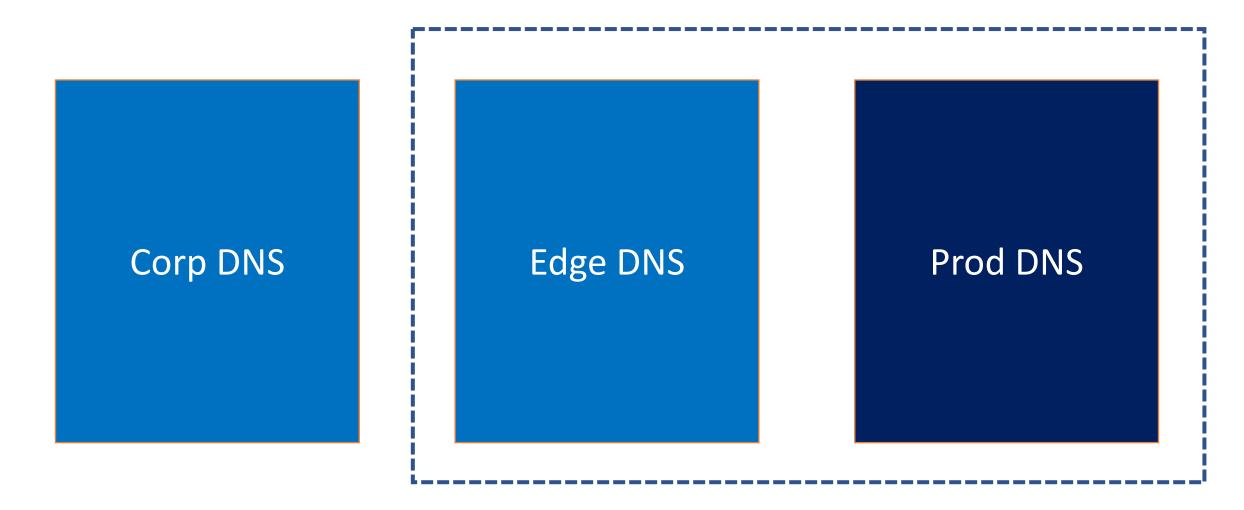




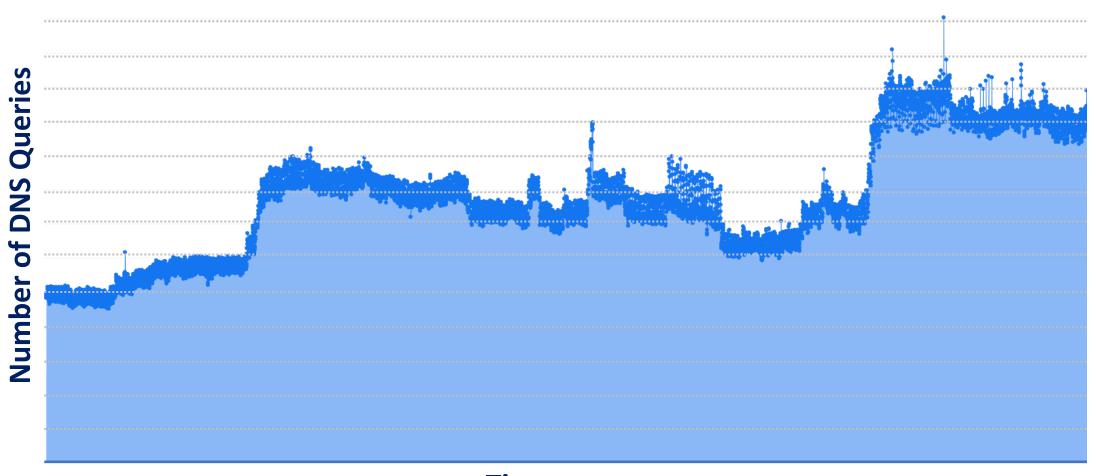
- Some Numbers
- A Peek at Our How DNS looks at Facebook
- Hitting the 4GB Wall
- Subdivision of Responsibilities
- Being a Good Neighbor
- Speeding up Propagation

