I think we've finally gotta line on these DNS problems.

IBM Internet Security Systems

Toward a Robust Domain Naming System

Michael H. Warfield

mhw@linux.vnet.ibm.com

02/07/2011

This talk will be about...

- **What the Domain Naming System (DNS) is.**
- **How it works.**
- **How it is attacked.**
- **How to protect it.**
- **How to deploy it so it is robust.**

This talk...

- **Will be for**
  - Those who run nameservers.
  - Those who depend on it to work.
  - (Anyone not covered at this point?)
- **Will not be**
  - Specific to any particular nameserver.
  - About how to set up your own zones and domains.

DNS – The Domain Naming System

- **DNS is a ubiquitous core protocol.**
  - Highly distributed, almost amorphous.
  - Highly redundant
  - Nameservers scattered all over the Internet.
  - UDP based with TCP used for large transfers.
- **DNS manages the majority of Internet naming.**
  - A lookup service translating names to resources.
  - It is to the Internet what a phone book is to telephones.

The Domain Naming System

- **Supports multiple types of resources records (RR).**
  - A – IPv4 Addresses, AAAA – IPv6 Addresses
  - MX – Mail Exchanger
  - CNAME – Canonical Name (name aliases)
  - NS – Nameserver
  - TXT – Text Record (catch pan for a lot of cruft)
- **Complete failure would bring the Internet to a halt.**
  - “Complete” failure is virtually impossible.
  - Localized failures are common.
  - Most failures are preventable.

Domain Name Components

- **FQDN – Fully Qualified Domain Name**
  - Complete name specification from host to root
  - Left to right hierarchy of higher zones from host to root.
  - Top Level Domains (TLD's) are the rightmost element
- **gTLD – Global Top Level Domains**
  - .com, .org, .gov, .net, .edu, .mil, .info, .biz, etc, etc...
- **ccTLD – Country Code Top Level Domains**
  - .us, .cn, .uk, .au, .eu, .jp, .tv, .fm, etc, etc...
- **Domains may have multiple layers of subdomains**
  - www.ip6.wittsend.com, mail.wittsend.atl.ga.us

Zones and Domains

- **A zone refers to a region of authority.**
  - A zone is represented by a start of authority and set of authoritative nameservers responsible for it.
- **A domain is a designation within a name string.**
  - Domains and subdomains are containers for FQDN's.
  - Conventionally, a name without the left most host name.
- **Zones and domains are often largely 1:1.**
  - A domain may be delegated across multiple zones.
  - A domain and subdomains may be in a single zone.
- **Terms are loosely used interchangeably.**

Reverse Zones

- **Specialized reverse lookup zones.**
- **Reverse zones map addresses back to names.**
  - IPv4 – inaddr.arpa domain
    - Reverse octets: 130.205.32.64 -> 64.32.205.130.in-addr.arpa
  - IPv6 – ip6.arpa domain
    - Reverse hex digits (16 hex digits reversed).ip6.arpa
    - www.ip6.wittsend.com has IPv6 address 2001:4830:3000:2:204:8ff:fe00:1151
    - 1.5.1.1.0.0.e.f.f.f.8.0.4.0.2.0.2.0.0.0.0.0.0.3.0.3.8.4.1.0.0.2.ip6.arpa domain name pointer www.ip6.wittsend.com.
  - PTR records point back to system and domain names.
- **Some services will refuse connections to clients which have missing or misconfigured reverse records.**

DNS Nameservers

- **Nameservers are roughly classified into broad groups.**
- **Root nameservers handle the top of the DNS.**
- **Authoritative nameservers provide resource records.**
  - May be a master or one of one or more slaves
- **Caching nameservers handle queries and store results.**
  - Caching servers may be forwarding or fully recursive or a mix.
- **Resolvers originate requests for applications.**
- **A nameserver can be a combination of some of above.**
- **Often same application, different roles, client and server.**

The Root Nameservers

- **The root zone is “.” (the right hand edge of a name).**
- **13 nameservers are responsible for the root zone.**
- **Many of the 13 are now “anycast addresses”.**
- **Many now also have IPv6 addresses.**
- **The root nameservers do no recursion and no caching.**
- **Each anycast nameserver has multiple instances.**
- **In reality, there are currently 255 discrete root servers.**
- **Attacks are mitigated by their numbers and distribution.**
- **The root can take care of itself.**

