# Charter - DoH Trial Readout

Jason Weil – Principal Engineer, Emerging Technology
Todd Medbury – Principal Engineer, Infrastructure Arch&Eng
Ricardo Meleschi – Principal Engineer, Provisioning & Appl

# DoH Initial Setup and Testing

# *Virtual setup speeds configuration and initial testing*

Virtual Load Balancers

- Configuration
- Health Checks

Virtual Machines

- Configuration
- Test Deploy

Test Responses through stack

Use client to look for *expected* response codes

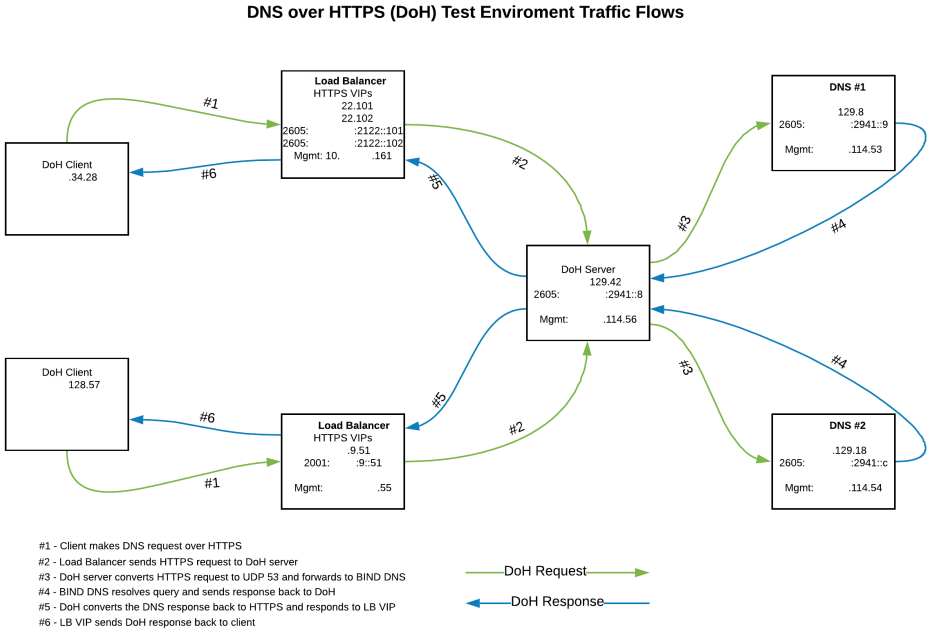


*This type of setup gives us confidence before physical hardware deployment. Everything is working, end to end.*



Charter COMMUNICATIONS

# Load Testing

# Virtual setup is deployed to physical hardware.

*How well is this going to work compared to our existing Do53 setup?*
*Let's apply some load.*

External testing was applied, but wasn't stressing system.
Let's add some more load via scripting.

- Test against a single endpoint to target one machine
- Test with differing approaches to show possible scenarios
  - Tested first with the list of lookups from my home DNS
    - As fast as possible (160 Threads)
    - Sequentially
  - Test second with one lookup, lots of threads, no throttle

# DoH Certificate Testing

# *Certificate type, strength will alter overhead for setup*

Tests were run for each  cert type.

Runs were sequential, this way, setup, query, and tear down are sequential, and will be additive, to show differences between cert types.

Runs of 10k, 100k, and 1M queries were run 10 times each for each cert type.

These were graphed, and a cert selected.