# **DNS Groupings**

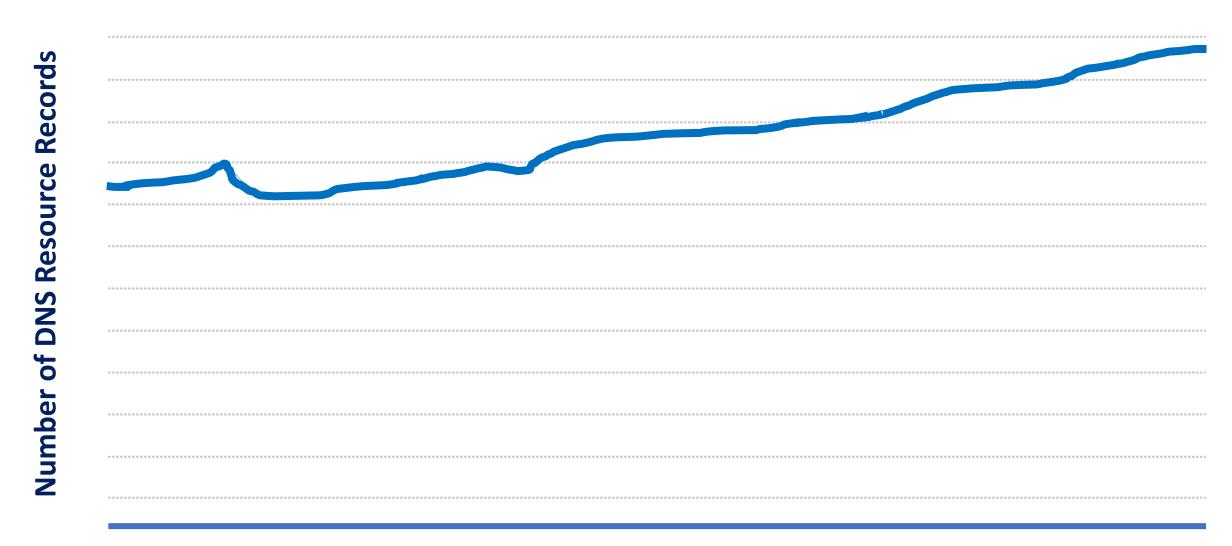


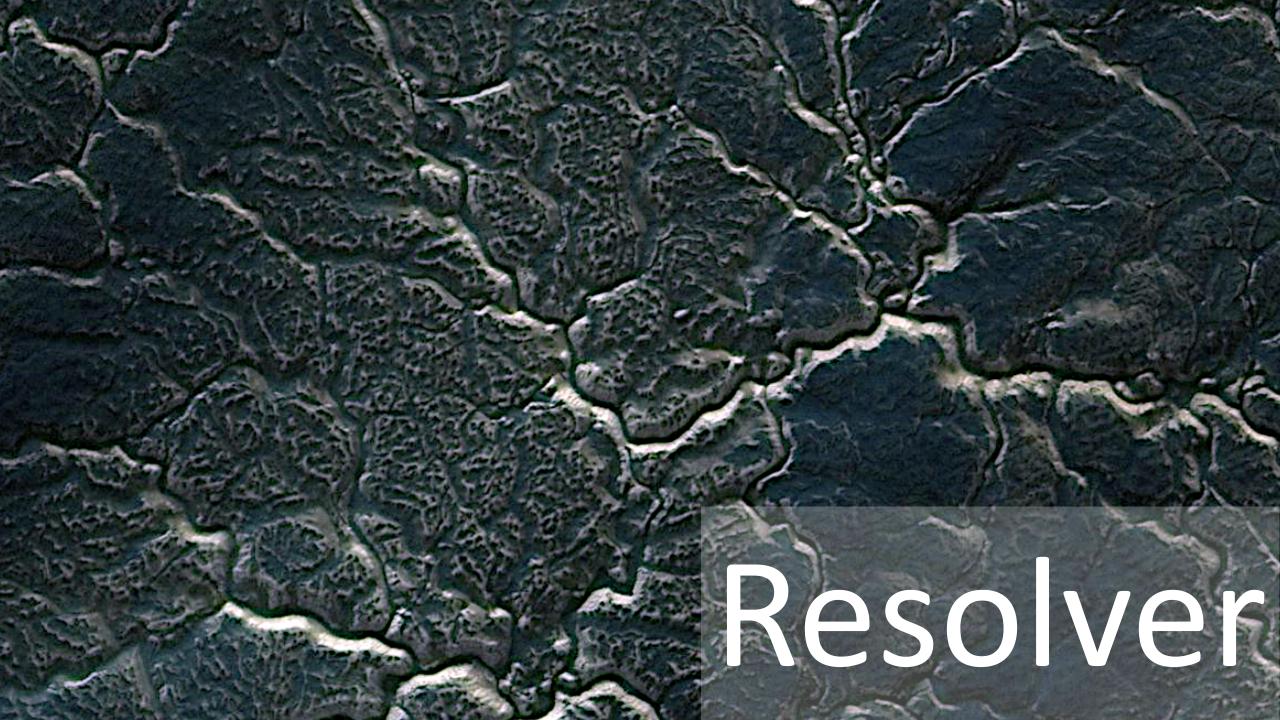


#### **Growth in Demand**

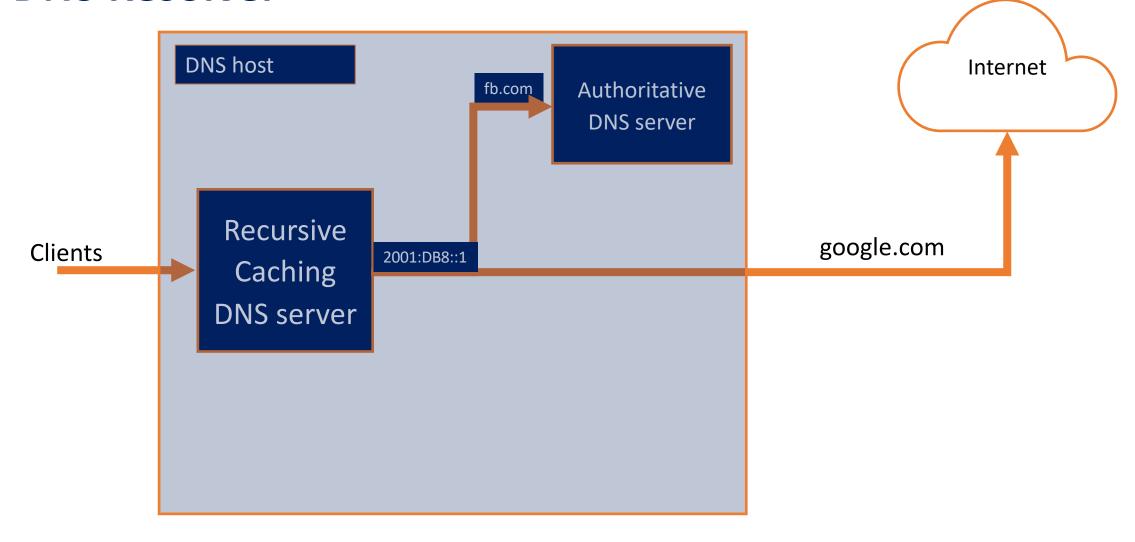


Time

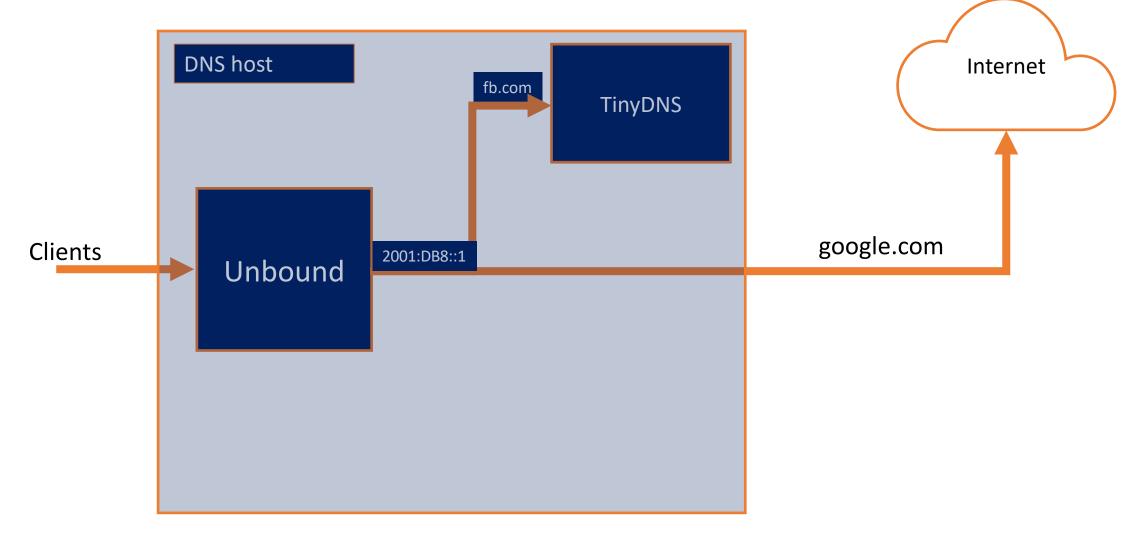




#### **DNS** Resolver

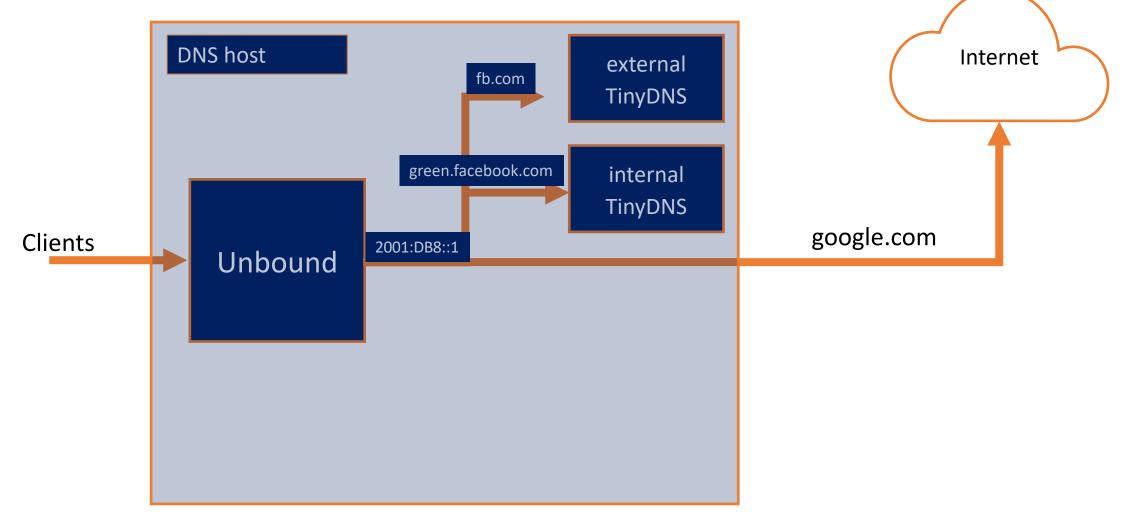


#### DNS Resolver – a closer look

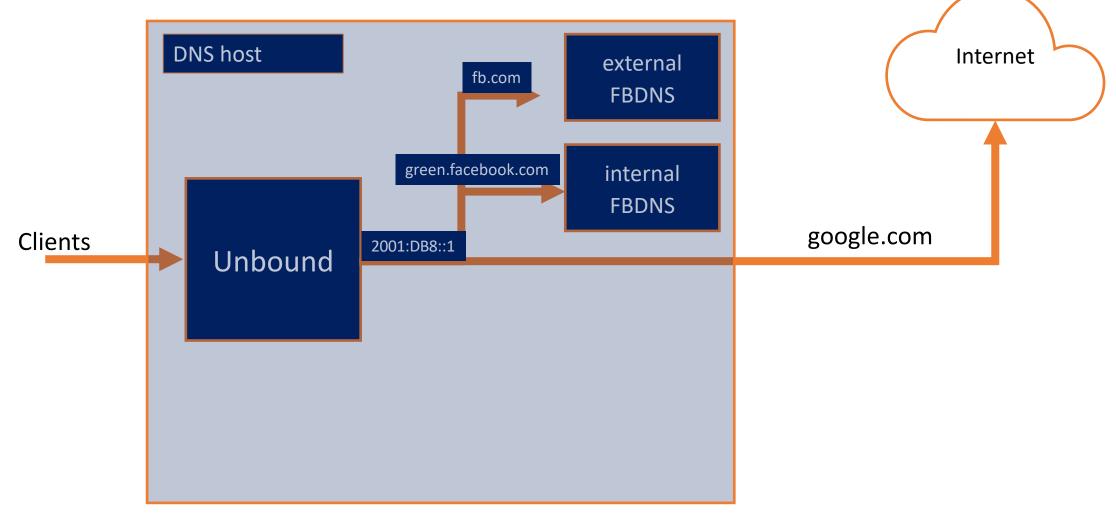




