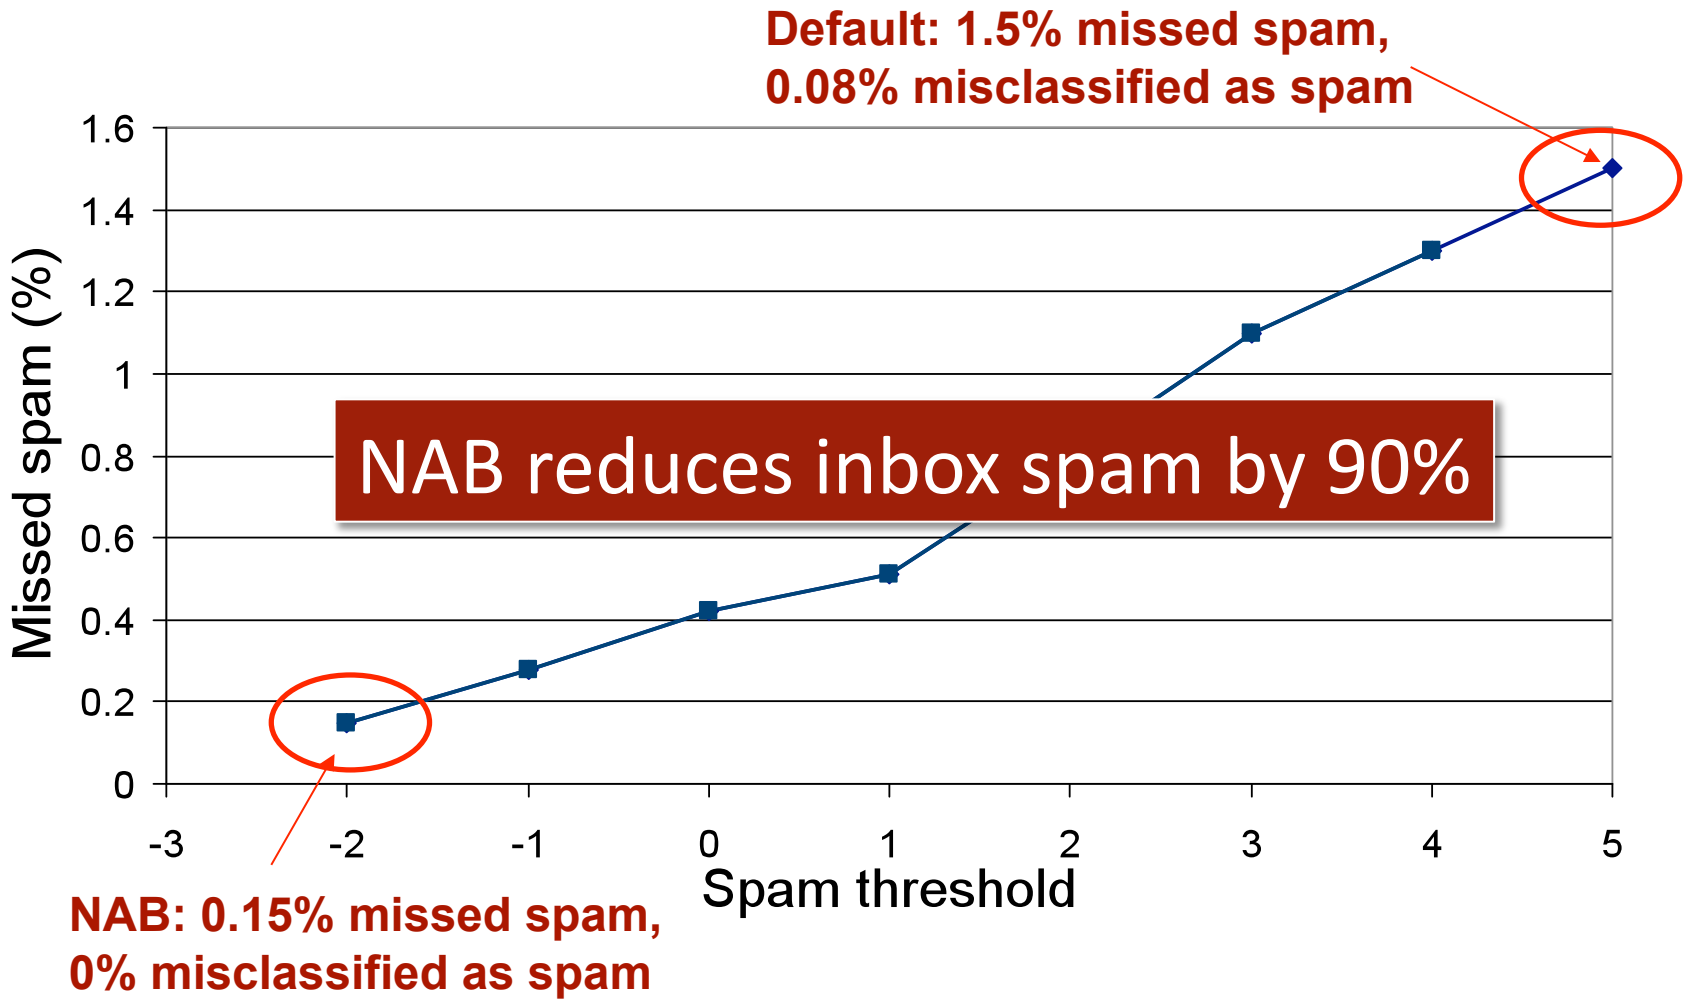

Email server overload mitigation

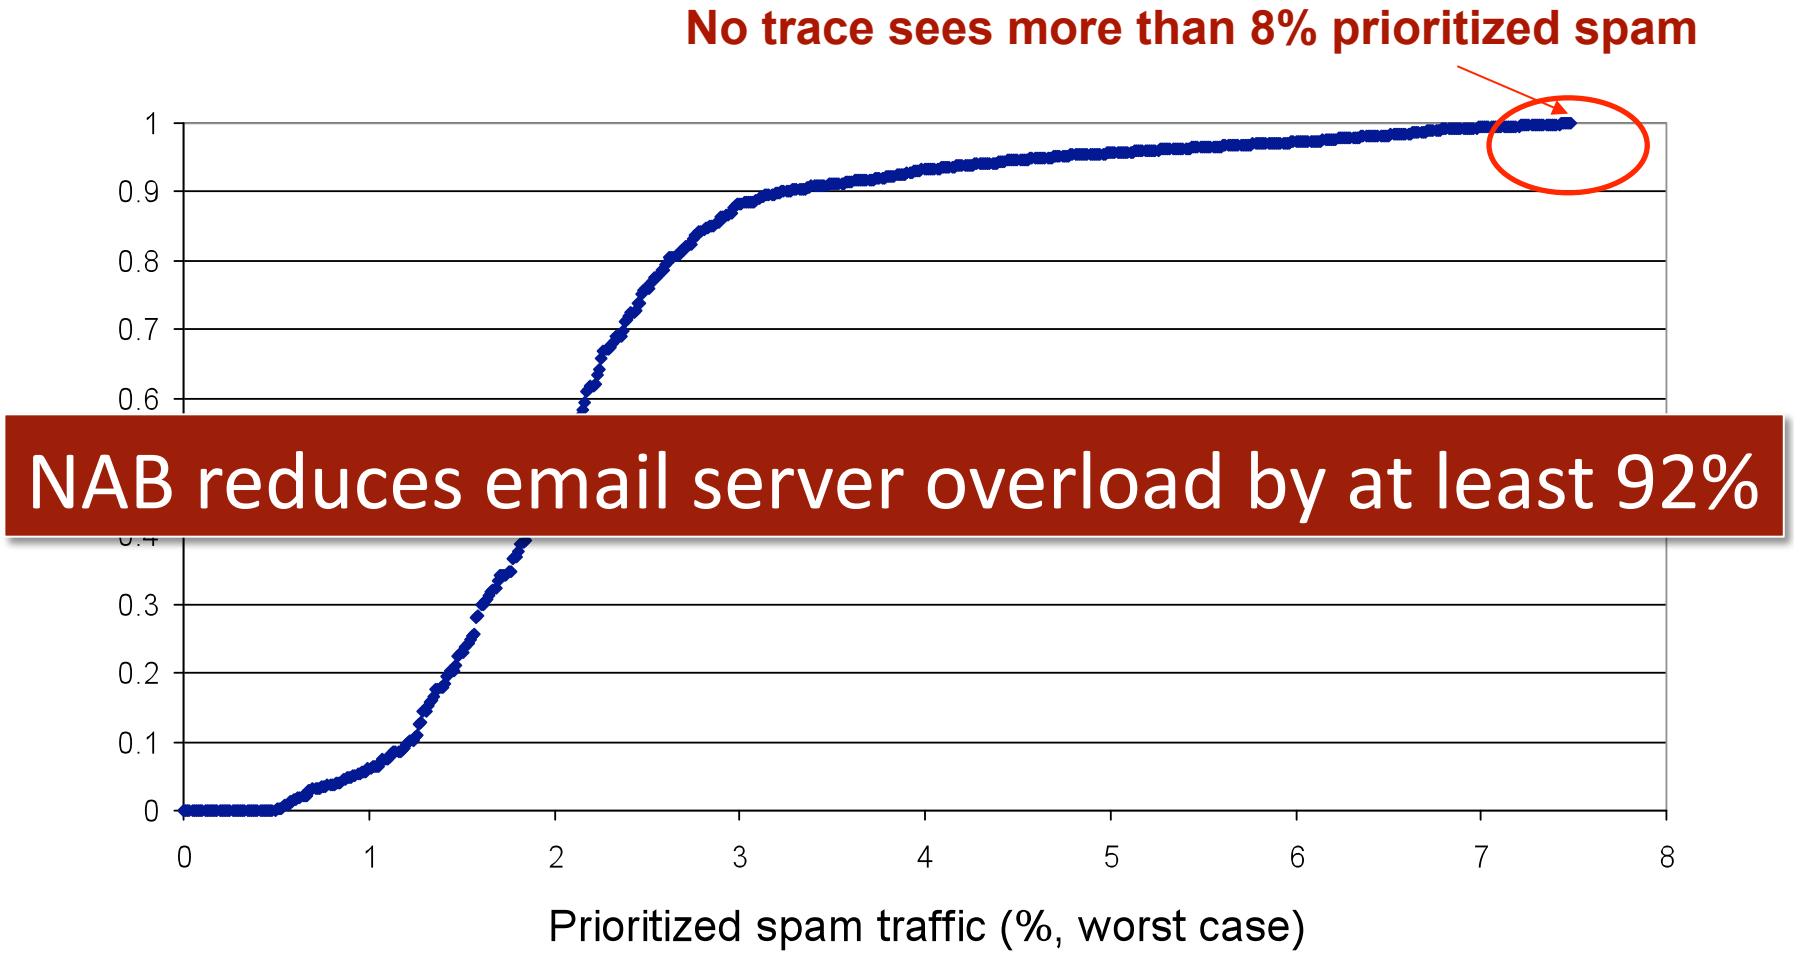

NAB reduces email server overload by at least 92%

DDoS mitigation

