Whats Wrong With The DNS

Duane Wessels
The Measurement Factory/CAIDA
wessels@measurement-factory.com

October 3, 2006

About This Talk

- A collection of all DNS-related problems and issues that we know about.
- Roughly in the following order:
  - Protocol (IETF)
  - Implementaiton
  - Operational (*NOG)
  - Registry/Registrar (ICANN)
- Please help fill in missing pieces and correct mistakes.

Whats Right With The DNS

- Gotten us this far
- Very robust
- Scales pretty well
- Extensible (to some extent)

Protocol Issues

The 512 Byte Limit

- RFC 1035 limits UDP message size to 512 bytes.
- Expandable via EDNS0, but its still not as widely deployed as we'd like.
- IP fragmentation is a bummer.
- Predictions are for larger DNS messages due to IPv6, DNSSEC, etc.

What can be done?

- ENDS0
- Encourage use of TCP? (might introduce more problems)

Cache Poisoning

- Some resolver implementations are vulnerable to simple cache poisoning.
- Bad news if resolver is open for recursion.
- Perhaps 10% of nameservers run Windows software, notoriously vulnerable to poisoning.

What can be done?

- Run up-to-date code
- On Windows, place a check in the "Secure cache against pollution" box.
- Deploy DNSSEC

Non-Existence of Non-Terminals

- Caching resolvers don't learn that *.local* is an invalid TLD.
- Negative caching doesn't help.
- Causes root server pollution.

What can be done?

- Change the protocol?
- DNSSEC?

UDP as Transport

- UDP's statelessness enables spoofing-based attacks.
- Also allows one-way communication (no route back), which may result in extra traffic from panicky resolvers.

What can be done?

- UDP based protocols should give up when they don't get any responses after some reasonable amount of time.
- Reconsider the use of TCP.

DNSSEC Hurdles

- The Net has been slow to adopt DNSSEC.
- Still don't know how to manage root keys and rollover.

What can be done?

- Continue hammering out these details.
- Sit around and wait for some high-profile cases to show the world why DNSSEC is needed.

IPv6

- IPv6 doubles the number of A* queries generated by some applications.
  - Resolvers have been doing AAAA/A6 queries for quite some time
- Recent concerns about Windows Vista:

  "If you adopt Vista, your DNS traffic is going to double ... You're going to see brownouts. All of a sudden, it is going to be mud season on the Internet, where things will just be kind of slow and gooey." — Paul Mockapetris

What can be done?

- Send A and AAAA queries in parallel to reduce latency?
- Invent a new query type that would return both A and AAAA RRs?

Blocking 53/tcp

- Network admins have been taught to fear zone transfers and block 53/tcp entirely.
- Caused problems when Network Solutions used TCP to identify valid users during an attack.
- May become more of a problem as messages increase in size (DNSSEC, IPv6, etc).

What can be done?

- Teach operators that DNS over TCP is nothing to fear.
- Proactively test nameservers for TCP reachability.

Phishing, Character Sets

- One phishing technique is to use character sets with similar-looking letters.
  - accents on vowels, etc
- See also internationalization

What can be done?

- Never click on URLs from untrusted sources?
- Stop sending and accepting HTML emails?
- Should browsers display something other than DNS names/URLs to indicate the source of content?
  - X.509 certificate details
  - whois records

Internationalization

- Ambiguity in the specs imply only ASCII letters allowed.
- Many of the world's languages have non-ASCII characters.
- China has already created their own TLDs using punycode names.

What can be done?

- Change the protocol
- Don't let people who only speak English design any more protocols.

Implementation Issues

Exploitable Bugs

- Popular nameserver software (BIND, Windows) has a history of exploitable bugs.
- Bugs may lead to cache poisoning, denial of service, etc.
- Made worse if nameserver is open to recursion.

What can be done?

- Run up-to-date code.
- Add ACLs for recursion.