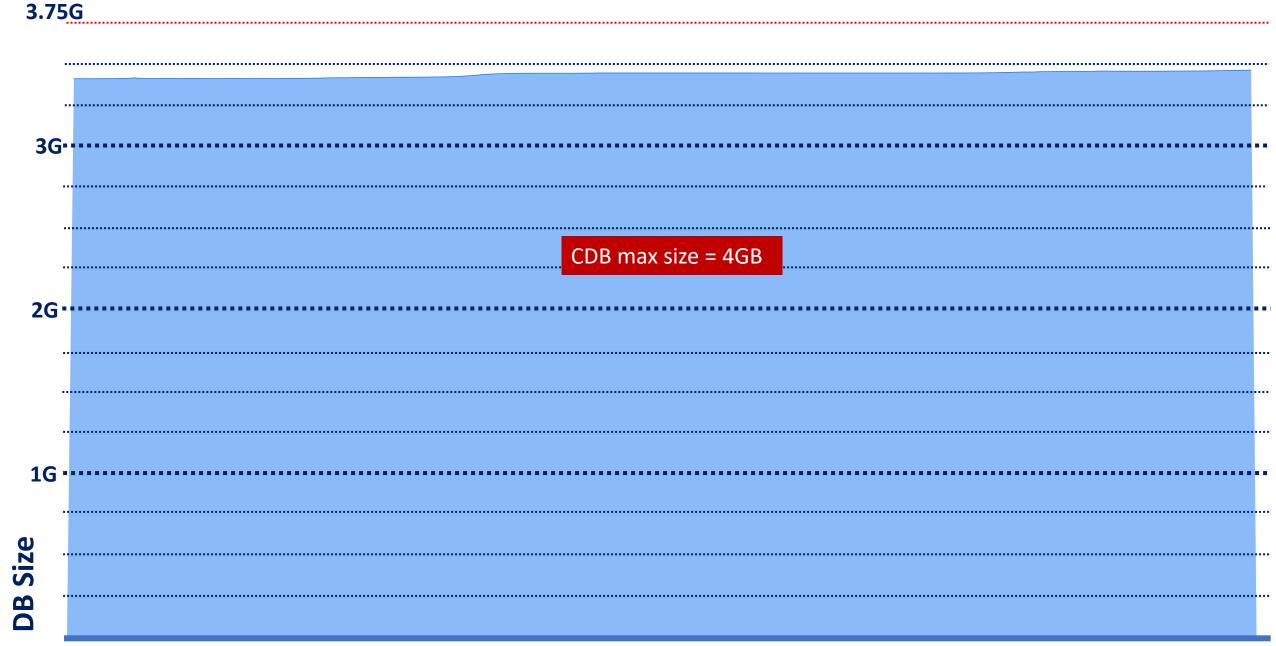
# **DNS** Resolver – sharding



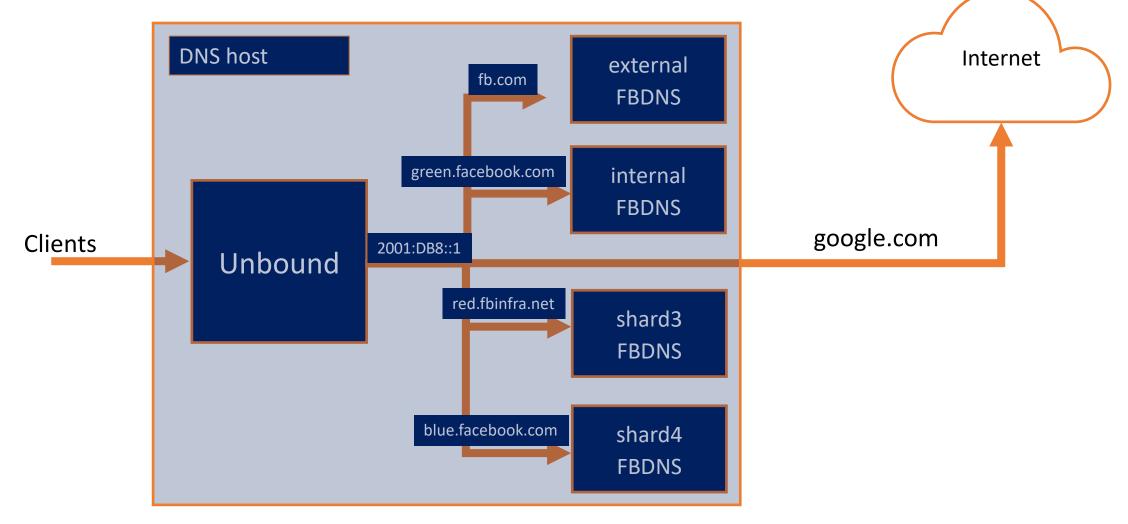
# **DNS** Resolver – sharding



#### Who Needs More Than 4GB?

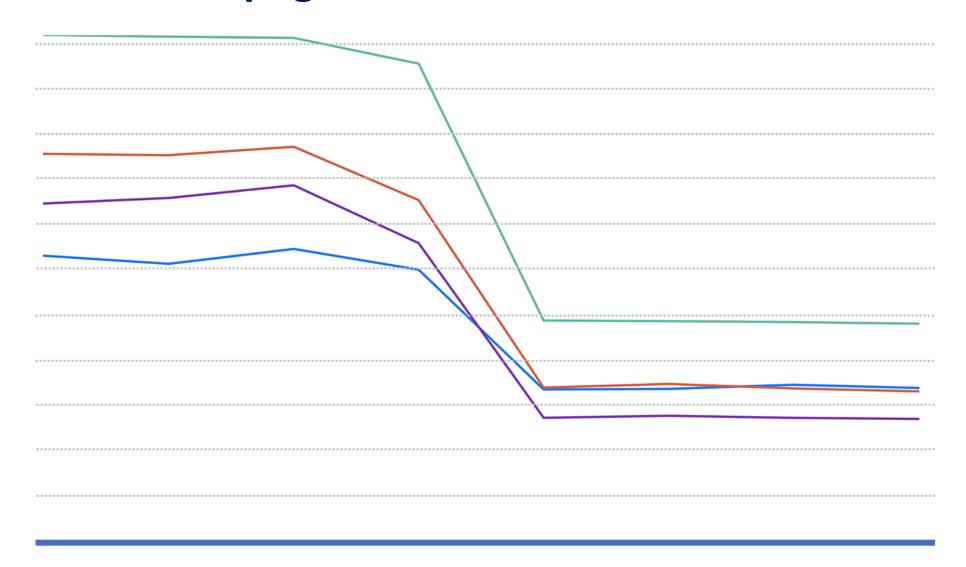


### Moar sharding

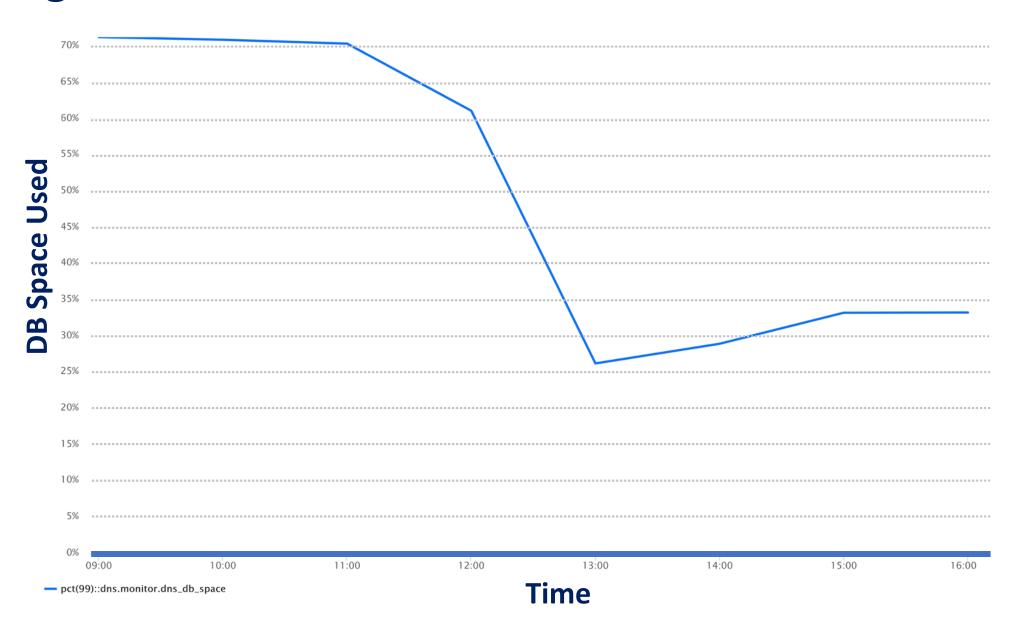


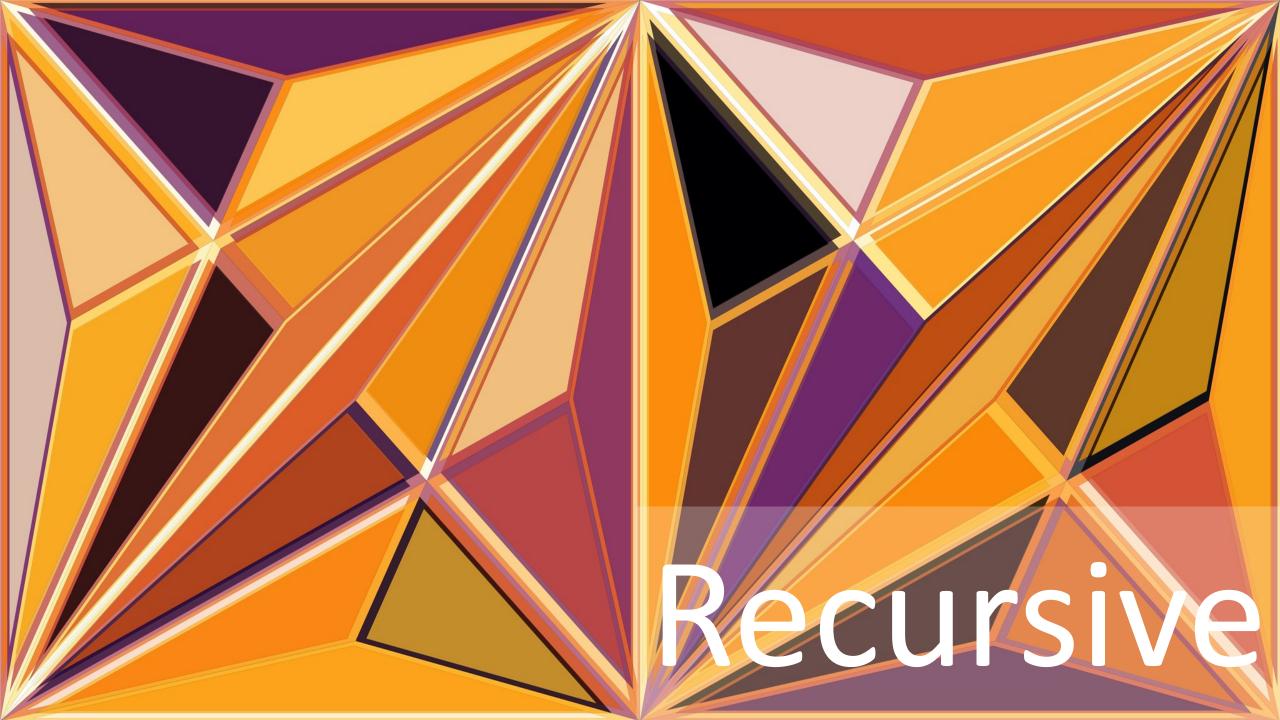


