

Threats to the Domain Name System

Cricket Liu, Infoblox

CLUE Meeting

Denver, Colorado

May 13, 2003

The Internet's Domain Name System

- The world's largest distributed database
- Comprises a hierarchical namespace of independently administered units called zones
 - Zones contain data in the form of resource records
 - Domain names are indexes to resource records
- Resource records map domain names to
 - Addresses
 - Mail destinations
 - More

The Internet's Domain Name System

- Name servers
 - Answer queries for records in zones they're authoritative for
 - Query other name servers for records in zones they're not authoritative for
- Resolvers
 - Act as intermediaries between applications and name servers
 - Translate applications' requests for resource records into DNS messages
 - Interpret DNS messages into data structures for applications

Why is DNS Important?

- Most Internet-based services rely on DNS
 - The World Wide Web
 - To let browsers map domain names in URLs to addresses
 - To let web servers map the addresses of browsers back to domain names
 - Electronic mail
 - To let mail user agents map domain names of mail servers to addresses
 - To let mail transport agents map domain names in email addresses to the names of mail servers
 - To let mail transport agents map addresses of sending mail servers to domain names
 - Telnet, FTP, instant messaging, streaming media
 - To let clients map the domain names of servers to addresses

DNS Infrastructure

- Thirteen root name servers
 - Spread around the world (sort of)
 - Each processing thousands of queries per second
 - Referring queriers to the name servers serving a particular top-level zone
 - For example, *com*, *net*, *org*, *au*, *uk*
- Thirteen authoritative name servers for *com* and *net*
- Spread around the world
 - Each processing thousands of queries per second
 - Referring queriers to the name servers serving subzones of *com* and *net*
 - For example, *nxdomain.com*, *infoblox.com*, *npr.org*
- Thousands of other name servers
 - Run by ISPs, companies, organizations, and individuals

The Root and *com/net* Name Servers

- 13 root name servers
 - North America (10)
 - Washington, D.C. area (6)
 - California (4)
 - Europe (2)
 - London
 - Stockholm
 - Asia (1)
 - Tokyo
- 13 *com/net* name servers
 - North America (8)
 - California (3)
 - Washington, D.C. area (2)
 - Atlanta
 - Miami
 - Seattle
 - Europe (3)
 - Amsterdam
 - Stockholm
 - London
 - Asia (2)
 - Hong Kong
 - Tokyo

Threats

- Security
- Scaling
- Misconfiguration and Single Points of Failure
- Attacks
- Politics
- Alternate roots
- Education

Security

- The Domain Name System wasn't designed with many security mechanisms
 - Most have been added only relatively recently
 - Transaction Signatures, or TSIG: RFC 2845, published May 2000
 - Some have been developed but haven't been widely implemented
 - The DNS Security Extensions, or DNSSEC: RFC 2535, published March 1999, *et al.*
 - Only partial support in BIND 8.3.4, the latest BIND 8 release
 - Support in BIND 9.2.2, the latest BIND 9 release
 - » Which itself isn't widely deployed yet
 - Almost no support in the Microsoft DNS Server

The DNS Security Extensions

- DNSSEC applies public key cryptography to DNS to allow maintainers of zones to “sign” zone data
 - Queriers can verify the authenticity of the data
 - Provides protection against spoofing
- But DNSSEC is problematic because
 - Signing a zone increases its size 5-7x
 - Maintaining a signed zone with current tools is cumbersome
 - Verifying signed resource records is computationally intensive
- This has slowed adoption of DNSSEC

The Status Quo

- In the mean time, administrators make do with few security mechanisms
 - Many administrators don't fully understand or use those that are available

Vulnerabilities

- Unfortunately, BIND name servers have also been the source of many vulnerabilities
 - Some of these have been quite severe, leading to root compromise

