Root Hints or Root Slaves

David Malone <David.Malone@nuim.ie>

25th July 2005

The idea

Better to take a zone transfer rather than query the root?

How could this be?

- Trade one big query against many little ones.
- Many bad queries get to root servers.
- Negative caching is good, but knowing all entries is better.

Debate on FreeBSD mailing list — decided to find out.

Hint configuration

```
zone "." {
        type hint;
        file "named.root";
};
```

Slave configuration

```
zone "." {
        type slave;
        file "s/root";
        masters {
                128.9.0.107;
                192.33.4.12;
                192.5.5.241;
        };
        notify no;
};
```

Method

- Busy server and quiet server.
- Restart BIND and run for week.
- Record queries with tcpdump.
- Study with some perl scripts.

Went more-or-less according to plan.

Results

bytes

packets

B

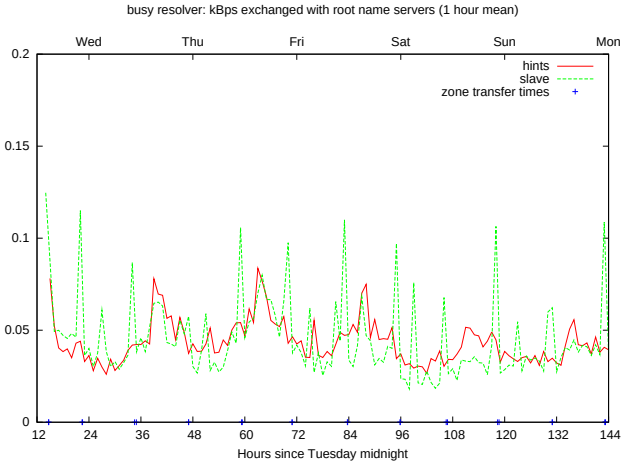

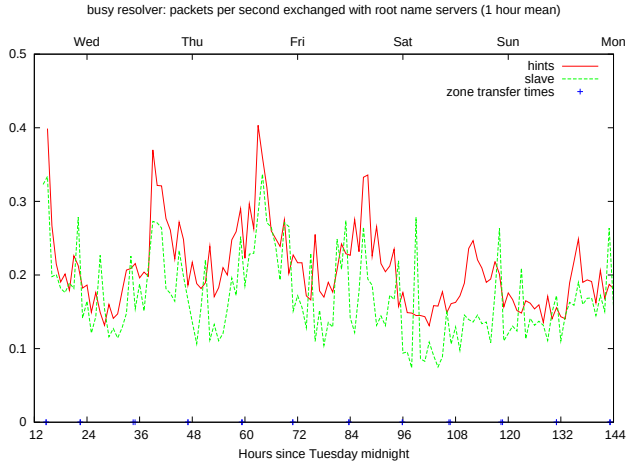

Q

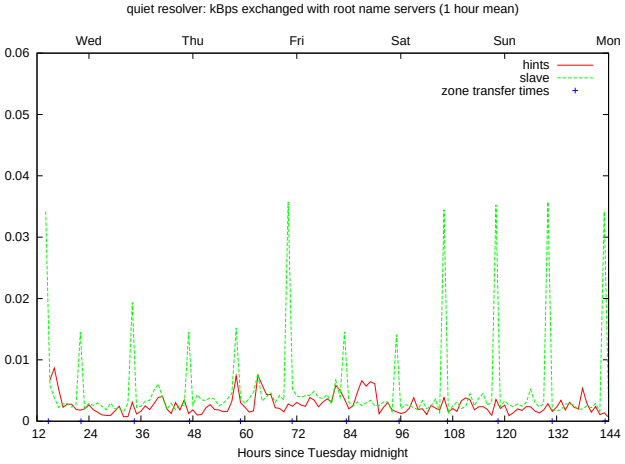

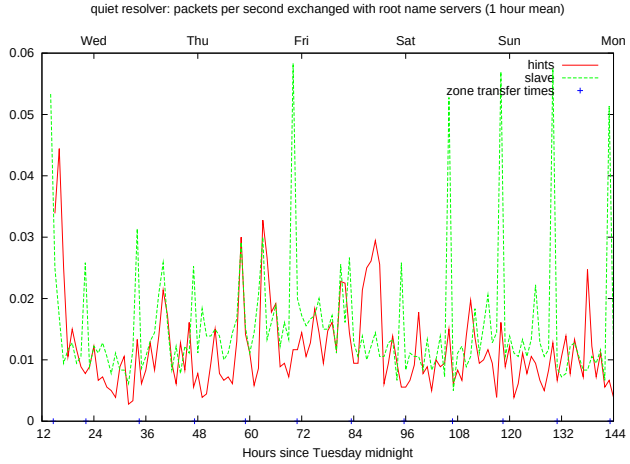

Volume Summary

	Quiet	
	Packets	Bytes
both	67%	152%
in	130%	207%
out	45%	43%

	Busy	
	Packets	Bytes
both	82%	103%
in	82%	106%
out	82%	85%

Quiet does not make enough queries to see a saving.

Breakdown (Busy)

	Queries			NXDomain		FormErr	
	hints	slave	ratio	hints	slave	hints	slave
A	11206	9158		1405	1	31	4
A....	315			314			
A6	31460	29418		33		24	15
A6....	268			265			
AAAA	720	157		555			
AAAA....	53			53			
ANY	2			2			
AXFR		26					
MX	558			492	12		
MX....	115	13		115			
MX[]	7899			7895		4	
NS	5					2	
PTR	660	647		7	9	1	3
SOA		1943					
SRV	24			24			
total	53285	41362	77%	11160	22	62	22

Wait time in seconds

Crude measure of wait time: $\sum \min(\text{response time}, 5s)$.

all types		without SOA
quiet hints	**quiet slave**	**quiet slave**