# **Faster Record Propagation**

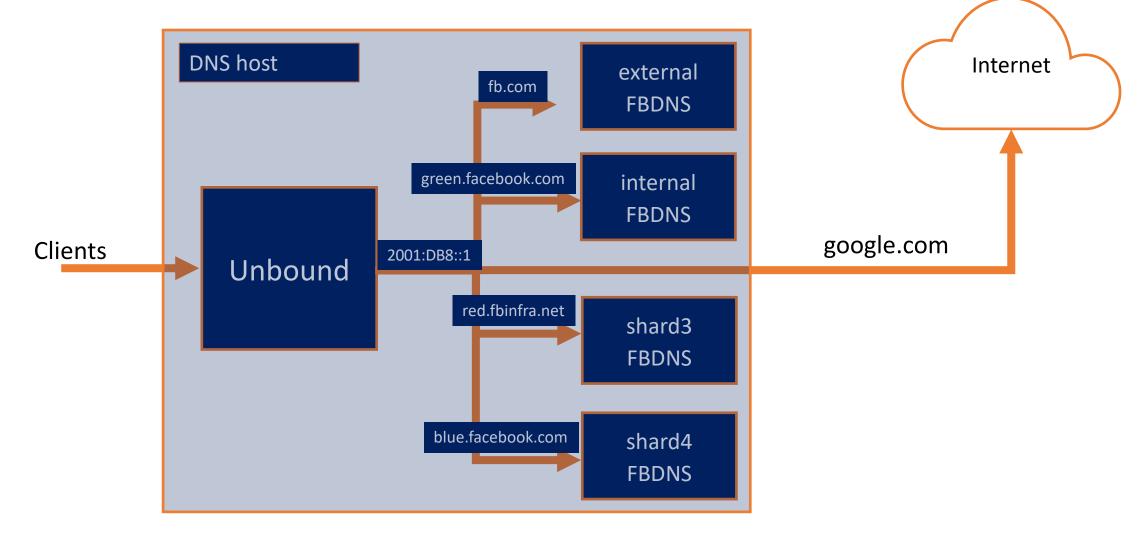


#### **Making Room on Disk**





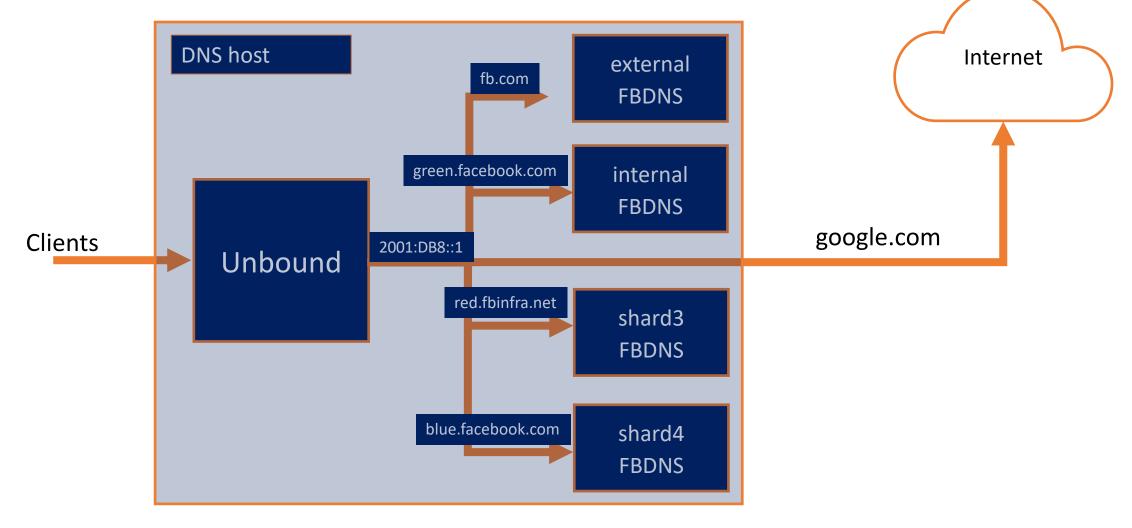
#### **A Closer Look at Recursion**



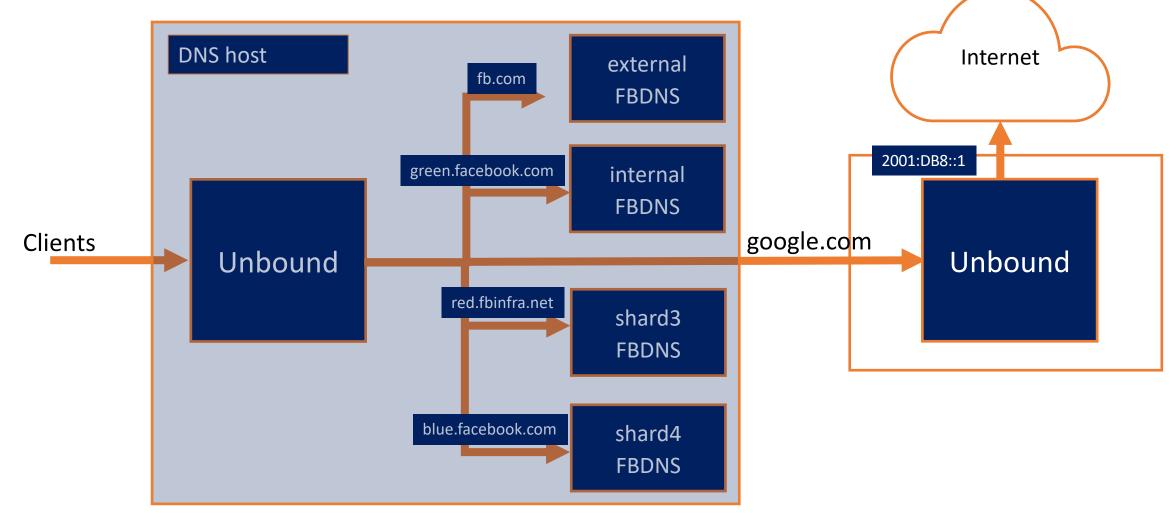
#### **Steering Queries**

```
$ cat /etc/unbound/intext/unbound.conf
name: "green.facebook.com."
    stub-addr: a.internal.ns.any.facebook.com.
    stub-addr: b.internal.ns.any.facebook.com.
name: "blue.facebook.com."
    stub-addr: a.shard4.ns.any.facebook.com.
    stub-addr: b.shard4.ns.any.facebook.com.
name: "red.fbinfra.net.
    stub-addr: a.shard3.ns.any.facebook.com.
    stub-addr: b.shard3.ns.any.facebook.com.
```