A Matrix of BIND's Vulnerabilities

version	zxf	sigdiv0	srv	nxt	sig	naptr	maxname	solinger	fdmax	complain	infoleak	tsig	libbind	openssl
4.8											+		+	
4.8.1							-						+	
4.8.2.1							-						+	
4.8.3							-						+	
4.9.3										+			+	
4.9.4										+			+	
4.9.4 p1										+			+	
4.9.5			-		+	+	+						+	
4.9.5 p1			-		+	+	+			+			+	
4.9.6			-		+	+	+			+			+	
4.9.7			-		-	+	+			+			+	
4.9.8			-		-	+	+			-			+	
4.9.9			-		-	+	+			-			-	
8.1			-		+	+	+	+	+	-			+	
8.1.1			-		+	+	+	+	+	-			+	
8.1.2			-		-	+	+	+	+	-			+	
8.2	-	+	+	+	+	+	+	+	+	-		+	+	
8.2 p1	-	+	+	+	+	+	+	+	+	-		+	+	
8.2.1	-	+	+	+	+	+	+	+	+	-		+	+	
8.2.2	+	+	+	-	-	+	+	-	-	-		+	+	
8.2.2 p1	+	+	+	-	-	+	+	-	-	-		+	+	
8.2.2 p2	+	+	+	-	-	-	-	-	-	-		+	+	
8.2.2 p3	+	+	+	-	-	-	-	-	-	-		+	+	
8.2.2 p4	+	+	+	-	-	-	-	-	-	-		+	+	
8.2.2 p5	+	+	+	-	-	-	-	-	-	-		+	+	
8.2.2 p6	+	-	+	-	-	-	-	-	-	-		+	+	
8.2.2 p7	-	-	-	-	-	-	-	-	-	-		+	+	
8.2.3	-	-	-	-	-	-	-	-	-	-		-	-	
8.2.4	-	-	-	-	-	-	-	-	-	-		-	-	
8.2.5	-	-	-	-	-	-	-	-	-	-		-	-	
8.2.6	-	-	-	-	-	-	-	-	-	-		-	-	
8.3.0	-	-	-	-	-	-	-	-	-	-		-	-	
8.3.1	-	-	-	-	-	-	-	-	-	-		-	-	
8.3.2	-	-	-	-	-	-	-	-	-	-		-	-	
8.3.3	-	-	-	-	-	-	-	-	-	-		-	-	
9.0.0	-	-	-	-	-	-	-	-	-	-		-	-	
9.0.1	-	-	-	-	-	-	-	-	-	-		-	-	
9.1.0	-	-	-	-	-	-	-	-	-	-		-	-	+
9.1.1	-	-	-	-	-	-	-	-	-	-		-	-	+
9.1.2	-	-	-	-	-	-	-	-	-	-		-	-	+
9.1.3	-	-	-	-	-	-	-	-	-	-		-	-	+
9.2.0	-	-	-	-	-	-	-	-	-	-		-	-	#
9.2.1	-	-	-	-	-	-	-	-	-	-		-	-	#

Vulnerable: '+', Not Vulnerable: '-', Feature does not exist: '#', Vulnerable (when linked with vulnerable library): '#'

Scaling

- DNS has scaled remarkably well
 - The design of DNS dates back to the early 80's
 - The design hasn't changed fundamentally since then
 - One of the design goals was to scale to accommodate the entire IP(v4) address space
 - And it does

A Perspective, Circa 1981

- “In the long run, it will not be practicable for every internet host to include all internet hosts in its name-address tables. Even now, with *over four hundred names and nicknames* [emphasis added] in the combined ARPANET-DCNET tables, this has become awkward.”

Source: David Mills, RFC 799

New Technologies Accelerate Growth

- IPv6
 - 2^{128} possible addresses
 - More address records
 - *Much* larger reverse-mapping domain
- ENUM
 - Using DNS as a global, integrated directory, mapping E.164 (global telephone) numbers to URLs (giving phone numbers, email addresses, etc.)
 - Potentially billions of domain names
- GPRS/3G wireless
 - Every mobile may have an IP address
- DNSSEC
 - Multiplying the size of existing zones

New Technologies Accelerate Growth

- IDN
 - Domain names get longer and (more?) opaque
 - More domain name registration and delegation