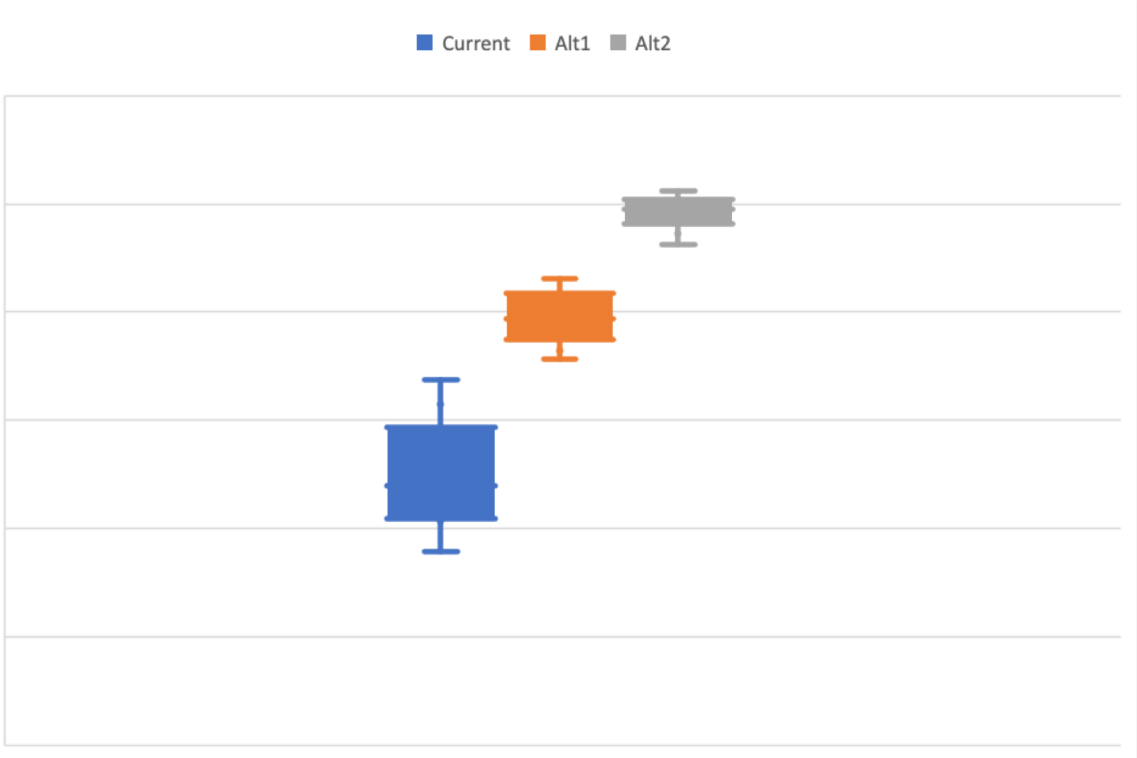


*Test. Test. Test. These will change by software used, hardware used and certificate type used..*



Charter COMMUNICATIONS

# Linux Server Tuning

# Linux stack tuning is important

Increase software max clients to expected levels!

Tune the stack – these need testing and environment knowledge.
#vim /etc/sysctl.conf (these are exapmles)

fs.suid_dumpable = 0
kernel.core_pattern = %e.%p.core kernel.core_uses_pid = 0
kernel.randomize_va_space = 2
kernel.sysrq = 0
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.all.accept_source_route = 0
net.ipv4.conf.all.log_martians = 1
net.ipv4.conf.all.secure_redirects = 0
net.ipv4.conf.all.send_redirects = 0
net.ipv4.conf.default.accept_redirects = 0
net.ipv4.conf.default.accept_source_route = 0
net.ipv4.conf.default.secure_redirects = 0
net.ipv4.conf.default.send_redirects = 0
net.ipv4.ip_forward = 0
net.ipv4.route.flush = 1
net.ipv4.tcp_syncookies = 1
net.core.default_qdisc = fq
net.core.rmem_max = 2147483647
net.core.wmem_max = 2147483647
net.ipv4.tcp_rmem = 4096 87380 2147483647
net.ipv4.tcp_wmem = 4096 87380 2147483647
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_sack = 1
#sysctl -p

Increase File descriptor limits! ( we used 1048576)

#vim /etc/security/limits.conf
# End of file

- hard core 0
- * soft nofile 1048576
- * hard nofile 1048576

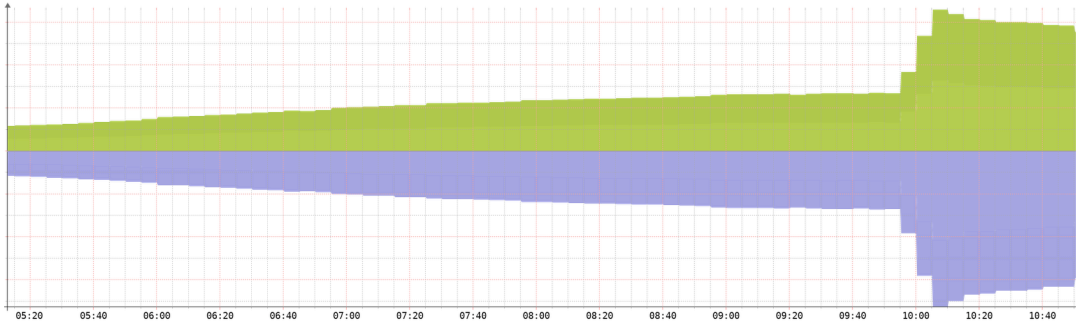


*These will be implementation specific – don't underestimate stack tuning!*



Charter COMMUNICATIONS

# DoH Trial Setup

# DoH Trial Infrastructure

- Two Independent DoH Clusters
  - California Cluster
    - *doh-01.spectrum.com*
  - Texas Cluster
    - *doh-02.spectrum.com*
- Load Balancers
  - Unique vendor per cluster
    - TLS sessions are terminated on the servers
    - LB Health Checks modified for DoH support
      - Pass = HTTP status 200 (OK)
      - base64 encoded URI check