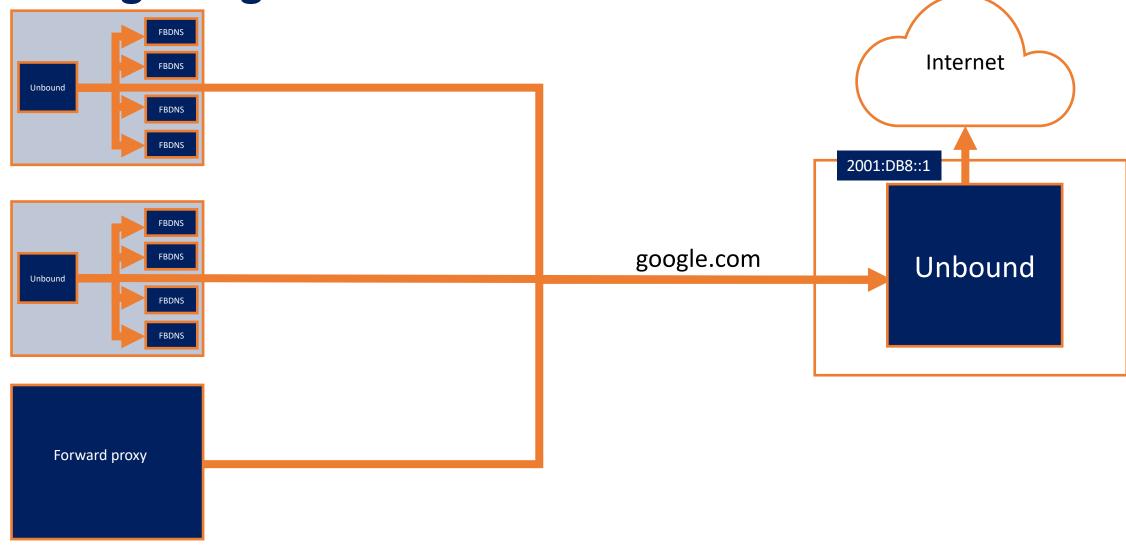
### **A Lot of Moving Parts**



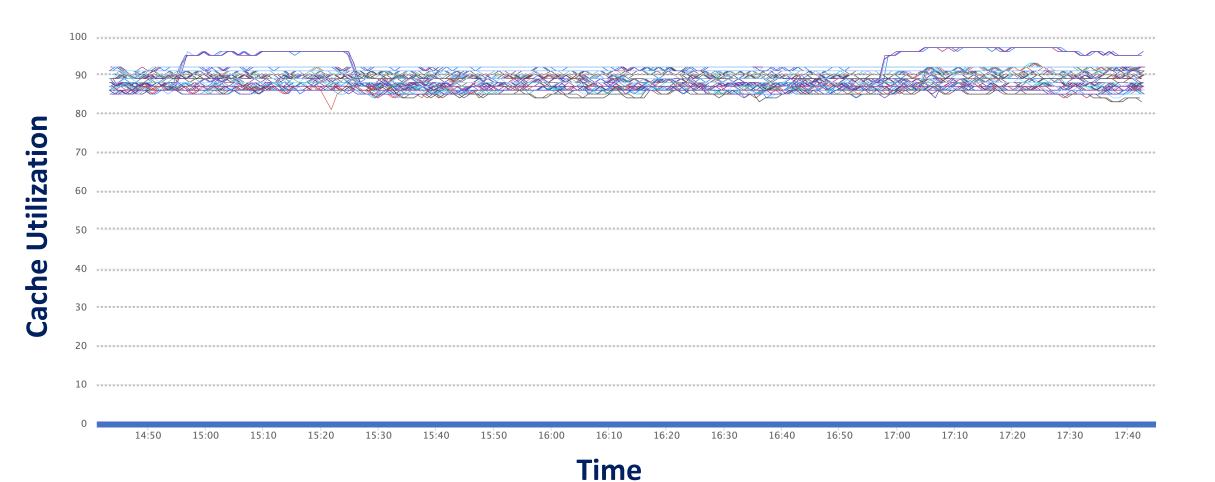
### **Dividing Out the Work**



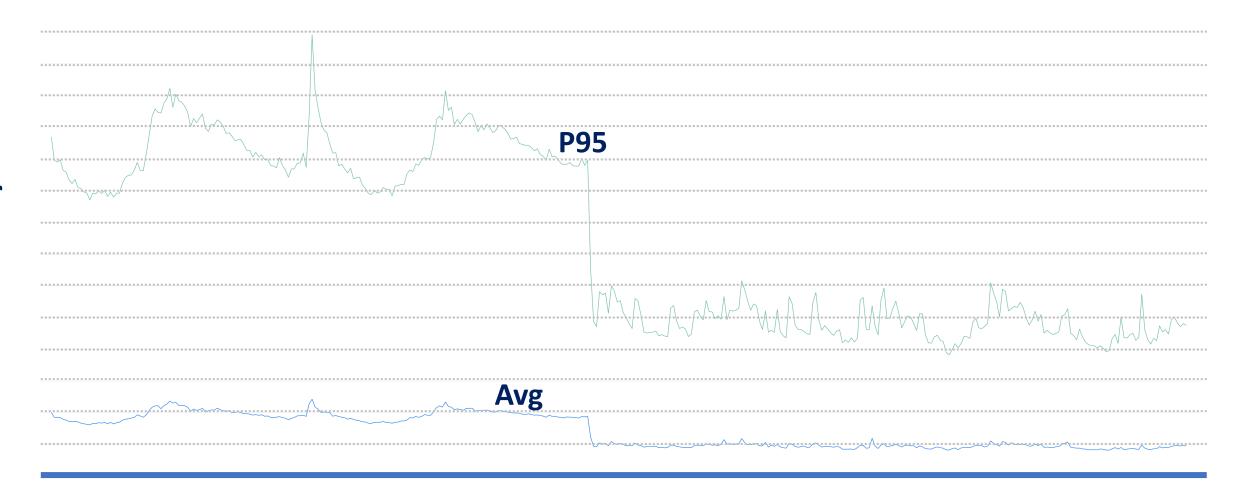
# **Stitching It Together**



#### **Cache Hits**

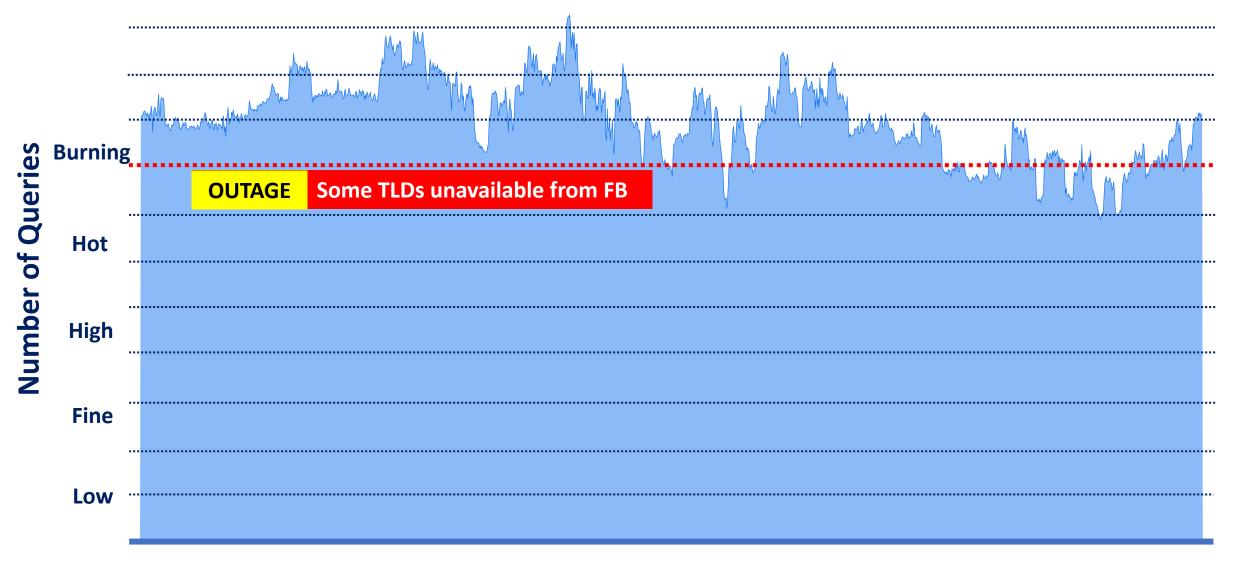


# **Speeding Things Up**



Time

# **Getting Blocked**



**Time** 

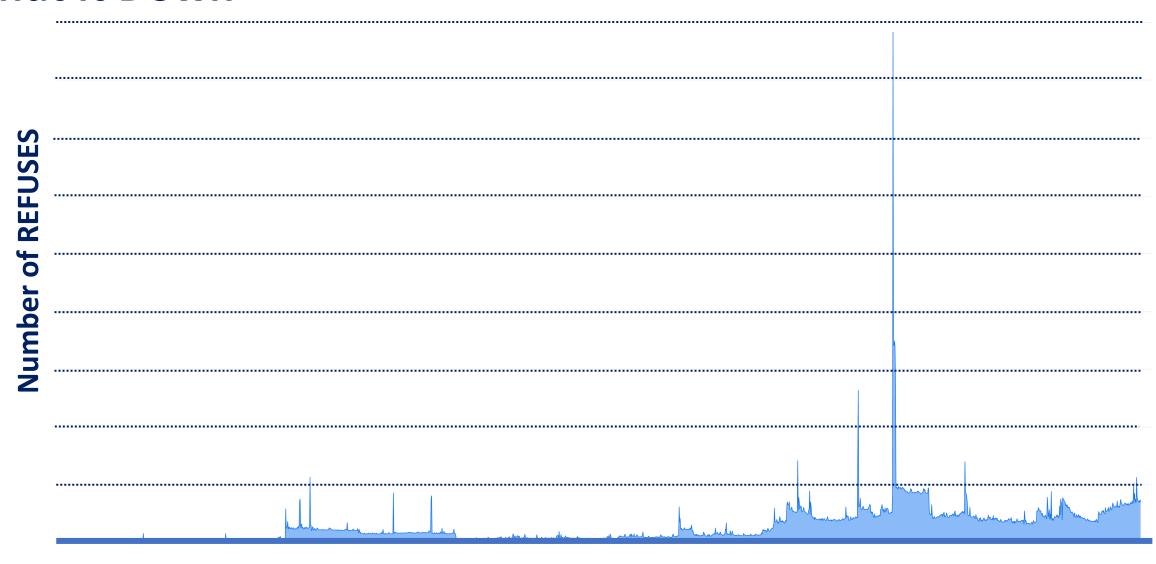
#### **Try Harder**

```
$ host -v test.foo
Trying "test.foo.first.search.path"
Trying "test.foo.second.search.path"
Trying "test.foo.third.search.path"
Trying "test.foo.fourth.search.path"
Trying "test.foo.fifth.search.path"
Trying "test.foo"
```

#### **No More Invalid Queries**

```
$ cat /etc/unbound/intext/unbound.conf
name: "com."
    forward-addr: a.internet.resolver.facebook.com.
    forward-addr: b.internet.resolver.facebook.com.
name: "blue.facebook.com."
    stub-addr: a.shard4.ns.facebook.com.
    stub-addr: b.shard4.ns.facebook.com.
```