Misconfiguration/Poor Implementation

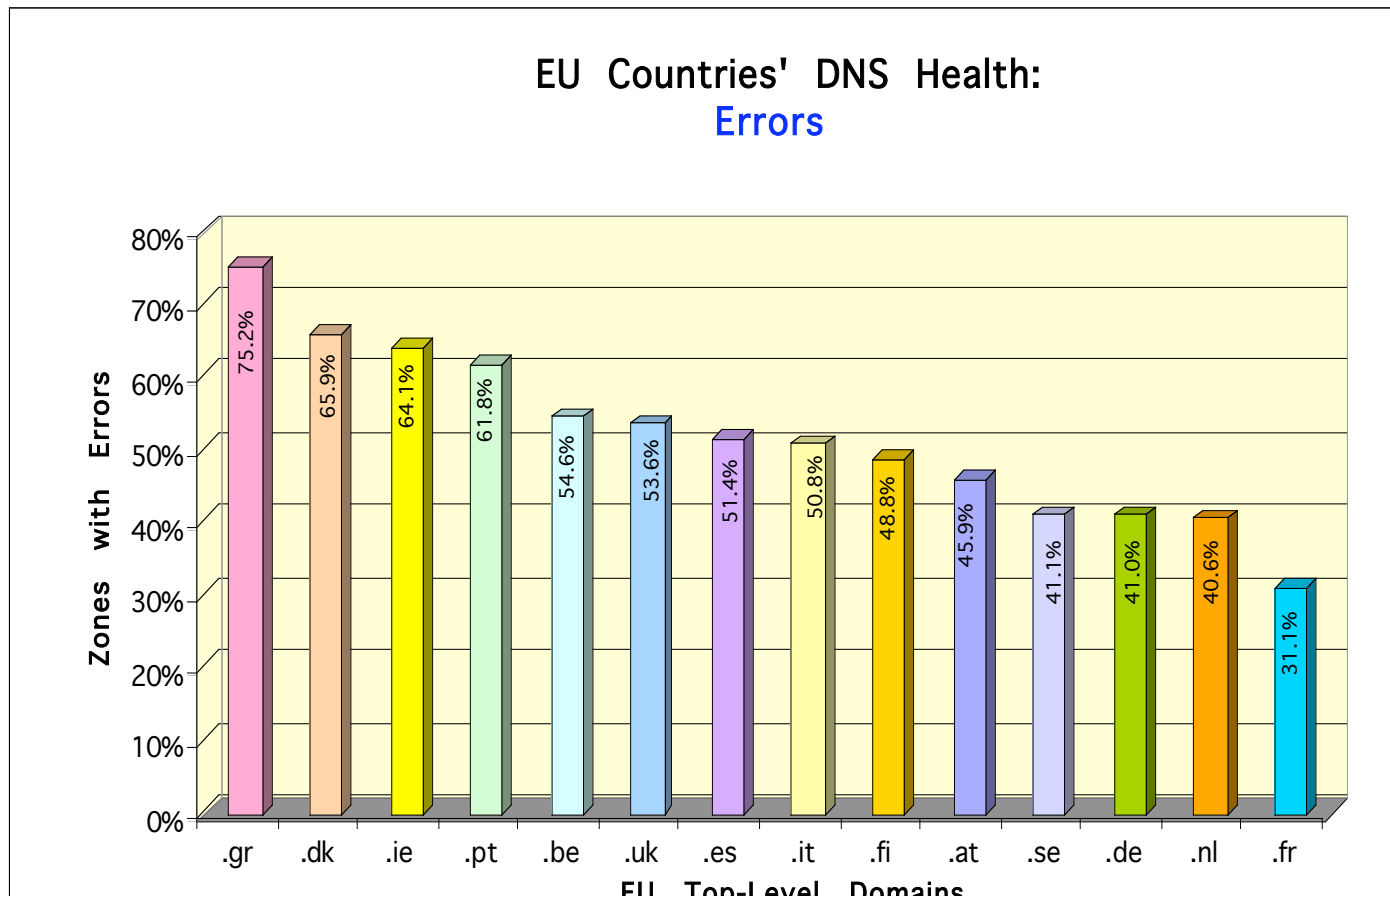
- These exacerbate problems caused by organic growth and the introduction of new technologies
 - Misconfiguration
 - Lame delegation
 - Mismatched name server information
 - Allowing RFC 1918 queries onto the Internet
 - Sending queries for the address of an address
 - For example, the address of the “domain name” 192.168.0.1
 - Sending queries for domain names that end in non-existent top-level domain names
 - Poor implementations
 - Repeatedly sending the same query
 - Not understanding or accepting certain errors or referrals

Men & Mice Study of EU TLDs

- Conducted November 20-21, 2001
- 2500 randomly selected zones
- Errors detected
 - Lame delegations: 23.9%
 - Unresponsive authoritative name servers: 21.3%
 - Mismatch of delegation data and zone data: 14.2%

Source: http://www.menandmice.com/6000/6350_eu_survey.html

Errors by Country




Single Points of Failure

- The Men & Mice EU TLD study also detected single points of failure
 - In particular, all authoritative name servers on the same subnet

Source: http://www.menandmice.com/6000/6350_eu_survey.html

Single Points of Failure by Country



Men & Mice Study of *com*

- Conducted August 2002
- 5000 randomly selected zones
- Errors detected
 - Single point of failure: 27.6%
 - Lame delegations: 20.2%
 - Unresponsive authoritative name servers: 17.94%
 - Mismatch of delegation data and zone data: 14.98%

Source: http://www.menandmice.com/6000/61_recent_survey.html

RFC 1918 Queries

- Queries from or about RFC 1918 networks
 - Neither should ever make it to the Internet
 - “Because private addresses have no global meaning, routing information about private networks shall not be propagated on inter-enterprise links, and packets with private source or destination addresses should not be forwarded across such links.”
 - “Indirect references to such addresses should be contained within the enterprise. Prominent examples of such references are DNS Resource Records and other information referring to internal private addresses. In particular, Internet service providers should take measures to prevent such leakage.”
 - Nor should dynamic updates in RFC 1918 reverse-mapping zones

