DNS Poisoning: Recent Developments

Notification and Discussion

David Dagon[1] Paul Vixie[2]

[1]Georgia Institute of Technology
Atlanta, Georgia

[2]ISC Redwood City, California

OARC Sept 2008 - Ottawa

Outline

- Our Motivation: DNS Poisoning Detection
- *Active* probes
  - Anecdotes of Kaminsky-class attacks
- *Passive* collection (SIE)
  - Recursive View: Local DNS Poisonings
  - Authority View: remote poisoning detection
  - Passive View: weather-map
- Conclusion

DNS Poisoning Risks

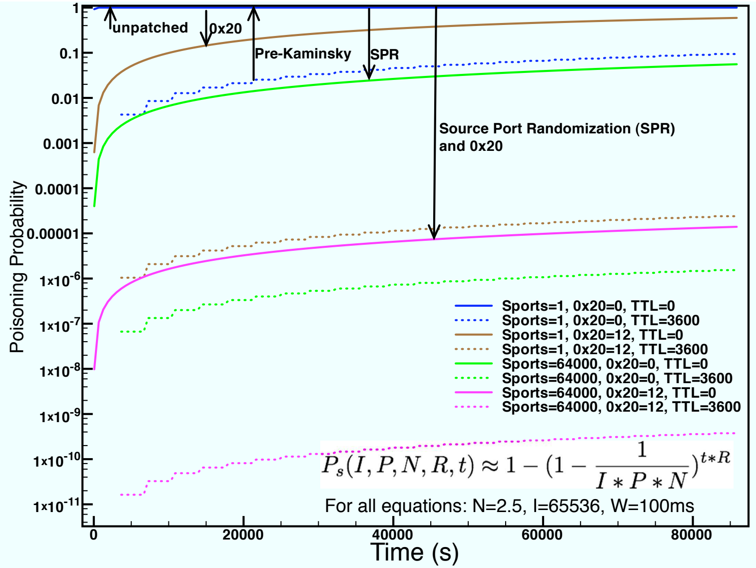

DNS Poisoning Risks

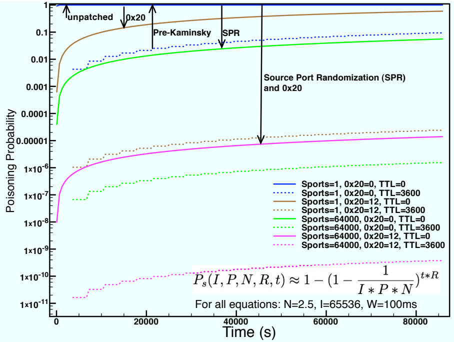

Salient points:

- The plot uses Herb's formula (tweaked for birthday)
- When TTL mattered, the risks were shown in the dotted lines
- After Kaminsky, TTL does not matter, and risks have shifted to the solid line
- Note that 16-bit resolution is nearly 100% poisonable in seconds
- The interim solutions (SPR, 0x20, etc.) can reduce risk somewhat.

Study Methodology

- Now that DNS exploits are "in the wild" we wish to detect them. How?
- Two general approaches:
  - Active probes (POPE)
  - Passive data collection (SIE)
- We present early results from these two large, infrastructure-intensive efforts

Study Methodology

- General architecture for poisoning detection using active DNS probes
  - Key idea: We can't hand verify everything; we need to build a high-quality filter of suspect results
  - We created a table of: large list of open recursives ($\approx$ 20M) and a large list of phishable domains (tens of thousands). We probed each host for each domain, repeatedly (observing TTL).
  - The cross product is enormous
  - Not every answer can be visited by a honeypot or hand analyzed. How do we find "interesting" results?

