DNS/DNSSEC Workshop

Anycast DNS

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

The problem

- Only 13 nameservers in DNS design
  - Design constraint of original spec as results fit in a single UDP packet in less than 512 bytes
- Changing the DNS system to fix this is not feasible
- Internet politics at work as well
  - Many countries felt left out as most root servers in USA

The solution

- Use the routing system to deploy more servers
- BGP4 is a Path Vector protocol
  - Uses the concept of an Autonomous System to make decisions about the best way to route packets
  - Shortest AS path wins

Routing relativity

- We can look at how the best route to Victoria University of Wellington is seen in different parts of the Internet
  - At each of these locations the route selected as 'best' has a very different AS path
- https://stat.ripe.net/widget/looking-glass
- Enter the IPv4 address: 130.195.0.0
- Repeat with: 2404:2000::/32

Anycast details

- At various locations we deploy DNS servers that serve up duplicate content and advertise their address into the global routing system
- Local traffic is kept local

Routing relativity again

- We can look at how the best route to l.root-servers.net is seen in different parts of the Internet
  - At each of these locations the route selected as 'best' has a very different AS path
- https://stat.ripe.net/widget/looking-glass
- Enter the IPv4 block: 199.7.83.42
- Repeat with: 2001:500:3::42/32

Root Server deployment

http://www.root-servers.org/

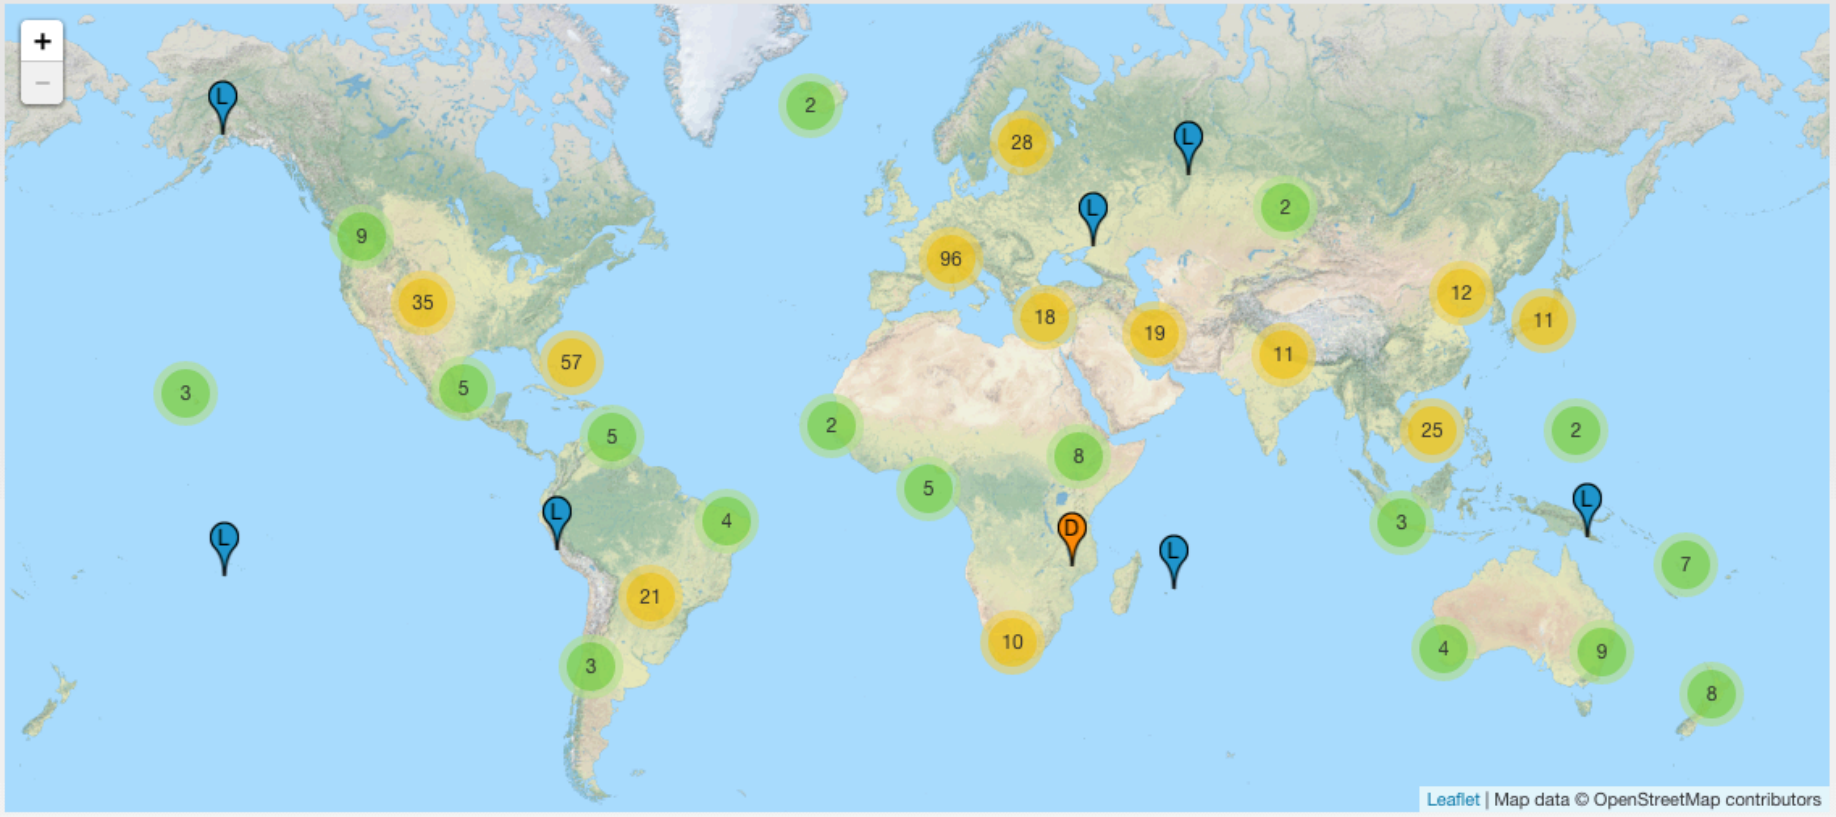