Root Name Server Map: http://www.root-servers.org/map

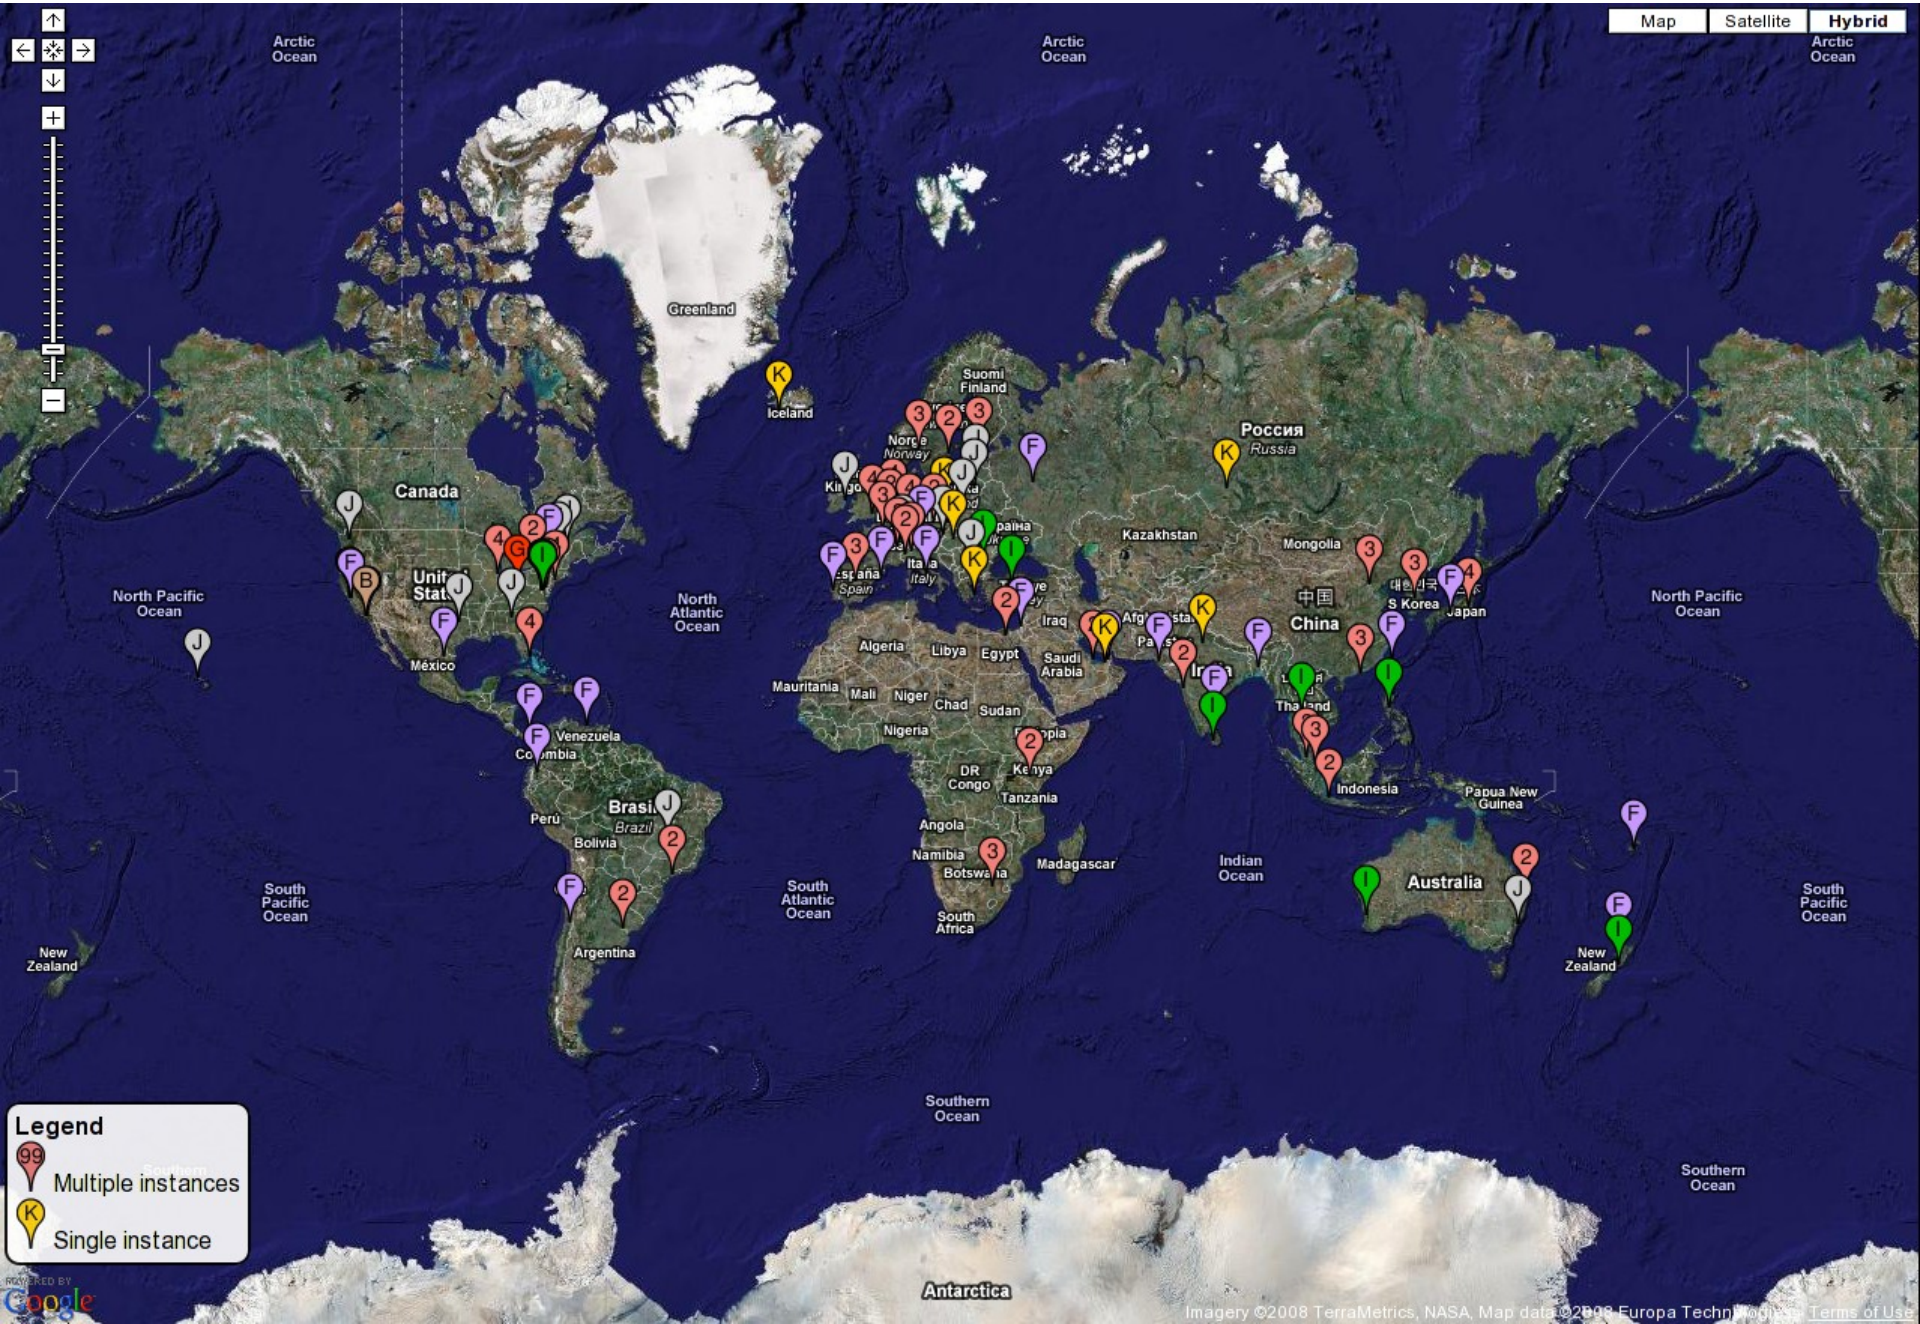

Authoritative Nameservers

- **Authoritative nameservers manage zones.**
- **Listed in a NS records in the root of that zone.**
- **Zones may be forward or reverse.**
- **One (or more) master(s) and multiple slaves.**
- **Slaves get zone data from the master or other slaves**
- **A compromised master impacts all the slaves.**
- **A compromised slave may be detected by verification.**
- **Any given server can be a slave for some zones and a master for others and they can support a mix of forward and reverse zones.**