AAAA NXDomain/Servfail

- Broken (but still too-widely deployed) nameserver software returns NXDomain or Servfail for AAAA when A record exists.
  - Older Windows nameservers
  - Load balancers
- Sendmail has "WorkAroundBrokenAAAA" feature.

What can be done?

- Do a better job testing and reading RFCs.

Cache Snooping

- Cache snooping can reveal information about your organization
- Send queries with Recursion Desired bit cleared to see if certain names are in the DNS cache.

What can be done?

- Use ACLs to deny queries from outside your organization.
- Separate caching and authoritative DNS services.
- Add some slight randomness to TTLs to increase the difficulty of temporal correlations.

Root Server Pollution

- 75–98% of root-server traffic is unnecessary.
  - Repeats
  - Invalid TLDs
  - No route to host
- Cruft (cache misses) floats to the top
- UDP does not provide feedback loop
- Encourages operators to over-provision and use anycast.

What can be done?

- Run *dnstop* on your network.
- Run recent versions of resolver software.
- Implement BCP38, double-check packet filter rules

Man-In-The Middle Attacks

- If you can sniff and spoof, you can lie to resolvers and poison caches.
- Or, if you can't sniff, you might be able to guess query names and query IDs.

What can be done?

- Developers should use good random query IDs.
- Operators should deploy anti-spoofing filters.
- Deploy DNSSEC

RFC1918

- Everyone uses RFC1918 addresses
- Noone creates RFC1918 in-addr.arpa zones.
  - ip6.arpa either?
- AS112 project exists to sink huge amounts of PTR and UPDATE traffic for RFC1918 space.

What can be done?

- Operators should configure reverse zones for their RFC1918 space.
- Implementations could answer for RFC1918 zones by default.

Resolver Domain Appending

- Stub resolvers can be configured to retry failed queries by appending a list of domain names.
- Allows people to type "www" in URL box instead of wasting extra keystrokes on "www.foo.tld'
- An incident reported to OARC last year showed brute force attempts to resolve all letter combinations within a large number of CCTLDs. i.e.:

```
148.168.190.28.53: 38578 [1au] A? aaig.com.ao.
148.168.190.28.53: 61215 [1au] A? aaig.com.co.ao.
148.168.190.28.53: 43359 [1au] A? aaig.com.aq.
148.168.190.28.53: 8109 [1au] A? aaig.com.co.aq.
148.168.190.28.53: 44889 [1au] A? aaig.com.nl.
```

What can be done?

- Educate admins and users that search lists are a bad idea.

Application Domain Appending

- Web browsers like to append TLDs to "names" that don't resolve.
- Combined with Resolver domain appending, a single, non-existent domain name can result in $O(n^2)$ queries leaving the caching resolver.

What can be done?

- Disable name completion features in web browsers.

Operational Issues

Lame/Bad Delegations

- Bad delegations increase lookup latencies and network traffic.
- Different levels of badness (lame, non-existent, non-responsive)
- Strict definition says that you need an answer to declare lameness.
- Smart resolvers will remember and exclude lame and non-responsive delegations.

What can be done?

- Utilize zone checking tools to check your delegations.

Open Resolvers

- Exploited in source-spoofed DDoS attacks.
- Combined with implementation bugs, allows outsiders to attack/crash your nameserver.
  - See CVE-2006-4095, SIG processing in BIND
- More susceptible to cache poisoning.
- Same old problem as open SMTP relays and open proxies.

What can be done?

- Add ACLs to your nameserver configuration.

Lack of Network Diversity

- In our survey's, about 25% of zones have all nameservers on the same /24.
- Granted, many of those zones may be unimportant.
- A router config error took all Microsoft DNS servers offline for 24 hours in January 2001.

What can be done?

- Ask a college to secondary for you.
- Use one of the many free/cheap secondary services.

Missing SOA

- So-called "parked" domains often do not have SOA records.
- Owners are too lazy to create separate zone files. Use fake parent zones instead.
- Related to cache poisoning.
- Examples: *mogoos.com*, *nailten.com*, *jessica.com*

What can be done?

