

The OpenDNS logo consists of the word "OpenDNS" in white, sans-serif font, centered within an orange rounded rectangle. The background of the slide features a faint, light gray map of a city grid.

**OpenDNS**

# Real World Impacts of EDNS Client Subnet

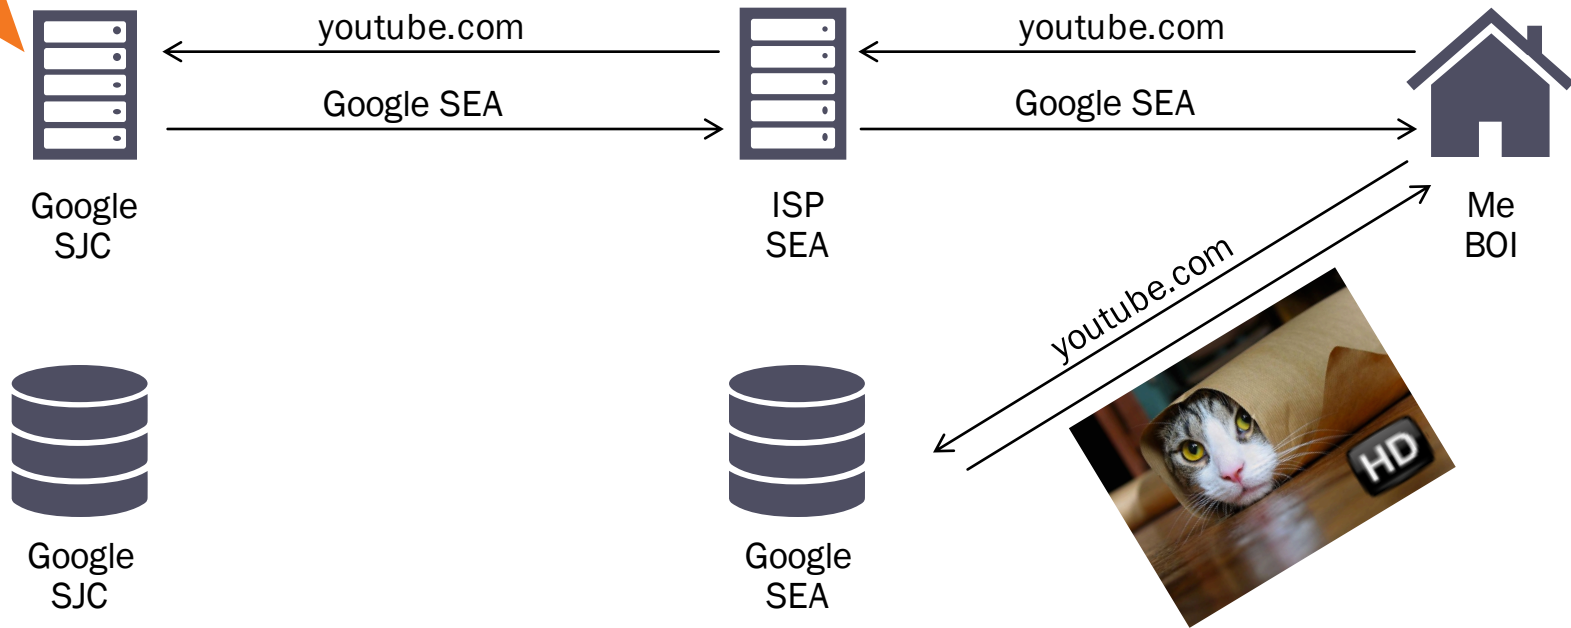
Brian Hartvigsen

Manager, Site Reliability Engineering

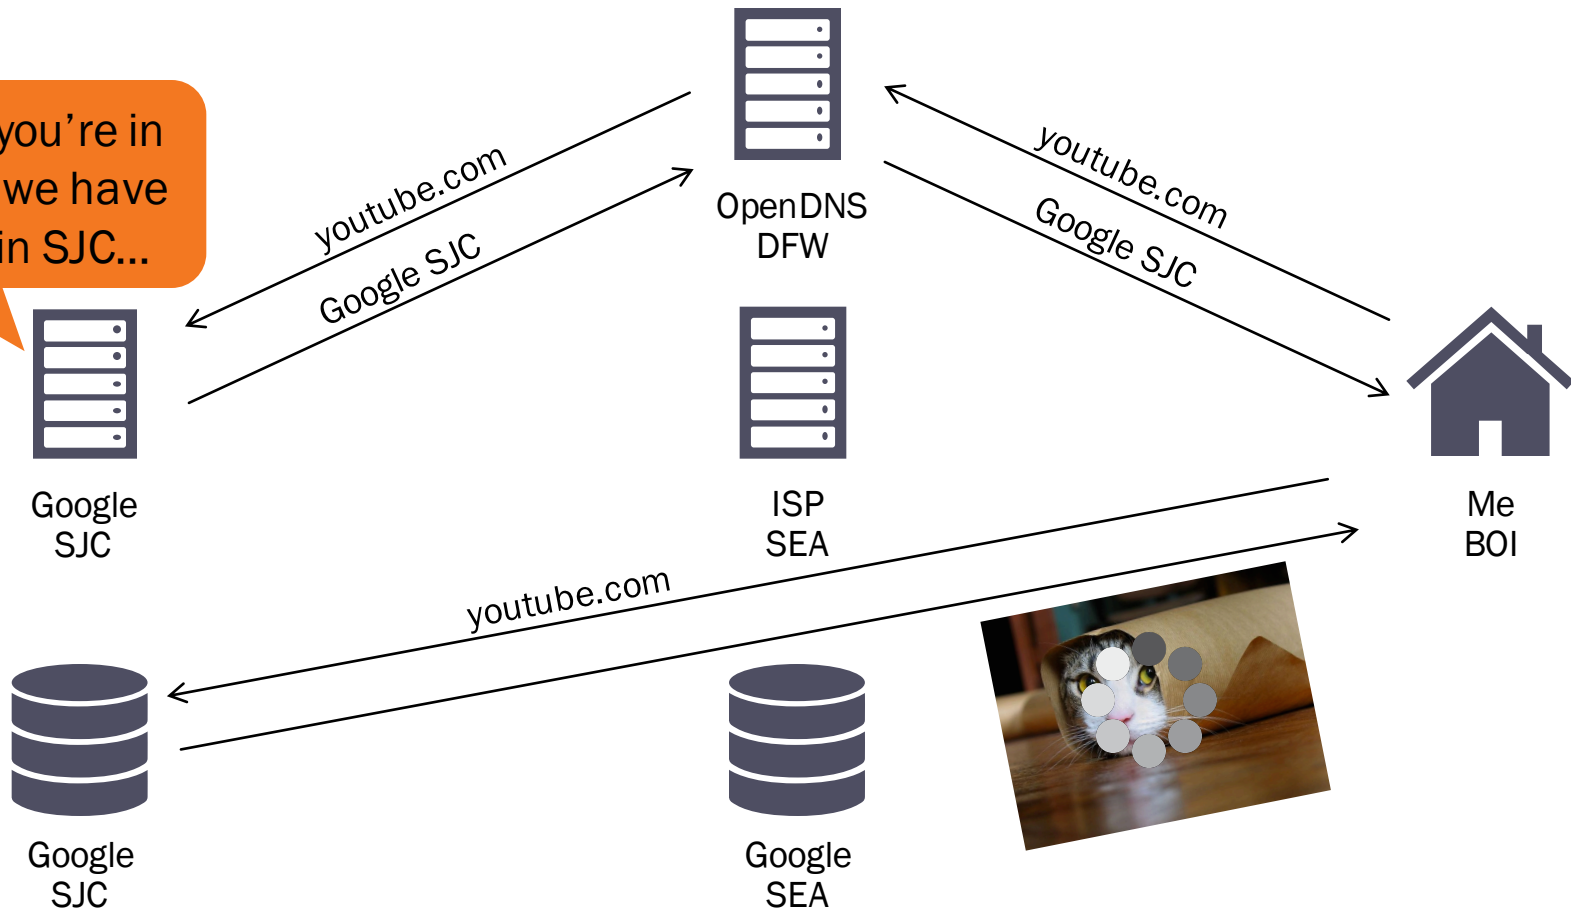
OpenDNS is  
now part of Cisco.