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Which l-root node am I using?

- traceroute L.ROOT-SERVERS.NET
  - The route your packets take to reach the nearest L root nameserver
- dig +norec @L.ROOT-SERVERS.NET HOSTNAME.BIND CHAOS TXT
  - Sends a query to the L root nameserver, asking it to reveal the name of the particular server which answers the query.

http://erg.wand.net.nz/amp/matrix.php/ipv4/latency/NZ/root+DNS

Source:	a.root-servers.net	b.root-servers.net	c.root-servers.net	d.root-servers.net	e.root-servers.net	f.root-servers.net	g.root-servers.net	h.root-servers.net	i.root-servers.net	j.root-servers.net	k.root-servers.net	l.root-servers.net	m.root-servers.net
auckland	192	126	133	195	11	139 ↓	186	212	11	13	274	2	193
catalyst	147	147	150 ↓	210	2	52	194	211	2	3	177	8	238
citylink	147	147	151	209	2	51 ↓	204	214	2	2	222	8	238
csotago													
fx-aknnr	135	132	139	224	13	41	191	198	13	12	210	0.8	189
inspire	143	144					196	207	3	11 ↓	218	5 ↑	191
massey-pn	277	136					210	221	3	11	247	10 ↑	247
maxnet	135 ↓	136					191	203	230	14	210	2 ↓	232
netspace	148	147	155	244	2	172	205	213	1	2	246	8	199
ns2b-digiweb	283 ↓	189	168	219	7	54	217	211 ↓	6	7	258	1	231 ↓
ns3a-avalon	146 ↓	143	152	220	1	172 ↓	204	212	1	2	244	8	195
ns3b-iconz	135	136	139	198	2	41	191	201	10	12	210	1	231
ns4a-orcon	141	132	142	206	12	182	187	237	12 ↑	12	186	17	183 ↓
rurallink	138	138	142	200	15	44	194	204	13	15	213	3	192 ↓
vuw													
waikato	139	168 ↑	143	197 ↓	4	45	193	210	14	14	214	4	196
wxc-akl	135	176 ↑	139	198	1	41	191	200 ↓	13	12	210	0.3	231