- Educate and encourage domain parkers to write correct zone files.

Complex Nameserver Dependencies

- According to "Perils of Transitive Trust in the Domain Name System," 46 separate nameservers are potentially involved in the resolution of names in the average zone.
- A compromise of only one of those dependent nameservers can result in hijacking, traffic redirection, etc.

What can be done?

- Use zone checking tools to understand dependencies for your own domain names.
- Deploy DNSSEC?

TTL Tradeoffs

- Lower TTLS $\Longrightarrow$ more flexibility, more traffic, less robust.
- Higher TTLS $\Longrightarrow$ less flexibility, less traffic, more robust.
- Experienced admins understand the TTL tradeoffs and can plan for changes well in advance.
- Some admins are caught off guard and may be forced to make a change when TTLs are high, resulting in service outages.
- Is the additional traffic from low TTLs really a problem for anyone?

What can be done?

- Plan in advance or live with the consequences.

Alternate/Expanded DNS Namespaces

- A number of "alternate" DNS namespaces already exist.
  - new.net
  - public-root.com
  - OpenNIC
  - cesidianroot.com
  - China
  - a number of others that have given up
- Motivated by money, desire for control, or fear/frustration of ICANN.
- RFC 2826: "To remain a global network, the Internet requires the existence of a globally unique public name space."

What can be done?

- Just say "No."
- Embrace the inevitable?

Chaos Class

- Many admins have been taught to fear *version.bind* and *hostname.bind* queries.
- Afraid that they will be attacked to exploit version-specific bugs.
- Many configure their nameservers to give out false or misleading answers.
- A determined attacker would probably just try all known exploits anyway.
- The *fpdns* fingerprinting software does a good job of reporting the approximate version.

What can be done?

- Realize that security through obscurity doesn't work.

NXDomain Interception

- Some ISPs change an NXDomain to an answer with A records pointing at a search engine with advertisements.
- Issues similar to sitefinder.

What can be done?

- "Opt out" by using third-party DNS resolvers.
- Run your own resolver locally.
- Deploy DNSSEC?

DNS As Load Balancer

- DNS is often used for load balancing.
  - DNS round-robin relies on resolvers to return RRs in random order.
  - GSLB gives out different answers depending on who asks.
- Usually requires low TTLs.

What can be done?

- Live with the consequences, if any.

Registry/Registrar Issues

Typo Squatters

- Money to be made capturing people who make typos
- Or mis-remember the domain name (*cough* www.ripe.org *cough*)
- In some cases, typos can lead to spyware/viruses.

What can be done?

- Trademark owners have to go after registrants of similar domains.
- Use bookmarks.
- Use spyware and virus protection.

TLD Wildcards

- A wildcard in a TLD zone confuses applications that expect NXDOMAIN.
  - See domain appending above!
- Sitefinder was focused on HTTP and required special hacks for SMTP, and other protocols?

What can be done?

- Use the *delegation-only* hacks in BIND.
- Deploy DNSSEC?

Domain Hijacking

- Possible to take over a domain name with a fake transfer request.
- *sex.com* (Oct 1995), *ebay.de* (Sept 2004), *panix.com* (Jan 2005), countless others.

What can be done?

- It seems like sufficient controls are already in place, yet this still happens.

Trademark Disputes

- DNS is tightly linked with trademark issues.
- Trademark owners have legal means to take DNS names from others.
- Trademark owners are encouraged to register their name in as many TLDs as possible.
- Choices for trademarks and company names are now determined by domain name availability.

What can be done?

- No idea

I'm Feeling Lucky

- Many people use Google rather than type in hard-to-remember domain names.
- And why not? Its safer and sometimes more convenient.
- Will Google make DNS obsolete?

What can be done?

- Use bookmarks?

References

1. Perils of Transitive Trust in the Domain Name System :: http://www.cs.cornell.edu/People/egs/papers/dnssurvey.pdf

2. DNS Cache Snooping :: http://www.sysvalue.com/papers/DNS-Cache-Snooping/

The End