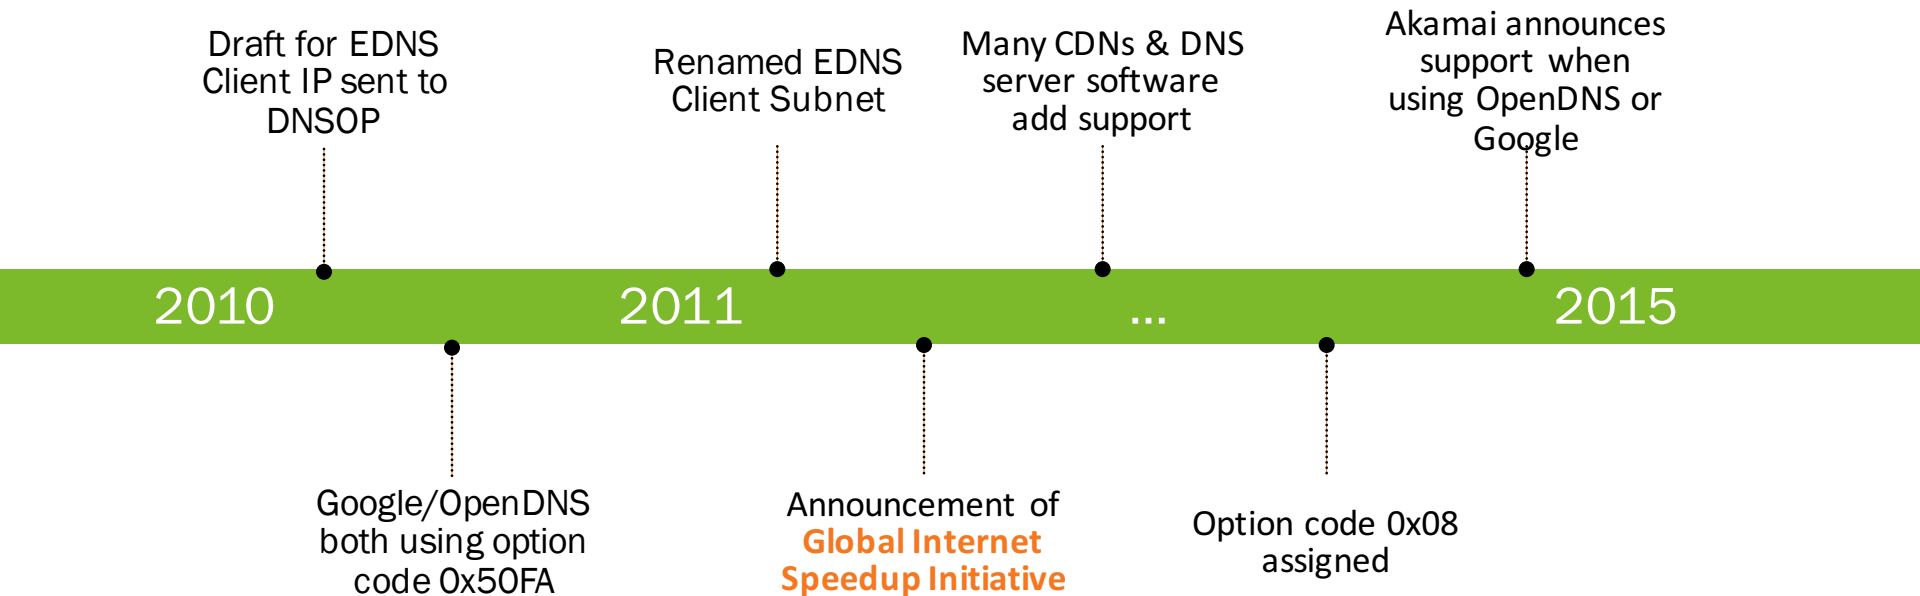
I see you're in SEA, we have a DC there...



