COMCAST

IPV4 AND IPV6 LATENCY ANALYSIS

FOR

ISOC LATENCY WORKSHOP, SEPT 25-26, 2013

SALIL BANERJEE
COMCAST CABLE

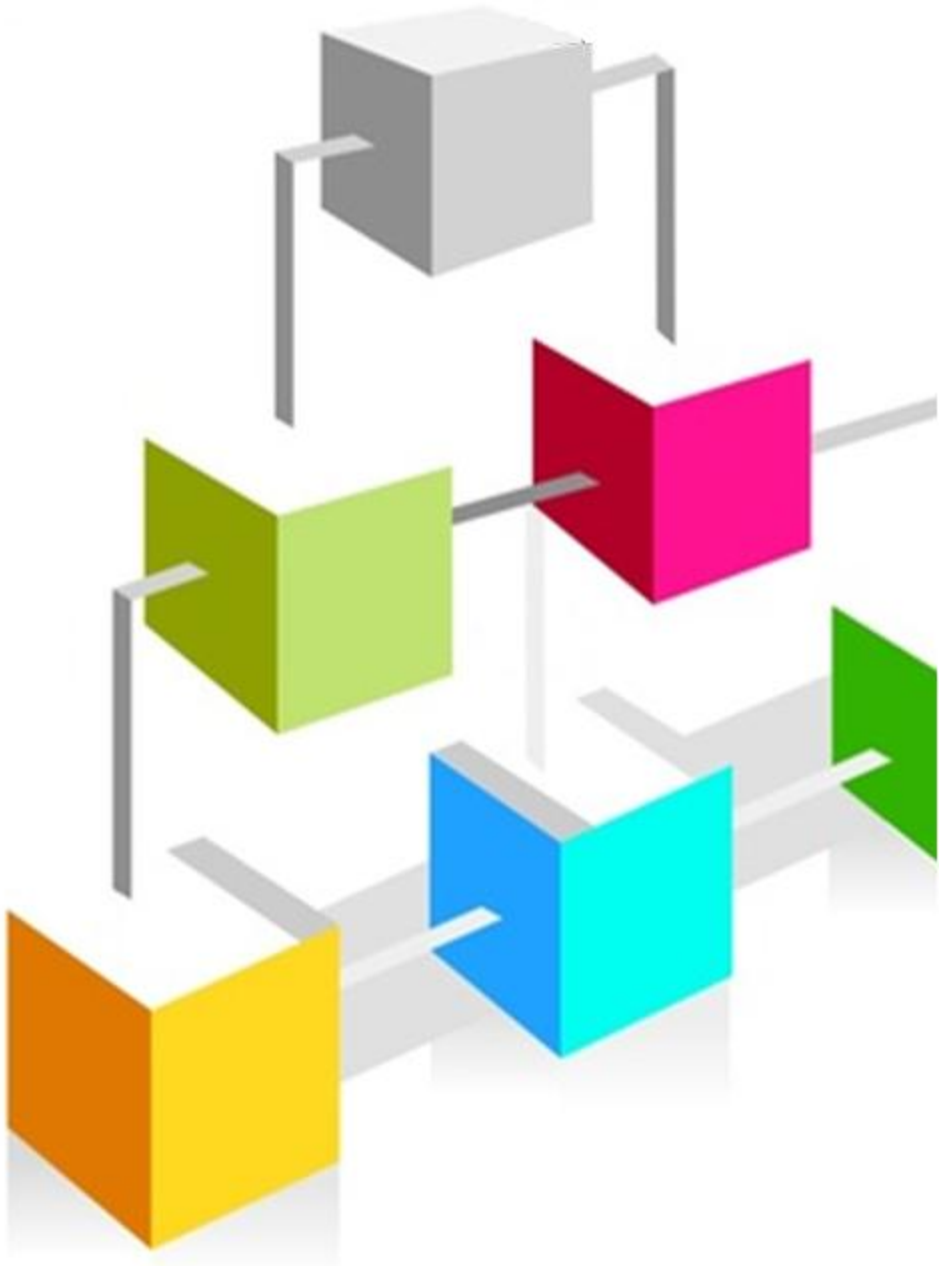

NE&TO

Comcast Latency Measurement Infrastructure

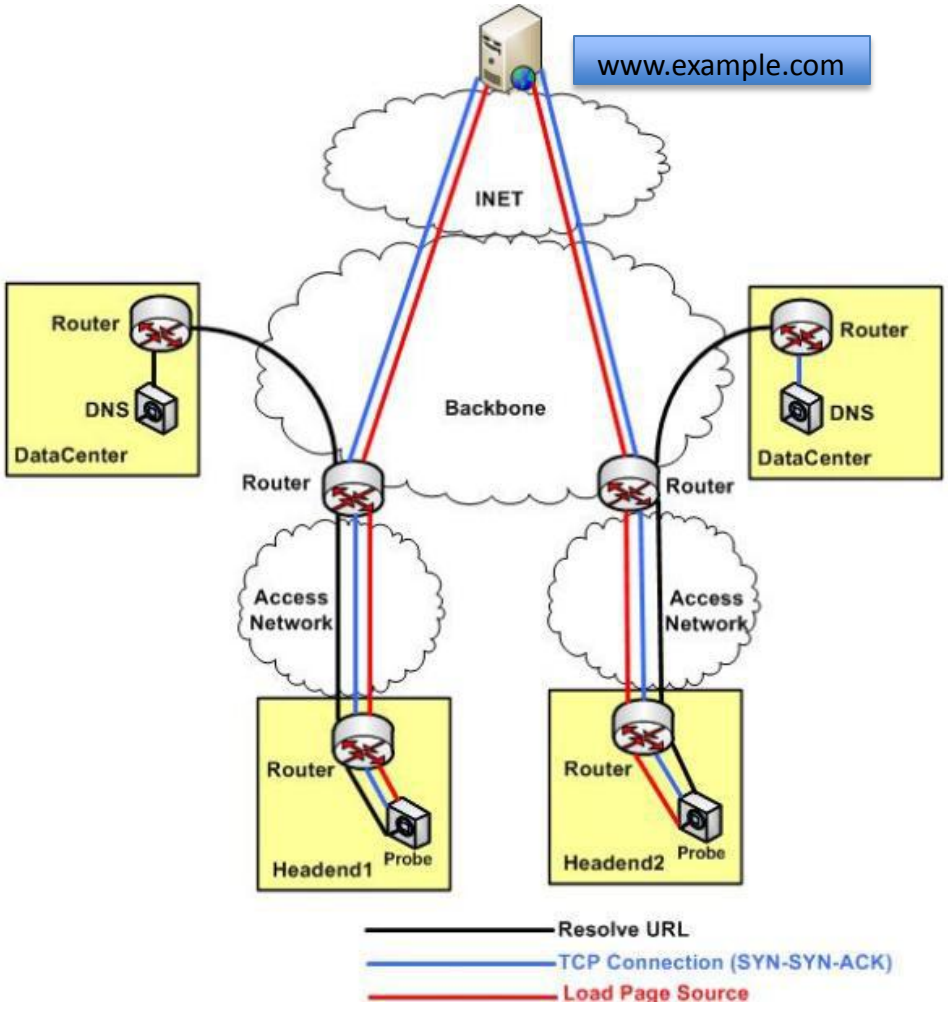

This chart shows the locations of the test probes in our network.

The probes run Unix Curl tool to initiate HTTP access to pre-configured websites.

The communications involved in the website access are shown separately in this diagram.

1. Black path for DNS resolution
2. Blue path for TCP 3-way handshake
3. Red path for web page load

Most of our focus is on TCP RTT

Latency Monitoring Diagnostic Systems

Website Latency Analysis (WLA) System

- Database and analysis system for probe-website RTTs
- RTT data every 5 minutes from 21 probes to several hundred websites
- Visual representations for high level and detailed views
- Provides thresholding/alerts for troubleshooting

Traceroute Analysis System

- Detailed traceroute (every hop) data for all probe-website pairs
- Traceroute data every 15 minutes; data retained for 2 weeks
- WLA identifies high RTT cases; Traceroute Analysis system helps troubleshoot