Caching Nameservers

- **Respond to resolver requests.**
- **Attempt to obtain responses from authoritative servers.**
- **May obtain responses from other caching servers.**
- **May forward queries to other known servers.**
- **May perform full recursive queries from the root up.**
- **Cache results and glue (extra address information).**
- **Optimize multiple queries by referencing cache.**
- **Caches expire based on a “Time To Live” (ttl).**

Fully Recursive Nameservers

- **These do most of the work of looking up names.**
- **Just one type of caching nameserver.**
- **Query the root servers for the Top Level Domain (TLD) servers.**
  - Root nameservers send back only redirects to the TLD's.
- **Query the TLD nameservers for desired zones.**
  - TLD nameservers may recurse to subdomains.
  - TLD nameservers generally return additional redirects.

Forwarding Nameservers

- **Special class of caching nameserver.**
- **Forwards requests to designated nameservers.**
- **Does not “walk the tree” from the root up.**
- **Generally does not query authoritative nameservers directly.**
- **Relatively immune to caching attacks.**
- **Dependent on forwarding service for poison protection.**
- **Servers can be both forwarding and fully recursive.**

Resolvers

- **Handle queries from applications.**
- **Initiate queries into the DNS.**
- **Forward requests to the caching nameservers.**
- **Do not perform full DNS tree walks.**
- **Only query configured servers (resolv.conf)**
- **May or may not cache results.**
  - Some applications may use caching resolvers (Mozilla).
- **Not generally subject to cache poisoning.**
- **Are single host based (not really a nameserver per se).**

Views

- **Are a feature of certain authoritative DNS servers.**
- **May present a different set of zone records depending on the origin of the query.**
  - Internal and external views of zones can then differ.
- **Have been used to consolidate subdomains into a flat domain.**
- **Have been used to consolidate servers.**
- **Add complexity to configurations.**
- **Can make validation and verification more difficult.**
- **Should be used with caution.**

Name Server Best Common Practice

- **There are several Best Common Practices (BCPs).**
  - RFC 1912 – Common DNS Operational and Configuration errors
  - RFC 2182 – Selection and Operation of Secondary DNS Servers
  - The Secure BIND Template
- **DNS servers should be dispersed.**
  - Topologically
  - Geographically
- **DNS servers should be redundant.**
  - Multiple slave authoritative servers
  - Multiple caching servers

DNS Attacks

- **DNS comes under attack frequently.**
- **Originally not designed with security in mind.**
- **Attacks can be made against domains or clients.**
- **On rare occasion an implementation flaw is attacked.**
- **Attacks tend to be localized.**
- **Attacks against the root nameservers have been largely ineffective.**
- **DNSSec can help armor DNS from poisoning attacks.**
- **A lot can be done short of DNSsec.**

DNS Failures

- **Most DNS problems are failures, not attacks.**
- **Most failures are configuration and human error.**
  - Time to Live results in a major error to failure disconnect.
  - Some errors do not show up for hours.
  - Some errors can take days to correct.
- **Many DNS deployments fail to follow best common practices.**
- **Even large corporations often fail to deploy their DNS infrastructure securely or robustly.**

The Microsoft Incident

- **January 2001**
- **Entire microsoft.com domain falls off the net.**
- **Triggered by a minor firewall configuration error.**
- **Slaves were cut off from the master and timed out.**
- **Configuration error occurred hours before failure.**
- **Time To Live (ttl) kept slaves available for hours.**
- **All the public slaves were on the same subnet.**
- **After initial recovery, the common subnet router was then DDoS'ed for days by attackers having fun.**