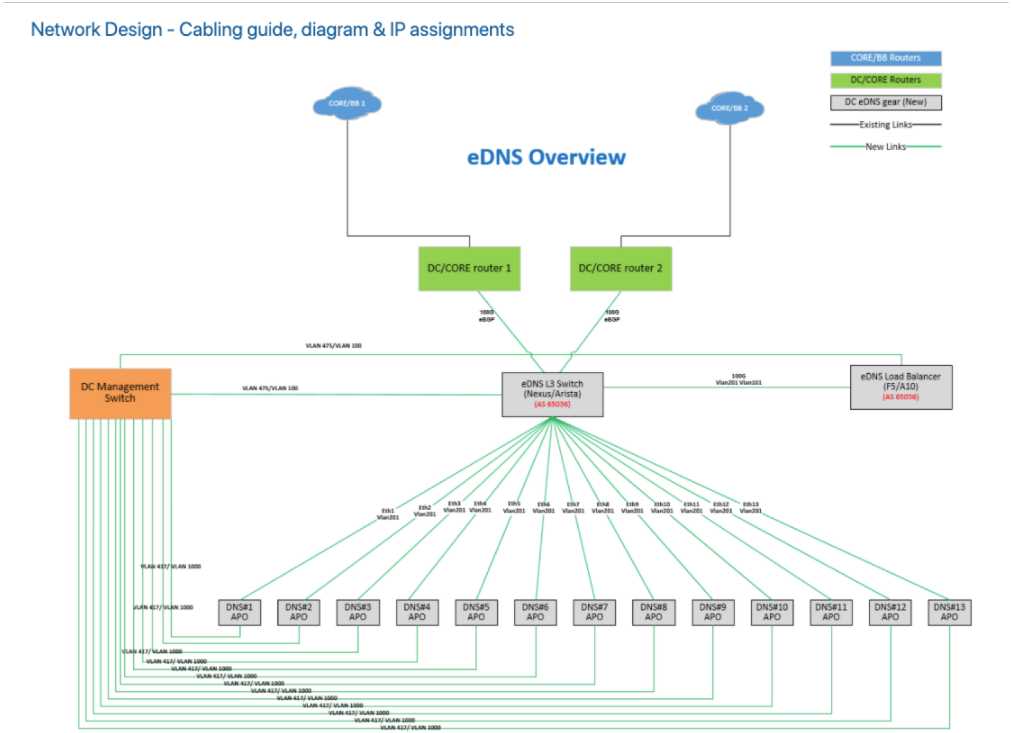


*Charter recently transitioned doh-01.spectrum.com and doh-02.spectrum.com trial to an anycast model.*



Charter COMMUNICATIONS

# Chrome (Same-provider auto upgrade) SPAU

# *Chrome auto-upgrade*

Charter began participating in Chrome 87

Chrome is driving significant DoH traffic useful for testing

Most DoH traffic is IPv6 due to RFC7804 requirements (eRouter for cable)

- L-10 DHCPv6 DNS_SERVERS
- L-11 RA Recursive DNS Server (RDNSS)

Auto-upgrade Do53 servers in Chrome

Do53 IPv4 - 209.18.47.61 & 209.18.47.62

D053 IPv6 - 2001:1998:0f 00:0001::1
2001:1998:0f 00:0002::1

IPv4 DNS WAN Config:
209.18.47.61
209.18.47.62

~~IPv4 DNS LAN Config:~~
~~192.168.1.1~~

**No DoH Upgrade**

IPv6 DNS WAN Config:
2001:1998:0f 00:0001::1
2001:1998:0f 00:0002::1

IPv6 DNS LAN Config:
2001:1998:0f 00:0001::1
2001:1998:0f 00:0002::1

**DoH Upgrade**

Chrome DoH Servers:
https://doh-01.spectrum.com/dns-query
https://doh-02.spectrum.com/dns-query

