DNSOPS WG IETF-67

SPF/Sender-ID DNS & Internet Threat

Douglas Otis

Doug_Otis@trendmicro.com

http://www.ietf.org/internet-drafts/draft-otis-spf-dos-exploit-01.txt

How does SPF/Sender-ID utilize DNS?

- SPF/Sender-ID use TXT (16) or SPF (99) RRs
- The SPF records may chain 10 or 11 times using Redirect, Include, and Exp macro mechanisms
- May publish as many IPv4 & 6 CIDRs that fit
- Each chained set may contain 10 macro-expanded mechanisms with domain name & CIDR overlays
- Mechanisms may resolve A, AAAA, MX, r-PTR or test existence of A or AAAA using Exist macro
- Mechanisms for MX or r-PTR may require 10 additional DNS transactions each. (10 x 10 = 100)

What names are checked?

- SPF records checks SMTP clients against both:
  - RFC 2821 EHLO hostname
  - RFC 2821 MAIL FROM email-address
- The scope for Sender-ID adds one:
  - RFC 2822 Resent-Sender email-address
  - RFC 2822 Resent-From email-address
  - RFC 2822 Sender email-address
  - RFC 2822 From email-address
- Some propose checking DomainKey/DKIM domains

SPF amplification (prior canvassing assumed)

~8k bits per message (small message)

2304 bit q + 2784 a = 5088 bits (example attack)

5088 bits x 100 = 508 kbits / name evaluation

508 kbits / 8kbits = SPF op baseline amplification = 64:1

1-~3 names evaluated per message

1-~20 recipients per message

1-~3 evaluation points per recipient

Low (64 kByte x 1 x 1 x 1) = 64:1 for 508 kbits/msg

Med (64 kByte x 1 x 2 x 2) = 320:1 for 2032 kbits/msg

High (64 kByte x 3 x 20 x 3) = 11,430:1 for 91.4 mbit/msg

Estimating SPF related DDoS Potentials

- Each second a session is blocked per 16,400 mailboxes
- Each second a message is tagged per 7,040 mailboxes
- Each second an untagged message per 27,400 mailboxes
- Millions of compromised systems send > 70% of spam
- > 80% of email is spam
- Bot-nets campaigns commonly operate 50k Bots in concert
- 230 kbit/sec In & 278 kbit/sec Out per SPF script op
  (29 kBytes In & 35 kBytes Out per SPF script op)
- > 4 SPF script operations per received message
- 100k Bots at 2 SPF ops with average of 10 recipients can produce 8 Gb/s In and 10 Gb/s Out at 1 msg/min at 0 cost
- Congestions avoidance is also circumvented in SPF libraries

DNS Label Components from SPF Macros

%{symbol | rev-label-seq | # of right-most labels | chars to “.” }

rev-label-seq = “r”

chars to “.” = “-” | “+” | “,” | “/” | “_” | “=”

# of right-most labels = 1 - 128

symbols :

- s = email-Address or EHLO
- l = left-hand side email-address
- o = right-hand side of email-address
- d = email-address domain or EHLO
- i = SMTP client IP address decimal octet labels
- p = r-PTR domain validated with IP address
- v = "in-addr" IPv4, or "ip6" if IPv6
- h = EHLO hostname

SPF/Sender-ID deployment

- ~3% of domains with MX RR publish SPF
- Fortune 100 & Top 20 domains:
  - 72% / 70% offer no SPF records
  - 6% / 10% SPF records fail Neutral
  - 13% / 10% SPF records fail Softfail
  - 9% / 10% SPF records fail Fail
- Abusive sources:
  - 77% offer no SPF records
  - 0.2% SPF records Pass
  - 14% SPF records offer Neutral results
  - 6% SPF records offer Softfail results
  - 2.6% SPF records offer Fail results

Hard to detect SPF enabled attacks

- Flood of DNS traffic from highly distributed sources
- Sources are within otherwise well managed domains
- Queries may exhibit large random names for:
  - Wildcard SPF RRs
  - Invalid address records
- Packet source/destination addresses are all valid
- Email logs do not explain the high level of DNS traffic
- Traffic originates from DNS serving access points & MTAs
- Concurrent with legitimate DNS traffic
- Could also be concurrent with suspicious poisoning traffic
- Could also be concurrent with a high level of DNS timeouts

Ideas for preventing SPF attacks

- Establish AUPs prohibiting use of SPF script libraries
- Return 0 answers for all SPF records
- Promote safe techniques with deterministic DNS overhead
- Authenticate the client before evaluating message content
- Avoid running scripts directly from unknown DNS or clients

Questions?