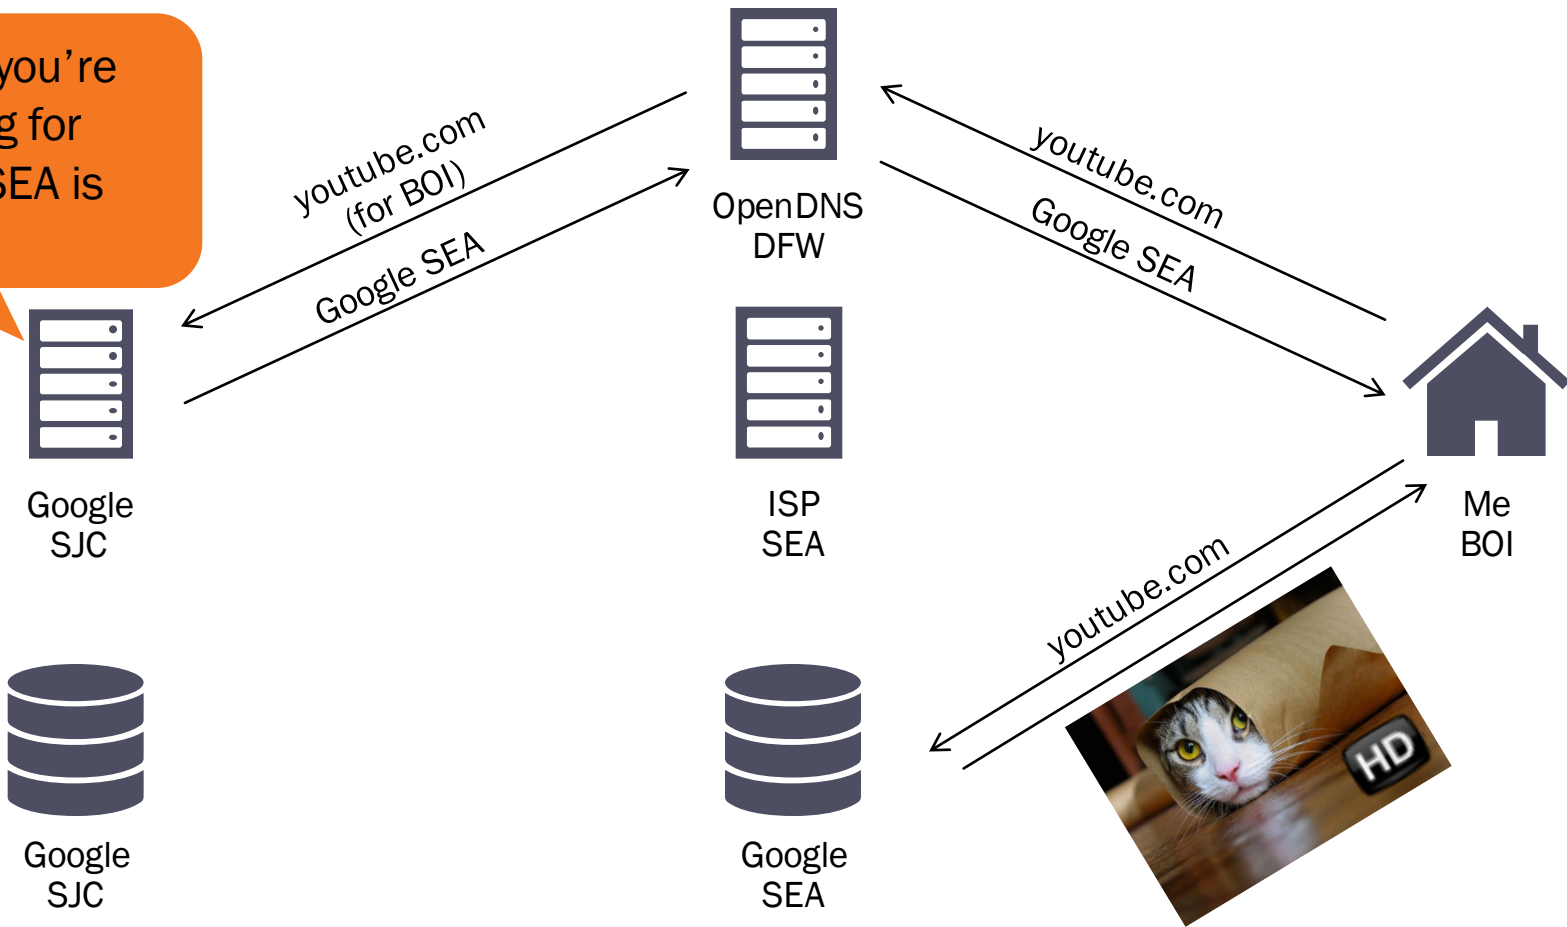
I see you're in DFW, we have a DC in SJC...