*Only IPv6 DNS Servers get auto-upgraded*



Charter COMMUNICATIONS

# Chrome DoH Trial generates testing load

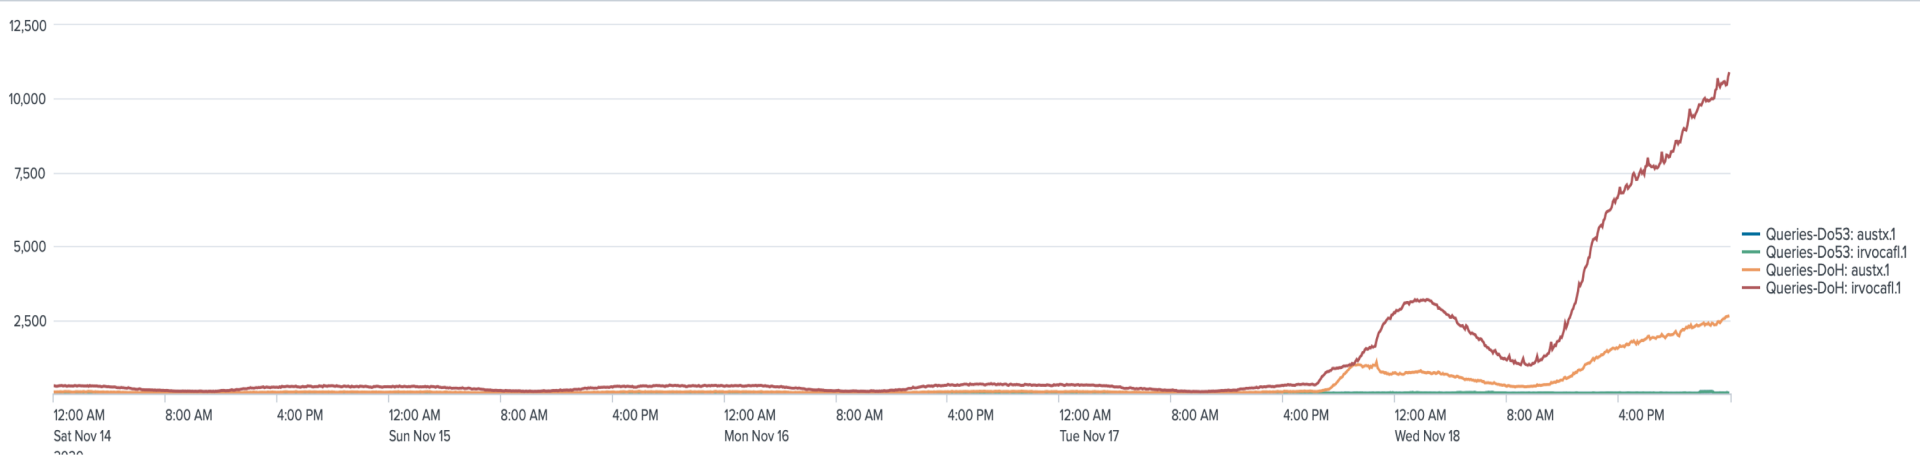


*Charter inclusion in Chrome DoH trial kicked off on November 18th*



Charter COMMUNICATIONS

# Load Test

## What happens when your resolver's CPU hits 100%?

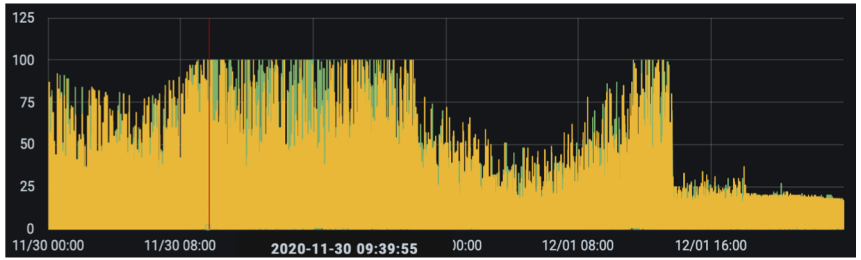


CPU Load

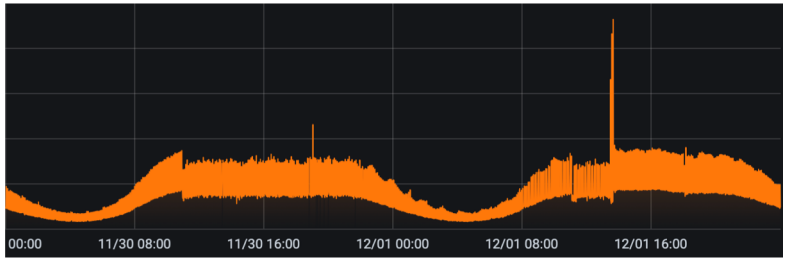


Queries per second

- CPU load reaches 100%
- Queries per second level off
- Network traffic spikes

*Spike in network traffic is due to aggressive client retries*


Network Traffic



Charter COMMUNICATIONS

# Future Testing

# Next Steps

*Performance analysis and improvements*

- Current Do53 platform supports only a fraction of DoH QPS
  - Explore hardware offload for TLS on recursive servers
  - Explore offloading TLS sessions to the load balancers

*New standards and deployment models*

- Evaluate and test impact of new DNS records
  - SVCB – Service Binding
  - HTPPS (formerly HTTPSSVC)
- Evaluate and test proposed resolver discovery mechanisms
  - https://datatracker.ietf.org/wg/add/documents/

# Thank You

## Questions?