103.1	205.7	56.8
busy hints	**busy slave**	**busy slave**
3700.3	2099.0	1806.4

Shouting helps?

```
13:12:12.858858 A? ns1.math.ubc.ca.
13:12:12.858921 A6? ns1.math.ubc.ca.
13:12:12.858981 A? ns2.math.ubc.ca.
13:12:12.859037 A6? ns2.math.ubc.ca.
13:12:13.119553 A6? DNS2.UCLA.EDU.
13:12:13.119988 A6? DNS3.UCLA.EDU.
13:12:13.120050 A6? ns0.domainregistry.ie.
13:12:13.120160 A6? RAIN.PSG.COM.
13:12:13.120379 A6? MERAPI.SWITCH.ch.
13:12:13.120435 A6? ARGUS.MORE.NET.
13:12:13.120765 A6? ns2.ucd.ie.
```

Why ask the root?

```
14:29:30.932372  [1au] A6? ns1.tigermail.com.
14:31:53.827204  [1au] A6? duende.cafenet.com.bo.
14:38:04.385035  [1au] A? ns2.att.net.co.
14:38:04.385865  [1au] A6? ns2.att.net.co.
14:38:04.386717  [1au] A? ns1.att.net.co.
14:38:04.387544  [1au] A6? ns1.att.net.co.
14:38:06.398633  [1au] A6? ns1.att.net.co.
14:46:19.623488  [1au] A? SEC02.ns.esat.net.
14:46:19.624328  [1au] A6? SEC02.ns.esat.net.
14:46:19.625332  [1au] A? AUTH02.ns.esat.net.
14:46:19.626193  [1au] A6? AUTH02.ns.esat.net.
```

Multiple SOAs

```
15:01:57.413132 > 128.9.0.107.53:  SOA? .
15:01:57.592471 < 128.9.0.107.53:  1/13/13 SOA
15:01:57.923128 > 192.33.4.12.53:  SOA? .
15:01:58.133155 < 192.33.4.12.53:  1/13/13 SOA
15:01:58.133914 > 192.5.5.241.53:  SOA? .
15:01:58.282999 < 192.5.5.241.53:  1/13/13 SOA
```

Pros

- Eliminates most bogus queries.
- Fewer queries, less waiting.
- May shorten waiting times.
- No limit on response size.
- More robust?
- Root zone can influence traffic with SOA.

Cons

- Traffic volume benefit not clear.
- TCP state a burden on servers?
- Bad TCP/anycast interaction?
- Hard to generalise.

BIND

- Why all the A6 queries?
- SOAs to closest master?
- Why queries outside ., arpa, in-addr.arpa, mil and root-servers.net?
- Repeated query problem.
- Automagic possible?

Thanks