#### **Shut It Down**



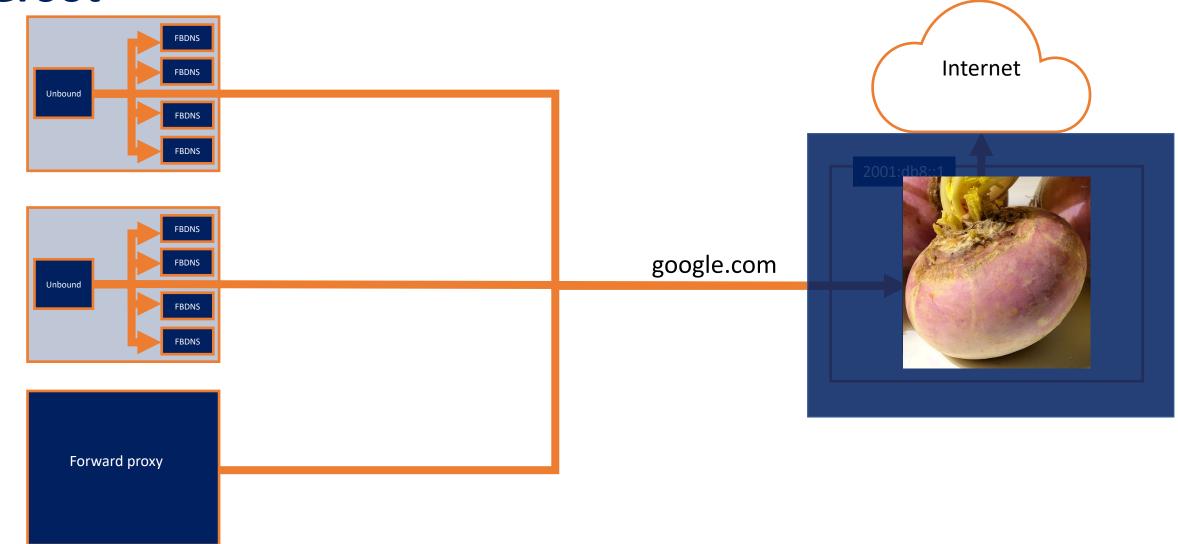
#### **Hearing Back Faster**



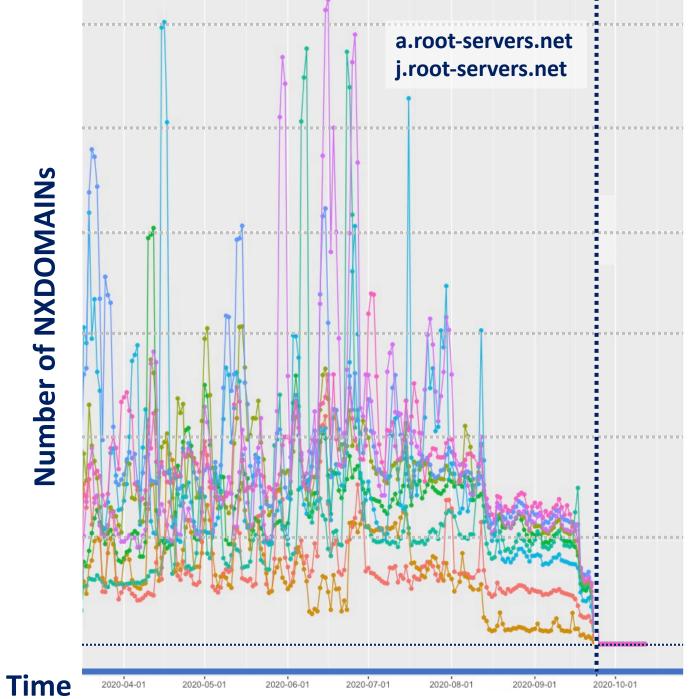
Ours is not to reason why, ours but to NXDOMAIN or REFUSE

```
search {list of domains }
$ time host -v invalid.host.in.lun1
                                                nameserver 2001:db8::1
                                                nameserver 2001:db8::2
Trying "invalid.host.in.lun1"
Host invalid.host.in.lun1 not found: 3(NXI options ndots:3
                                                options timeout:3
Received 113 bytes from 2001:db8::1::#53 coptions attempts:4
Received 113 bytes from 2001:db8::2::#53 in 72 ms
real 0m0.083s
$ time host -v invalid.host.in.lun1
Trying "invalid.host.in.lun1"
Host invalid.host.in.lun1 not found: 5(REFUSED)
Received 38 bytes from 2001:db8::1::#53 in 0 ms
real 0m0.012s
```