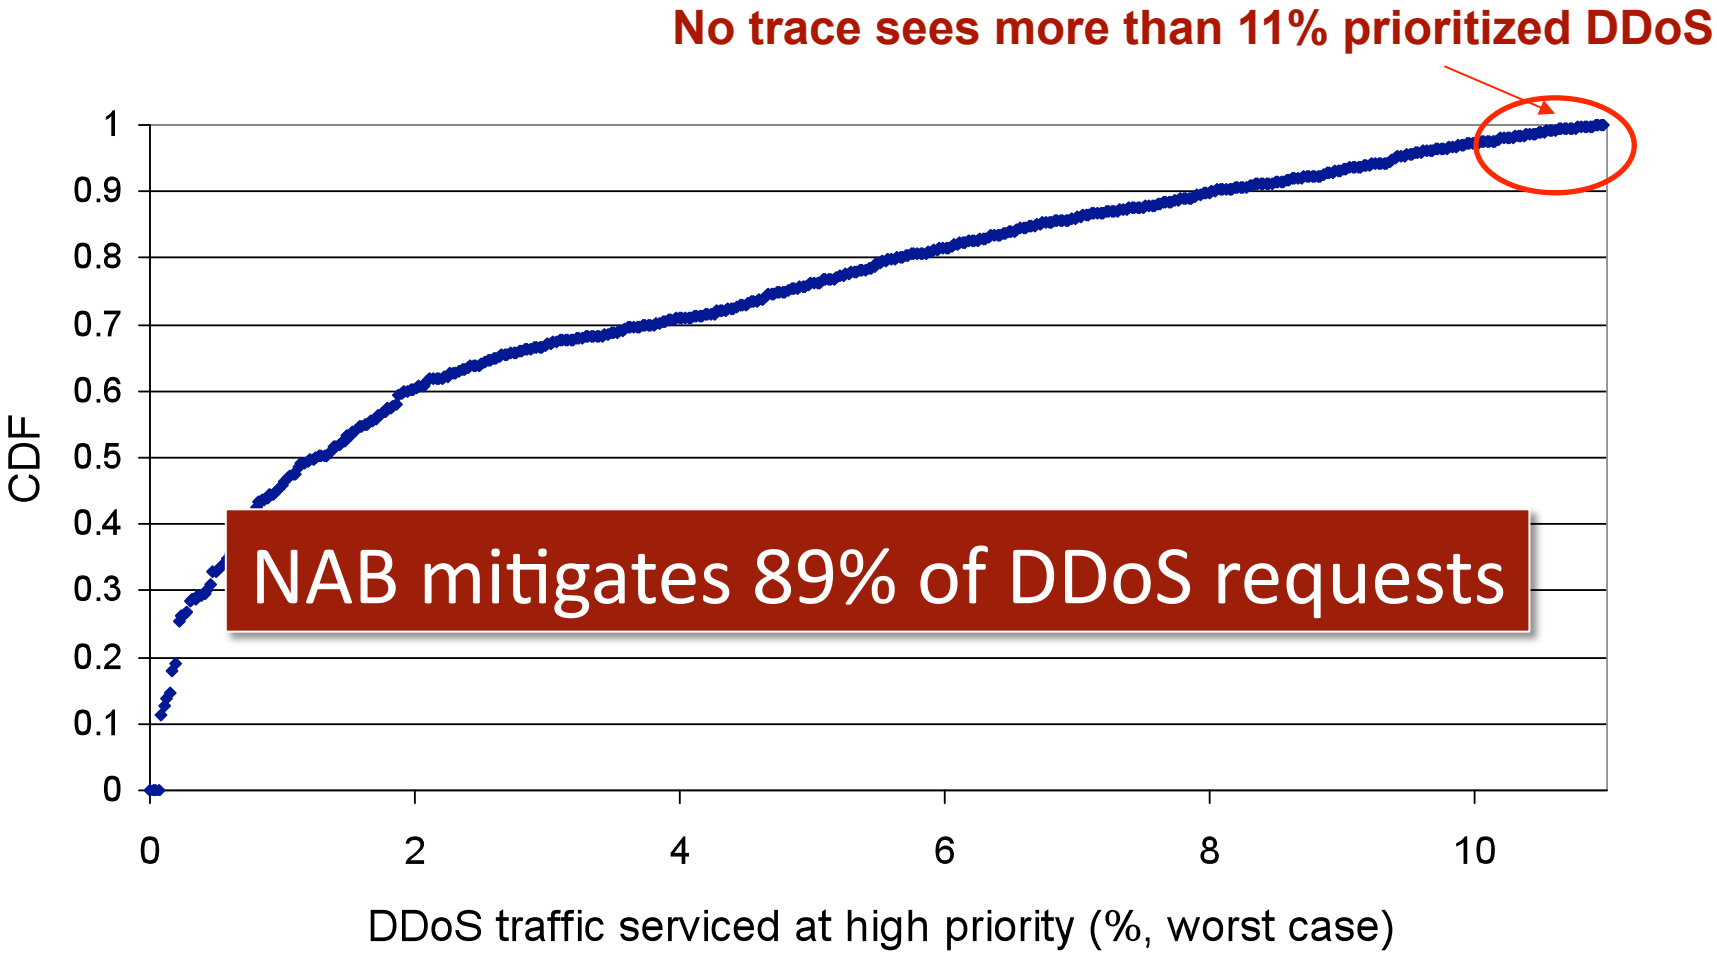

Click-fraud mitigation

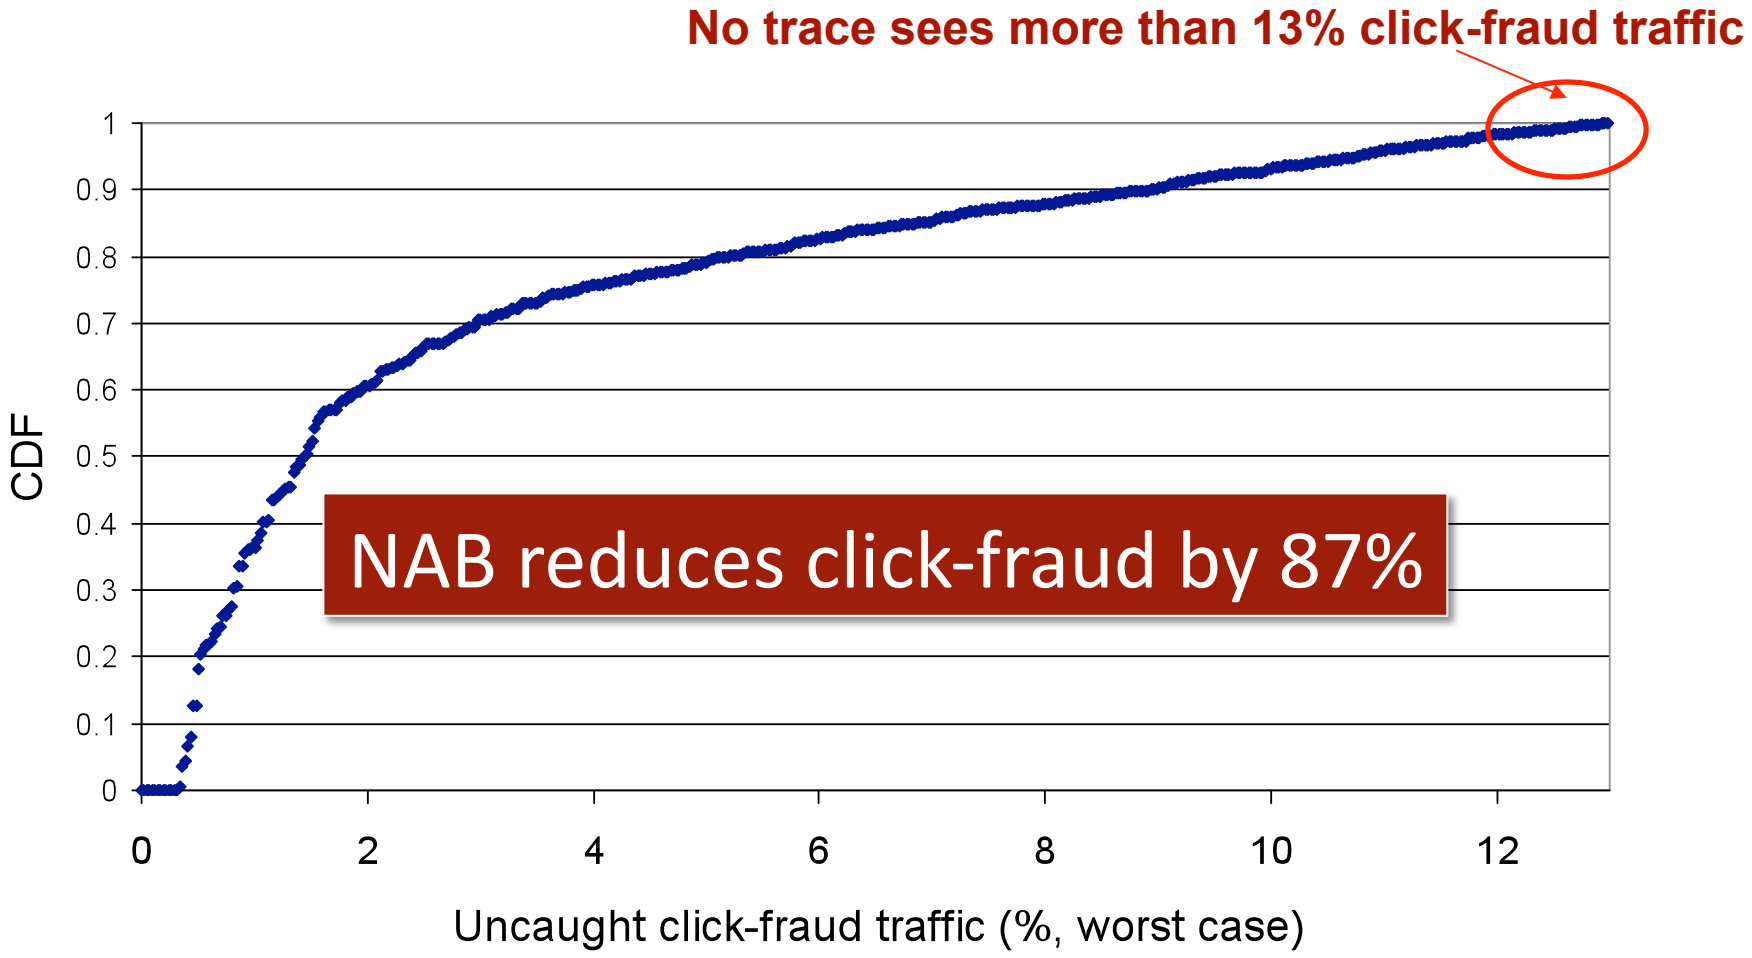

Related work

- Human activity detection
  - CAPTCHAs [Ahn et al.,'03]
    - Susceptible to man-in-the-middle attack
  - Nexus [Williams et al.,'08]
    - Not for commodity OSes
- Mitigating spam, DDoS, click-fraud
  - Spam: Occam [Fleizach et al.,'07], SPF, DKIM
  - DDoS: Path validation, bandwidth-as-payment
  - Click-fraud: Syndicators, clickable CAPTCHAs
  - Mostly specialized, share little commonality

Conclusions

- NAB: Improves service availability in the presence of botnets
  - Even on botted hosts, users get ~ 100% service
    - No blacklisting
  - De-prioritize or drop up to 90% bot traffic
- Automatic content- and machine-specific attestations
- Single abstraction, support for application-specific verifier policies
- Future work: Attestation without virtualization