The AT&T Incident

- **December 2007**
- **Major DSL provider loses their caching nameservers.**
- **Failure was reported to be due to a single failed router.**
- **Outage impacted large portion of Southeast US.**
- **Customers using the DSL nameservers were impacted.**
- **Customers running their own nameservers actually had improved service.**
  - (Where did everybody else go?)
- **Network connectivity not impacted.**

DDoS Attacks on the Root DNS Servers

- **There have been several attempts to DDoS the root.**
- **Initial attacks resulted in very limited success.**
  - Some localized outages reported.
- **Several root nameservers switched to anycast.**
  - Most of the root nameservers now use anycast addresses.
- **There are now hundreds of nameservers servicing the 13 root nameserver addresses.**
- **More servers are on the way.**
- **No doubt, more attacks are on the way as well.**
- **Some accusations of China abusing their root.**

Reflection Attacks

- **DNS may return multiple records for small requests.**
- **Responses may be much larger than queries.**
- **DNS is UDP based and can be easily be spoofed.**
- **Reflection attacks spoof queries to recursive servers.**
- **Recursive queries reference malicious servers.**
- **Malicious servers return massive responses.**
- **Larger distributed responses overwhelm targets.**
- **Basic resource amplification attack.**
- **Don't allow global access to recursive servers.**

DNS Cache Poisoning

- **DNS cache poisoning works by feeding false information into DNS server caches.**
- **Poisoning can be primary responses or glue records.**
- **Poisoned results remain until cache is flushed (ttl).**
- **Original attacks exploited predictable query id's (QID).**
- **Attacks tend to be localized against specific targets.**
- **ISP DNS servers can impact a broad client base.**
- **One attack fed false .com glue records hijacking .com!**
- **Modern servers highly randomize QID's.**

Kaminsky Attacks

- **Dan Kaminsky demonstrated a method of poisoning nameservers even with random QID's.**
- **QID's are only 16 bits.**
- **He optimized queries and spoofed responses.**
- **Took advantage of predictable query source ports.**
- **Old configurations uses 53/udp as the source port.**
- **Randomized source ports help mitigate this.**
- **Attacks are noisy with large numbers of packets.**
- **Some attacks still possible even with highly randomized query source ports.**

Firewalls and NAT and DNS

- **Firewalls and NAT devices can impact DNS.**
- **NAT devices may map ports to predicable ports.**
  - Exposes even patched hardened servers to Kaminsky attacks.
- **Firewalls may be misconfigured.**
  - Old static rules only allowed 53/udp <=> 53/udp or worse.
  - Misconfigured firewalls may force misconfigured DNS.
- **Stateful firewalls can detect and stop most DNS cache poisoning by triggering on off-port responses.**
  - Creative use of the iptables “recent” target can detect and block attacks attempting to guess DNS source ports.

Domain Hijacking

- **This is the act of taking over an entire domain.**
- **Generally requires registrar compromise or error.**
  - Social engineering is real popular for this!
  - Lurkers are waiting for domains to expire.
- **They change DNS servers designated as authoritative.**
- **They can “man in the middle” (MITM) your DNS.**
- **Impacts all accesses to a domain.**
- **Not much to be done at the DNS level.**
- **Keep domain registrations secure and up to date.**

DNS Covert Channels, Tunnels, and VPNs

- **Covert channels and tunnels do exist over DNS.**
- **Can be malware beacons and even covert VPNs.**
  - Beacons can be very low and slow and hard to spot.
  - Beacons can leak data and retrieve commands.
- **Can be a very simple tunnel with existing software.**
  - OpenVPN on port 53/udp – doesn't play well with cachers.
- **Can be very sophisticated specialized tunnels.**
  - DNScat - Works through cachers but provides no routing.
  - Iodyne – Full VPN including routing.
- **Most sites don't lock down open DNS access!**

Google, China, and Aurora

- **Goggle accused China of hacking their systems.**
- **The attack was an Advanced Persistent Threat, APT.**
- **Google logs and saves all DNS queries and responses!**
- **Aurora, the APT, was first detected in DNS logs!**
- **Google rolled back logs to determine “patient zero.”**
  - Compromised by spear phishing attack months earlier.
- **Rolling forward, they caught more versions of Aurora.**
- **All DNS requests had to go through Google servers.**
- **Datamining of DNS logs was vital to their response.**