Study Methodology

- We therefore built “POPE”: A DNS Monitoring Infrastructure
  - Based on “King: Estimating Latency Between Arbitrary Internet End Hosts” (IMC 2002)
  - (Why call is POPE? Because popes are slightly better than kings.)
  - Uses RTT deltas to find “interesting” things
  - Theory: no poisoning would result in *improved* DNS service times
  - Statistically unusual measurements trigger heavy-weight (e.g., honeypot) analysis

Methodology

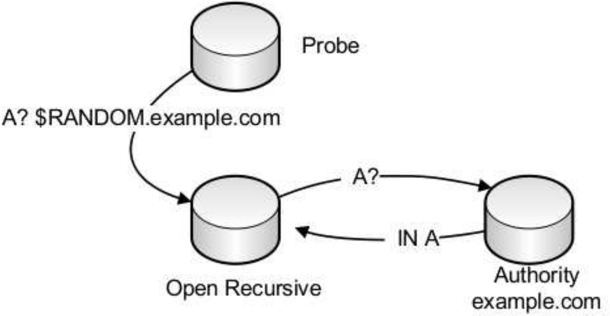

Methodology

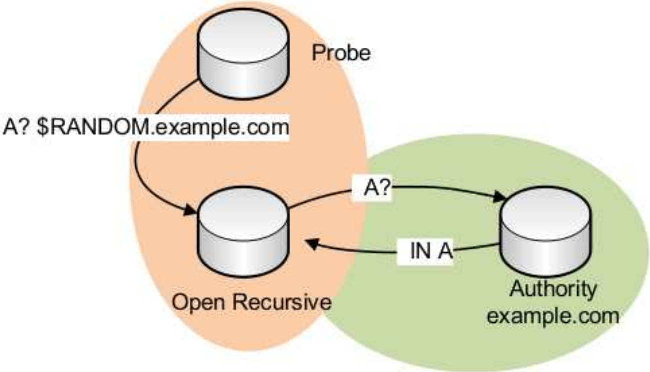

DNS Poisoning Risks

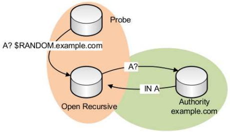

Salient points:

- We ask for a non-existent child record; the recursive fetches
- We ask *again* for the non-existent record; the recursive answers from cache
- We subtract the RTT times, to find the estimated RTT between the recursive and authority. We do this for 20M recursives x 10K authorities
- We develop statistical distributions of RTT from open recursives to authorities
- *Theory:* A poisoning would in most cases increase the RTT
- Changes in RTT time are handed off to heavy-weight honeypots, and ultimately expensive hand analysis.

DNS Poisonings

- Using POPE-style measurement, one can observe “interesting changed RTTs”
  - Further investigation can confirm poisoning
  - 16-bit resolvers returning phish records, however, suggest Kaminsky-class poisonings
  - We still do not know the *actual* cause, of course
- Example follows
  - Focus on US financial zones, and a few other zones
  - Localized trends observed; look at the *delta* over time, more than the y-axis dimension.

DNS Poisonings

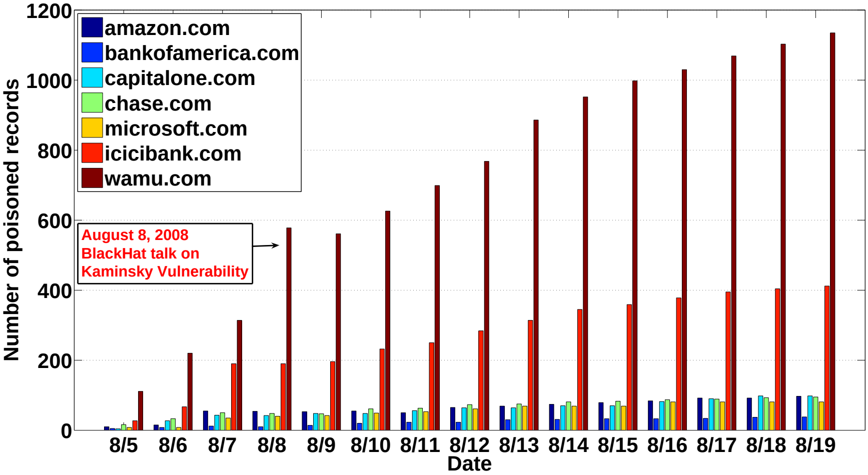

DNS Poisonings

Zone	Successful Poisonings	Poisoned Sub-zones	Unique IPs in Answers
amazon.com.	944	4	11
bankofamerica.com.	351	1	25
capitalone.com.	960	3	18
chase.com.	947	2	27
microsoft.com.	827	4	13
icicibank.com.	4416	7	11
wamu.com.	11050	6	24

(Note: “malicious” IPs were hand verified)

Passive Detection

- Active probes of course have a “cost”
  - RFC 1262 needs a refresh
  - Numerous abuse@ complaints
  - Obviously, we can’t do Internet-wide monitoring of poisoning using iterative probes
  - At best, active probes give us hints of where to look
- What can we do instead using passive data collection?