#### **Groot**

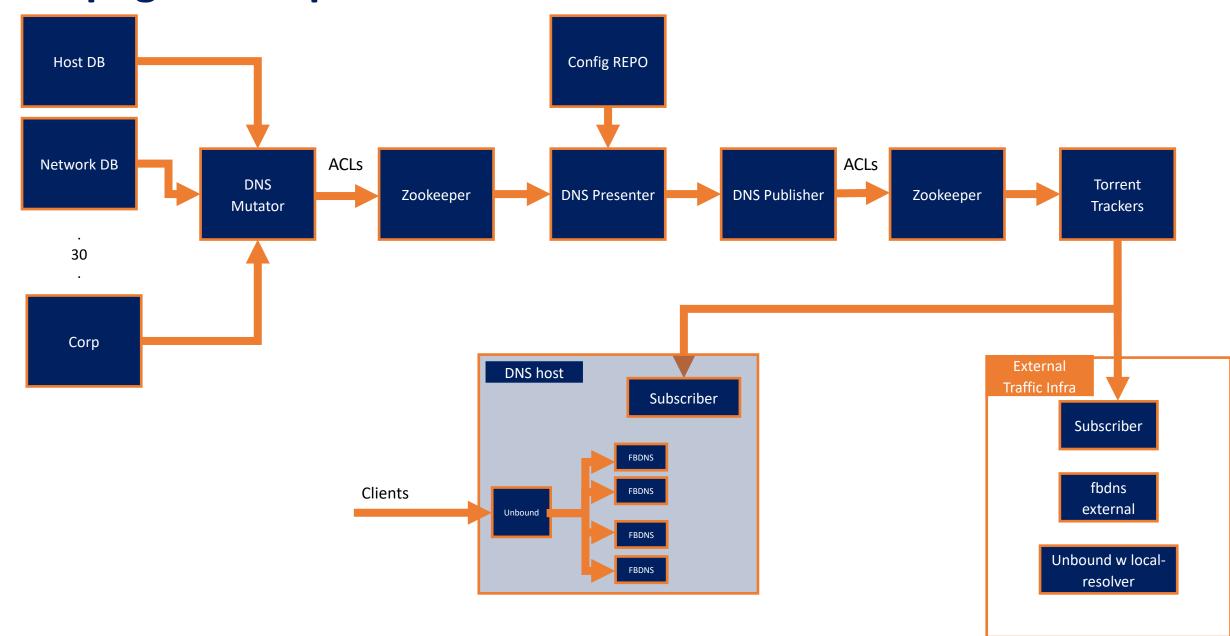




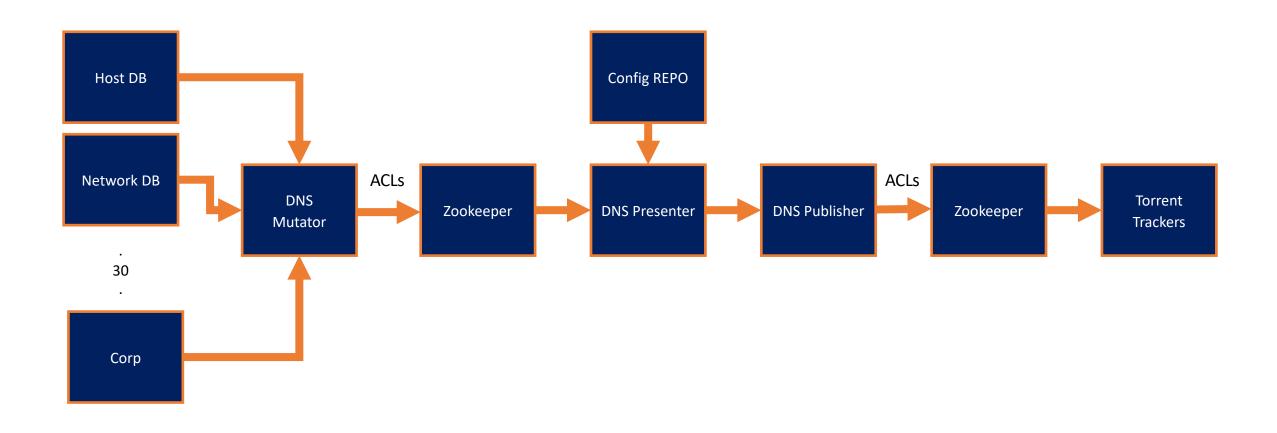




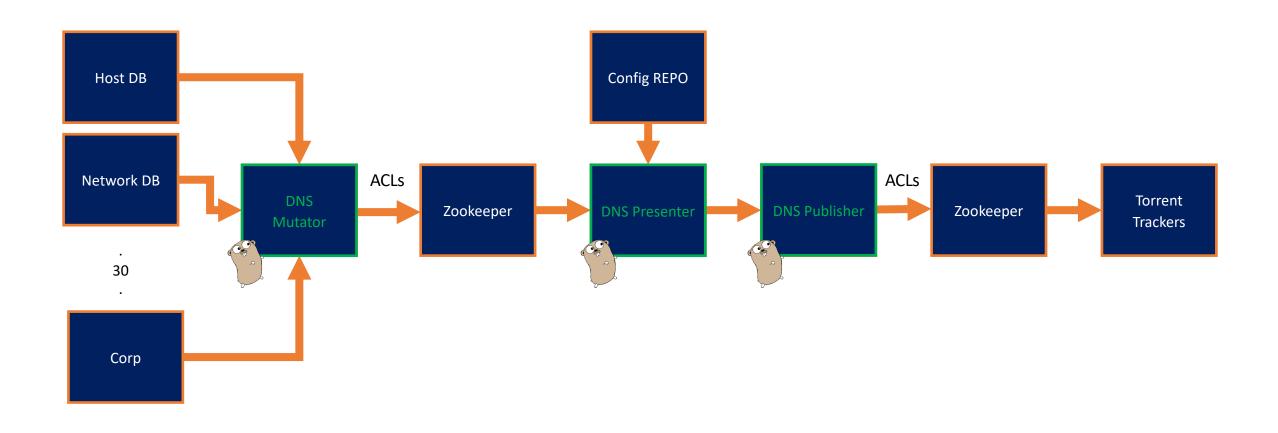
# **Propagation Pipeline**



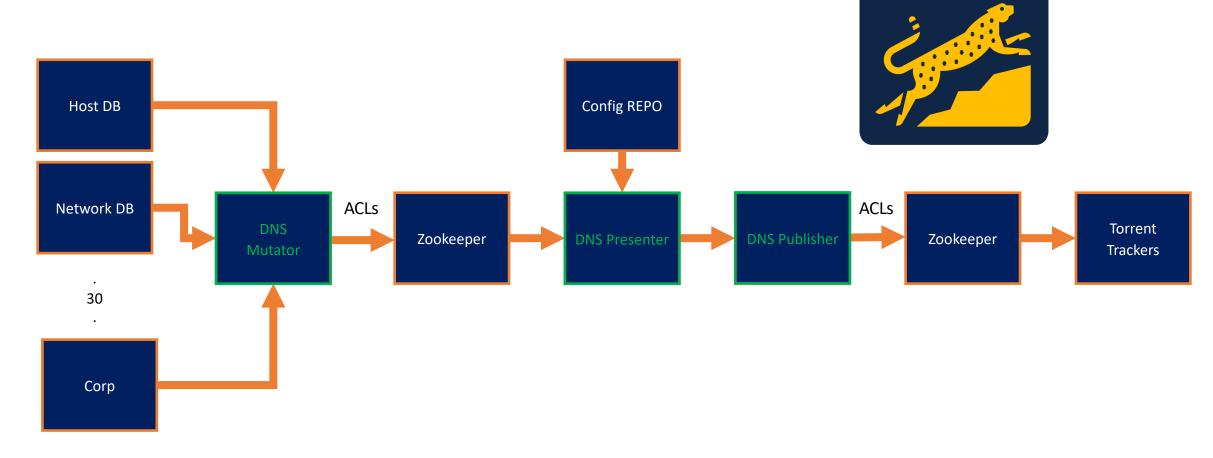
# **Speeding Up the Pipe**



# **Speeding Up the Pipe**



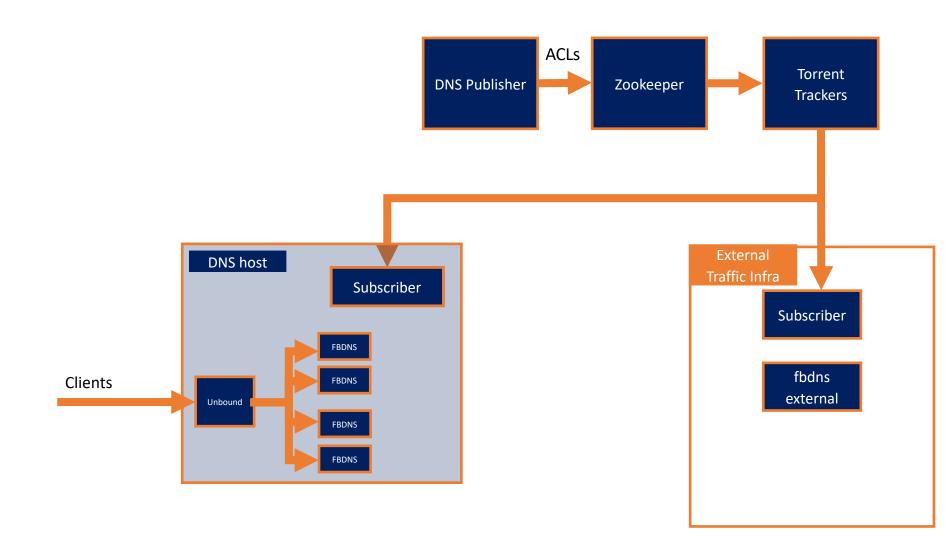
# **Speeding Up the Pipe**



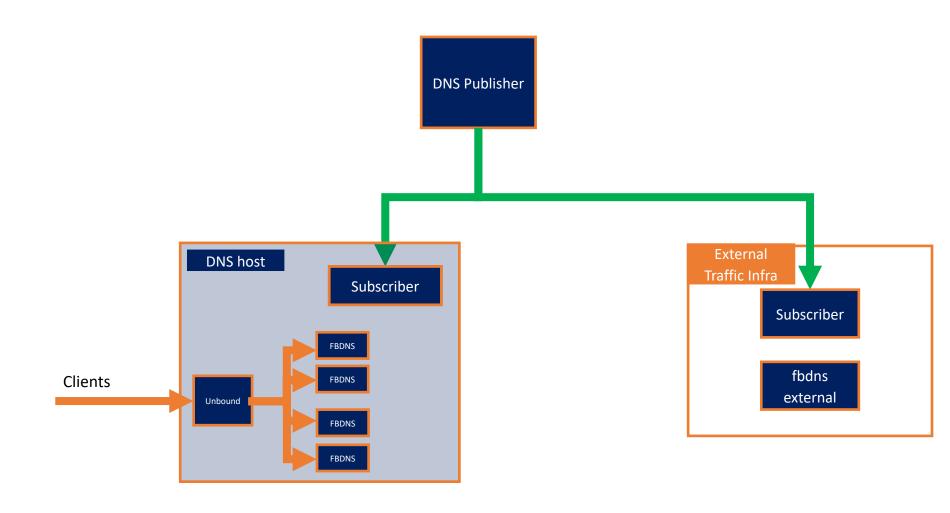
### **Stuck Pipes**

**OUTAGE** 

DNS propagation down, subscribers stuck on torrent download



#### **Trackerless**



# **Final Takeaways**

- Expectations of DNS uptime are high
- Rapid growth and scale create demands on DNS that call for a flexible approach to ensuring the service remains reliable.
- An incremental approach pays dividends.
- Go loves DNS

# facebook Thank You

# facebook