Suspicious DNS traffic

- **Heuristics – Rules of thumb, not hard rules!**
  - External domains returning local or private addresses.
  - Queries for many hosts in a single external domain.
  - Queries for large host names in small domains.
  - Huge replies to simple queries.
  - Requests or replies to foreign addresses for no reason.
  - Very small ttls in responses.
  - Attempts to query DNS bypassing local cachers.
  - Correlations between any of the above!
- **Some can be legitimate with legitimate applications.**

Legitimate Beaconing

- **There are legitimate uses which look like beaconing.**
- **DNS is highly efficient and reliable for short fast queries.**
- **DNS often works even through closed WiFi application layer gateways.**
- **Example applications:**
  - Realtime Blackhole Lists, RBLs
  - Hashtable query lists
  - Security whitelists and blacklists
  - “LoJack” style location service for mobile devices?

DNSsec

- **DNSsec adds digital signatures to zones.**
- **Signatures are cryptographically secure.**
- **Signatures keys are hierarchical and registered.**
- **Several keys are required for each zone.**
- **Some keys are periodically regenerated.**
- **Zones must be updated and re-signed periodically.**
- **Secure DNS is non-trivial to deploy but not bad.**
  - OpenDNSsec helps this with a “bump on the wire” method.
- **DNSsec enables trustworthy keys and sigs in DNS.**

Recommendations – Authoritative Side

- **Split your DNS service functionality.**
- **Isolate your DNS service from other services.**
- **Distribute authoritative nameservers.**
- **Do not expose master nameservers to the outside.**
- **Restrict zone transfers.**
- **Keep master / slave configurations coherent and simple!**
- **Set up an independent validation server.**
- **Run frequent verifications (slave/slave & master/slave).**
- **Monitor your registrations for tampering.**

Recommendations – Client Side

- **Restrict outside access to caching nameservers.**
- **Place caching nameservers behind stateful firewalls.**
- **Test and upgrade NAT devices and properly configure.**
- **Fix misconfigured firewalls.**
- **Consider a secure forwarding service like OpenDNS.**
- **Test any outside servers you rely on or forward to.**
- **Run your own caching nameservers.**
- **Distribute caching nameservers for performance and redundancy.**

General Recommendations

- **Enable logging where possible.**
- **Log to remote servers.**
- **Capture DNS traffic and archive it.**
- **Mine that data for suspicious behavior.**
- **Deploy and support DNSsec where possible.**
- **Keep DNS server software up to date!**
- **Marshall your DNS**
  - Block outbound 53/udp and investigate attempts to bypass.
  - Use only designated DNS servers inside security perimeters.

Conclusion

The Domain Name Service is a fundamental core protocol upon which the entire Internet depends. DNS is an old protocol but is exceptionally resilient.

There are many ways to enhance the resiliency and robustness of a DNS deployment. The result is a stable and highly maintainable DNS that is resistant against attack.

References

- DNS Best Practice Resources
  - http://www.infoblox.com/library/dns_resources.cfm
- Secure BIND Template
  - http://www.cymru.com/Documents/secure-bind-template.html
- RFC 1912: “Common DNS Operational and Configuration Errors”
  - ftp://ftp.rfc-editor.org/in-notes/rfc1912.txt
- RFC 2182, “Selection and Operation of Secondary DNS Servers”
  - ftp://ftp.rfc-editor.org/in-notes/rfc2182.txt
- Microsoft Web site outages highlight DNS as single point of failure
  - http://www.infoworld.com/articles/hn/xml/01/01/26/010126hndnsfailure.html
- DSL outage hits AT&T in Southeast
  - http://www.cnn.com/2007/TECH/12/04/att.outage.ap/index.html
- OpenDNSsec:
  - http://www.opendnssec.org

Thank you very much!

Thank you very much for attending. These slides are available in several forms at the following URLs:

http://www.wittsend.com/mhw/2011/RobustDNS.odp

http://www.wittsend.com/mhw/2011/RobustDNS.pdf

http://www.wittsend.com/mhw/2011/RobustDNS.ppt

Special thanks to:

Introduction Comic: Chad Carpenter's Tundra

http://www.tundracomics.com

IBM Internet Security Systems

Toward a Robust Domain Naming System

Michael H. Warfield

mhw@linux.vnet.ibm.com

02/07/2011