nominet

DNSSEC Support in SOHO CPE

OARC Workshop
Ottawa
24th September 2008

nominet

or: “How not to write a DNS proxy”

Study Details

nominet

"What is the impact of DNSSEC on consumer-class broadband routers"?

- Joint study between Nominet UK and Core Competence
- Core Competence funded by Shinkuro, Inc., under contract from ISOC, ICANN, Afilias
- Conducted July and August 2008
- Expansion of .SE's previous study

Devices Tested

nominet

- 4 SOHO Firewalls
- 12 Dual Ethernet “Gateways”
- 8 ADSL Routers

- Selected based on market share and popularity
- All tested with “out of the box” configuration as far as possible
- NAT-PT and DHCP

nominet

Test Environment

- Queries sent with “dig” and custom Perl scripts using Net::DNS
- Packets captured on both LAN and WAN side of unit under test with Wireshark/tcpdump
- Queries sent both directly to the unit under test (“proxy mode”) and through the unit to the upstream RDNS (“routed mode”)
- Upstream DNS on a private network, with a fake “root”
- RDNS and ADNS running Bind-9.5-P1

Test Environment

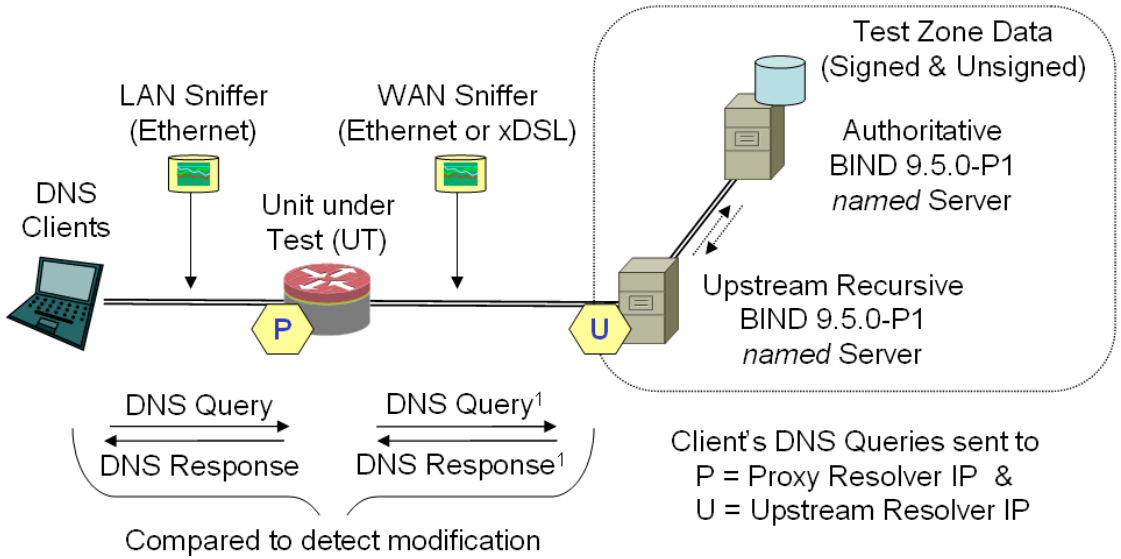

DHCP Behaviour

- 15 devices put their own (LAN) IP address in their DHCP server's "Domain Name Server" option
  - But 9 of those 15 have no way to change the DHCP settings
- A further six devices put the upstream address in, but only once the WAN link is up ("chicken and egg" problem)
- The remaining three don't proxy by default

nominet

Proxy Behaviour #1

Devices that were “dumb” about DNS tended to do better than “smart” devices, but only so long as they did the rest of UDP/IP correctly:

- Fragment reassembly was a big problem
  - Some fragments black-holed
  - Some sent from the wrong Source IP
  - Typically evident in packets near the WAN MTU

Proxy Behaviour #2

nominet

Many implementors only appear to have read (some of) RFC1035, and no subsequent RFCs:

- Responses truncated at 512 bytes (without setting TC)
- Responses having TC flag cleared in transit
- Packets dropped in either direction when CD=1 or AD=1
- EDNS0 packets black-holed or rejected
- No support for failover to TCP
- QIDs not random [NB: this is for future study]

NAT-PT Behaviour

nominet

- Half of the devices tested had poor source port randomization in their NAT-PT logic
- Most (if not all) of those pick source ports sequentially
  - Risk of cache poisoning attacks not mitigated
- When combined with poor QID selection, severe risk of exposure to normal response spoofing attacks

nominet

Results

- Only six of 24 devices were mostly compatible with DNSSEC “out of the box”
- 18 of the 22 devices that actually do DNS proxying had limitations on packet size (512 bytes or ~MTU)
- 6 of those 22 had incompatibilities that effectively prevent use of “proxy mode” for DNSSEC
- However all devices handled DNSSEC correctly when using “routed mode”
- Only one device could proxy DNS over TCP

nominet

Unaffected Configurations

Anything using “route” mode:

- Fully validating local recursors
  - NB: some still prone to cache poisoning
  - Potential high load on authority servers
- Clients with hard-coded settings
  - NB: some clients (e.g. Mac OS X) make it hard to ignore DHCP settings. They default to adding the hard-coded list to the DHCP settings, not replacing them.

Good news!

That covers most configurations that would be used by more technically sophisticated users

Affected Configurations

nominet

Anything that uses DHCP to get DNS settings and where:

- the response is a large RRset (containing DNSSEC records or otherwise); or
- the server returns unexpected flags (c.f. Bind 9.4.1 bug found in the .SE study); or
- the client is a security-aware stub
  - Is this a likely deployment model for desktop DNSSEC?
  - Could a client detect whether the proxy is “good”, and failover to fully recursive otherwise?

nominet

Study Follow-up

- IETF draft - BCP for how to write a proxy
- Vendor fixes?
- Research on the quality of PRNGs for
  - Source Port ID
  - Query ID
- Fuzzed queries and responses - can we actually crash the routers?
- “fpdns” Mk II - for identifying RDNS?
- How common is it to run a recursor behind NAT?