Passive Detection

- Three general detection positions (based on *monitoring location*):
  - Local: Recursive View (Technologies: ICMP(3,3), Excess answers that do not match queries, etc.)
  - Remote: Authority View (Technologies: ICMP(3,3))
  - Omniscient: high-level view from SIE
- We'll present some early results from analysis of passive collection

Detection and Response: Passive Collection

- SIE
  - Large collection of above-the-recursive AA-bit DNS traffic
  - Aggregated at 2 (soon 3) switch centers
    - Data is replayed and then deleted
    - No logging, no storage allowed; just real-time analysis
  - Detects more than just poisoning
  - “Channel” metaphor:
    - E.g., changed NS, flux, etc.
    - Allows re-broadcast and sharing of analysis

Planned SIE Switch Centers

SIE Data Rates

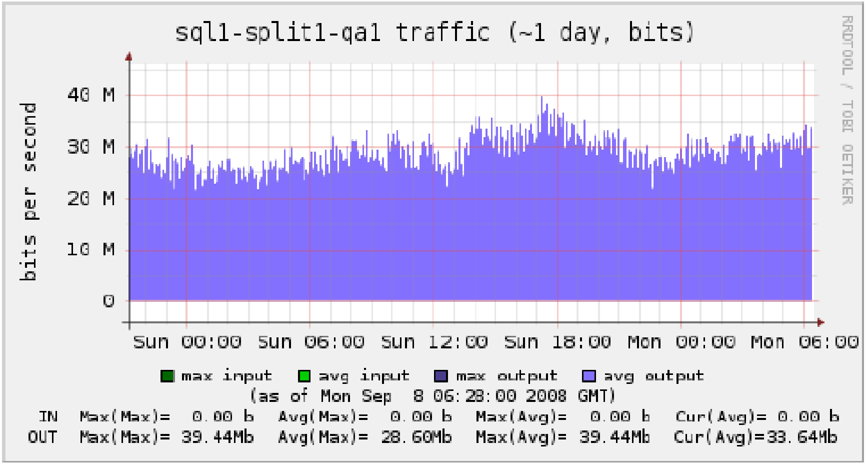

SIE Analysis Example #1

- “Bootstrapping”: using one known bad
- E.g., consider “channel 8”’s list of changed NS records. Let’s look for fluxing fake “antivirus” sites

```
ncaptool -n - -fvmg - dns "regex=(antivirus)" 2>

[116 nf -] 2008-08-28 03:22:27.163092000 [000000
   [85.17.45.51].53 [66.28.28.210].48236 \
   dns QUERY,NOERROR,63378,qr|aa|ra \
   1 ns2.antivirus-xp-08.net,IN,A \
   1 ns2.antivirus-xp-08.net,IN,A,60,85.17.45.51
   2 antivirus-xp-08.net,IN,NS,600,ns2.antivirus
   antivirus-xp-08.net,IN,NS,600,ns1.antivirus-x
   2 ns1.antivirus-xp-08.net,IN,A,60,85.17.45.51
   .,CLASS4096,TYPE41,32768,[0]
```

SIE Analysis Example #3: IRS eFile

- IRS' eFile is often targeted by phishers
- One scheme: a phish website is hosted on a fluxing botnet
- SIE allowed for rapid identification of *all* RRsets for phish eFile site
- Key: SIE also allowed for exploration of IP $\longrightarrow$ domain mappings
- New, still-dormant eFile phish sites were thereby found
- IRS investigator's finding: "The domains associated with both the site and the nameserver's were dropped less than 15 minutes later. Zero victims."

Conclusion

- DNS Poisoning is a “refreshed” tool in the attacker’s kit
  - Many instances observed; anecdotes suggest correlation with Kaminsky-class attacks
- We can use active probes ... to a point
- Passive collection/analysis will likely provide a more scaled detection
- Detection/sensor needs are a current priority