fx-aknnr to b.root-servers.net

	1 Hour (average)	24 Hour (average)	7 Day (average)
Latency (ms)	132	135	135
Packet Loss (%)	0	0	0

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

PCH
Packet Clearing House

Unicast Attack Effects

Traditional unicast server deployment...

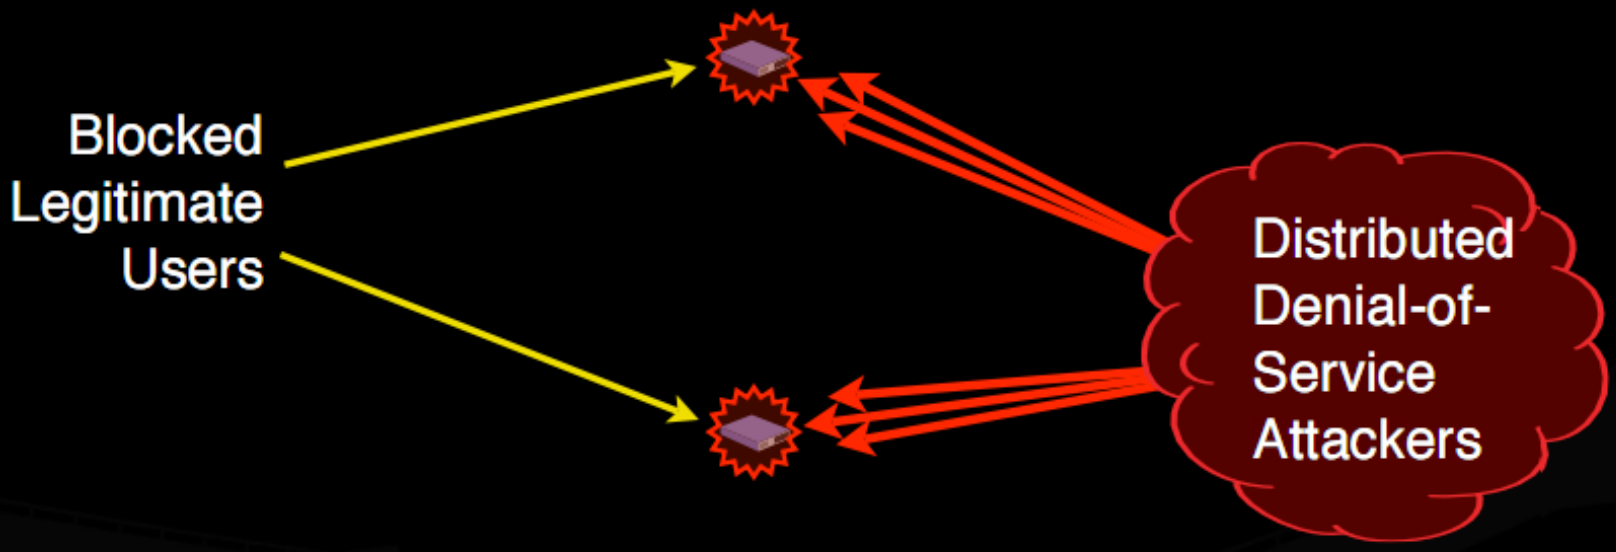

...exposes all servers to all attackers,
leaving no resources for legitimate users.