Source: RFC 1918



Oops

- On the Internet, the reverse-mapping zones that correspond to RFC 1918 address space are delegated to two name servers, *blackhole-1.iana.org* and *blackhole-2.iana.org*
- According to Bill Manning, who runs them, *one* of those name servers received 120,000 queries per second during the latest SirCam outbreak
- According to CAIDA studies, one of these name servers received 51.4M dynamic updates in 86.5 hours
 - 10,000/minute
 - 165/second

Sources: Posting to USENET newsgroup comp.protocols.dns.bind,
<http://www.caida.org/outreach/presentations/2002/nanog0210/duane.pdf>

“DNS Damage”

- Research by CAIDA (Brownlee, claffy and Nemeth, and later by Wessels), monitoring the mix of queries received by *f.root-servers.net*, showed enormous numbers of repeated queries and other useless traffic



Source: <http://www.caida.org/outreach/presentations/2002/nanog0210/duane.pdf>

Other Garbage In

- 220 IP addresses represented 50% of the queries
- 15% of the (24 hour!) sample's queries came from one company
 - 63.4% of those were for A RRs for *[a-m].root-servers.net*
 - 14.6 of those were for the company's own domain names
- 2.22% of queries were recursive
 - 6.26% of queriers

Source: <http://www.caida.org/outreach/presentations/2002/nanog0210/duane.pdf>

The Net Effect

- The root name servers process many, many times more traffic than they need to, strictly speaking
 - For *f.root-servers.net*, over the sample period, 50x!

Why Not Add More Name Servers?

- The maximum size of a DNS message over UDP is 512 bytes
 - Guess how many name server records and corresponding address records fit into 512 bytes

Working Around the 512-byte Limit

- EDNS0
 - Allows queriers to “advertise” the ability to accept a larger, UDP-based DNS message
 - Supported by newer BIND 8 and all BIND 9 name servers
 - See RFC 2671
- Shared anycast
 - Using the same IP address for multiple, physical root name servers
 - The ISC and APNIC may begin using this scheme with *f.root-servers.net*
 - See <http://www.ietf.org/internet-drafts/draft-ietf-dnsop-ohta-shared-root-server-02.txt>

Attacks Against Name Servers

- DDoS attack against the root name servers, October 2002
 - A flood of ICMP traffic designed to deny root name service
 - Successful against as many as seven of the 13 root name servers
- Exploitation of TSIG buffer overrun by li0n worm, March-April 2001
- DDoS attack against Microsoft's name servers, January 2001
 - A follow-up to Microsoft's router misconfiguration mishap
- Kashpureff's *www.internic.net* cache poisoning attack, July 1997
 - Exploiting name servers that provided recursion to anyone

Politics

- Nobody seems to like ICANN
 - Not representative
 - Too political
 - Too opaque
- Some TLDs are beginning to question ICANN's sovereignty
 - Or at least their funding of ICANN
- Many notables would greatly circumscribe ICANN's role or dismantle it completely
 - Randy Bush
 - Dave Farber
 - John Gilmore
 - Peter Neumann

A World Without ICANN?

- Without a popularly recognized body to coordinate (not run or rule) the Internet's namespace, balkanization threatens
 - This way lies madness

Alternate Roots

- Several groups have set up alternate sets of root name servers
 - For example, the AlterNIC, OpenNIC, Open Root Server Confederation
- These load a root zone that delegates to a superset of the ICANN-recognized TLDs
 - For example, *.faq*, *.geek*, *.porn*
- Should these gain widespread acceptance, users of a particular set of root name servers will be able to resolve a different set of domain names than users of another set

Root Name Server Quality Control

- “Yes, some of those machines you saw stacked up in the bedroom were Alternic root servers.”

Source: Eugene Kashpureff, interview published in “urly indicator,” Alternic.com, 8/18/2000

Education

- When I wrote *DNS and BIND* with my friend Paul in the early 90's, DNS was a black art
 - *And it still is*
- There aren't enough people who understand DNS thoroughly
 - Hostmasters and postmasters
 - System and network administrators
 - Theorists, strategists and pundits

“Only You”

- To minimize the burden you place on “high-level” name servers (such as the roots)
 - If you use RFC 1918 address space, set up the corresponding zones on your name servers
 - Make sure your firewalls allow DNS messages from port 53 to high-numbered ports on your name servers
 - Or you won’t get responses
 - If you use Active Directory or Windows 2000/Windows XP’s network registration features, make sure that your dynamic update and query traffic remains local
 - Your name servers must be authoritative for a zone with the same name as the name of your Active Directory domain

“Only You”

- To minimize the risk of your becoming a victim
 - Eliminate single points of failure in your DNS infrastructure
 - Review your name servers’ configurations and the contents of your zones
 - Use publicly available tools such as *dnswalk*
 - <http://www.visi.com/~barr/dnswalk/>
 - Make sure your name servers are authoritative for the reverse-mapping zones that correspond to all of your networks
 - Get educated
 - Buy one of the many good books on DNS
 - Take a class on DNS