# A Short History




I see you're asking for BOI, SEA is good!



# A Couple Notes on OpenDNS' implementation

- Whitelist Only
- Not accepted from client side
- Fixed scope, manageable via configuration


# The Setup



- 
- Mirror production traffic to 2 resolvers
  - Identical traffic
  - ~2 weeks of data
  - Enable ECS on only 1
  - Monitor
-


# Start from the user's perspective

- Biggest measure is RTT
- ~5000 qps
- ~85% <25ms
- ~95% <300ms






- Still ~95% <300ms
- But only ~80% <25ms
- More jitter in RTT



# Cache Hits vs Cache Misses


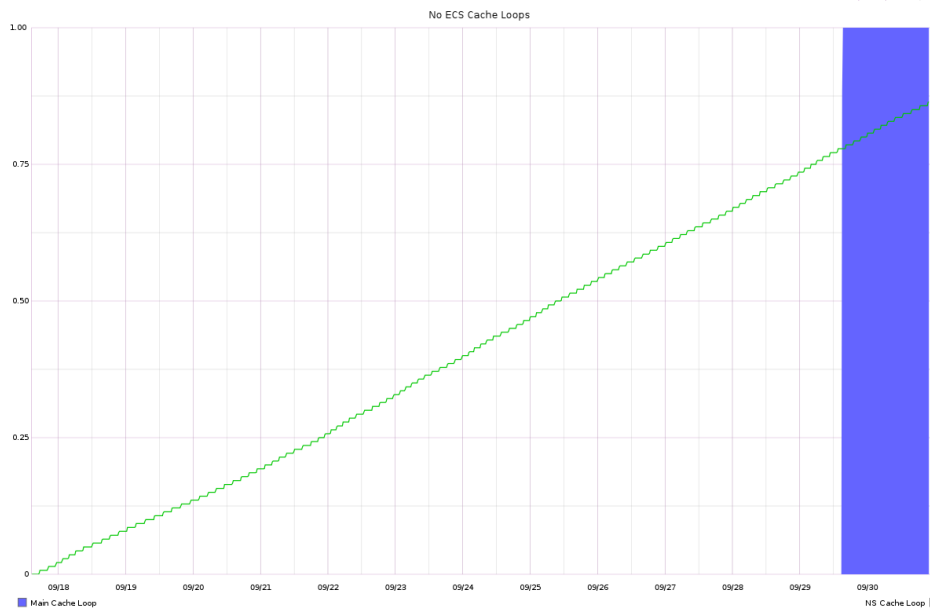
- Cache Miss is no data **or** record with non-matching ECS info
- Every failed lookup =  $2 + n$  number of misses
  - $n$  = number of possible ECS enabled records that won't match (max 16,777,215)
- Store ECS ability as a special record in the cache
- Every successful lookup == 2 cache hits (special record + actual record)

Cache Hits vs Misses




# Cache Churn

- Fixed cache size per resolver
- Allocate all available memory
  - Bi-directional linked list
- Measure cache “cycle”
  - How often we run out of room in the cache to store something



# Moving upstream



# Bumps in the Road

- Malformed responses
- Common one was getting back additional 0 bits in address
  - Remember to chop it to NETMASK

# Bumps in the Road

- Different A records for nameservers w/ ECS
  - example.com NS ns1.example.com
  - ns1.example.com A 192.0.2.53 (for 1.2.3.0/24)
  - ns1.example.com A 192.0.2.153 (for 1.2.4.0/24)
  - ns1.example.com A 192.0.2.253 (no ECS)
  - Which records should you use when resolving www.example.com for 1.2.3.0/24?
- Discussion and clarification around this is welcome