PCH
Packet Clearing House

Anycast Attack Mitigation

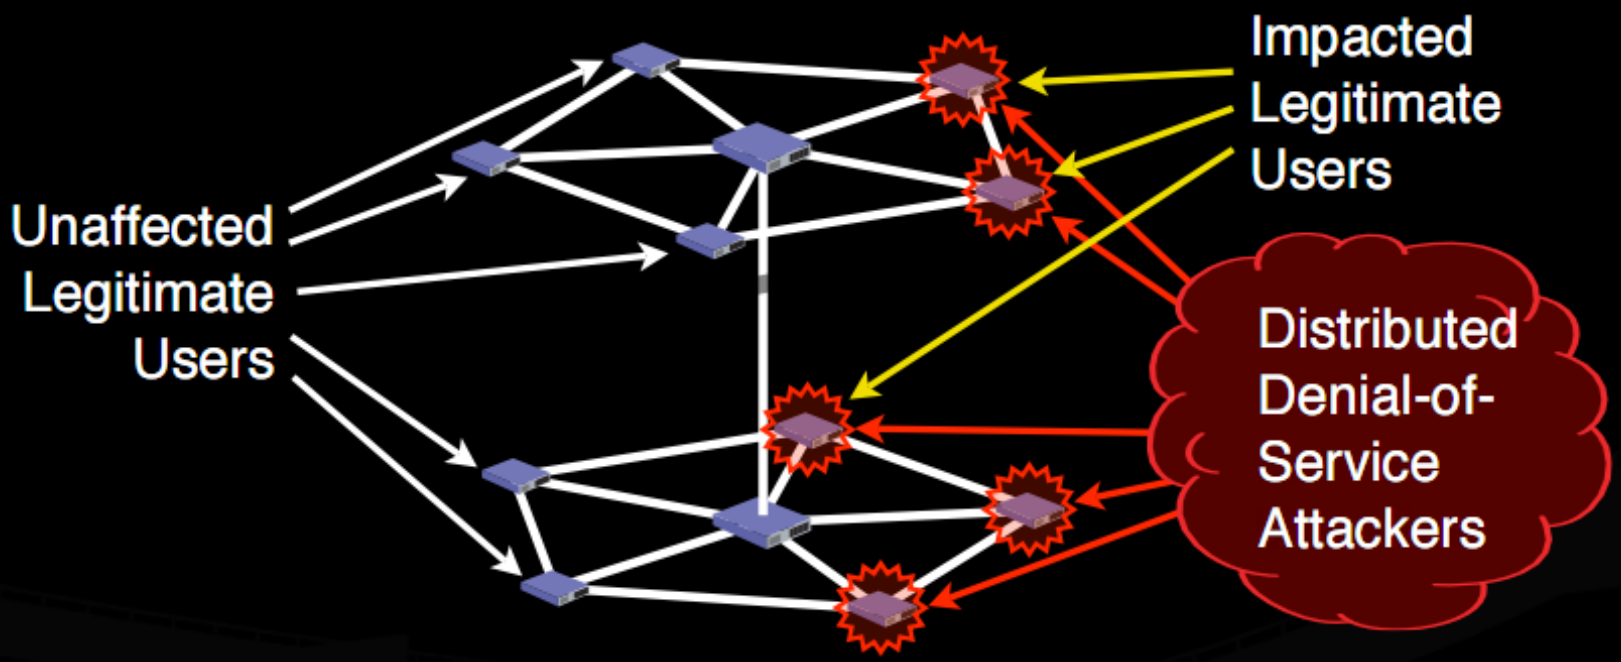

The benefits

- Reduce single points of failure in name server instances
- Faster response
- Reduce DDOS impact on DNS service
- Simple scaling for users

The trade offs

- More complexity in management of servers
- Need out of band channel to update information and gather statistics
- Widely deployed as benefits seem to outweigh downsides

Some reading

- http://en.wikipedia.org/wiki/Anycast
- f-root
  - https://www.isc.org/f-root/f-root-resources/
- l-root
  - http://www.dns.icann.org/

Further reading

- http://ftp.isc.org/isc/pubs/tn/isc-tn-2003-1.txt
- http://www.pch.net/resources/papers//anycast/Anycast-v07.pdf
- http://www.sanog.org/resources/sanog8/sanog8-dns-service-architecture-gaurab.pdf
- http://www.ietf.org/rfc/rfc3258.txt