RTT Monitoring and Analysis Systems

WLA Screen

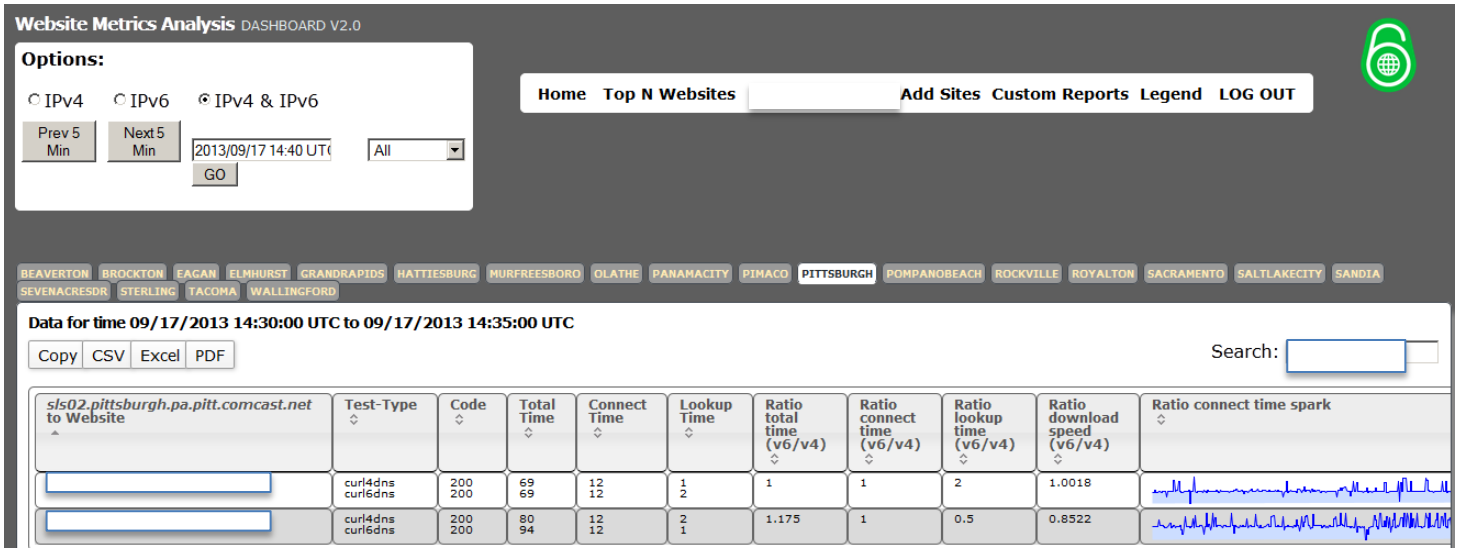

Traceroute Analysis Screen

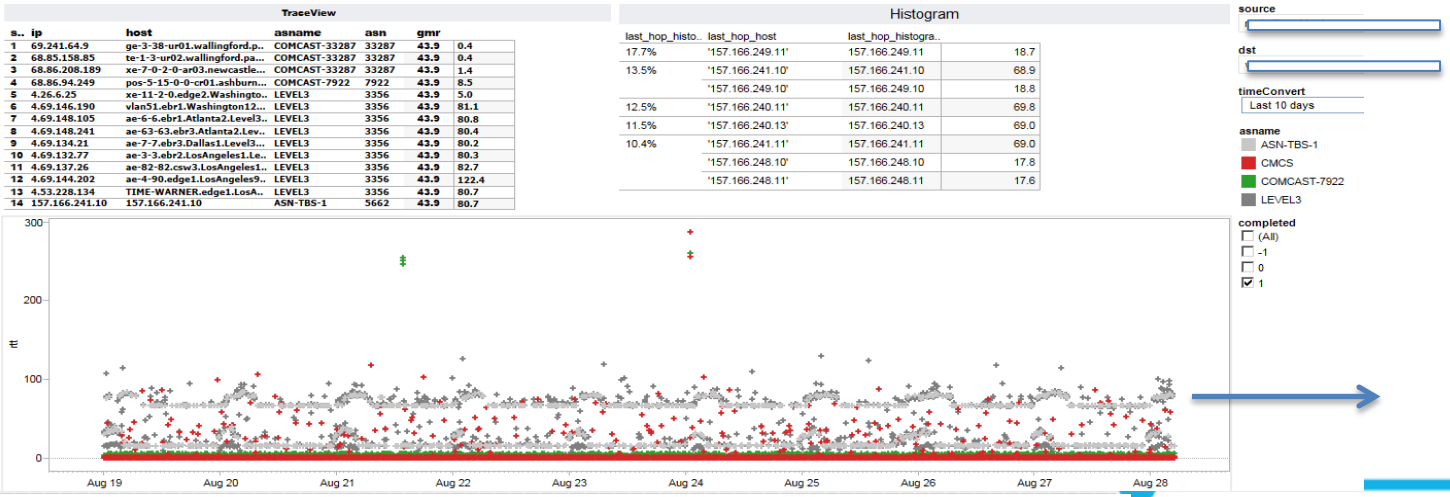

RTT Samples

Characterization of Comcast RTTs – IPv4

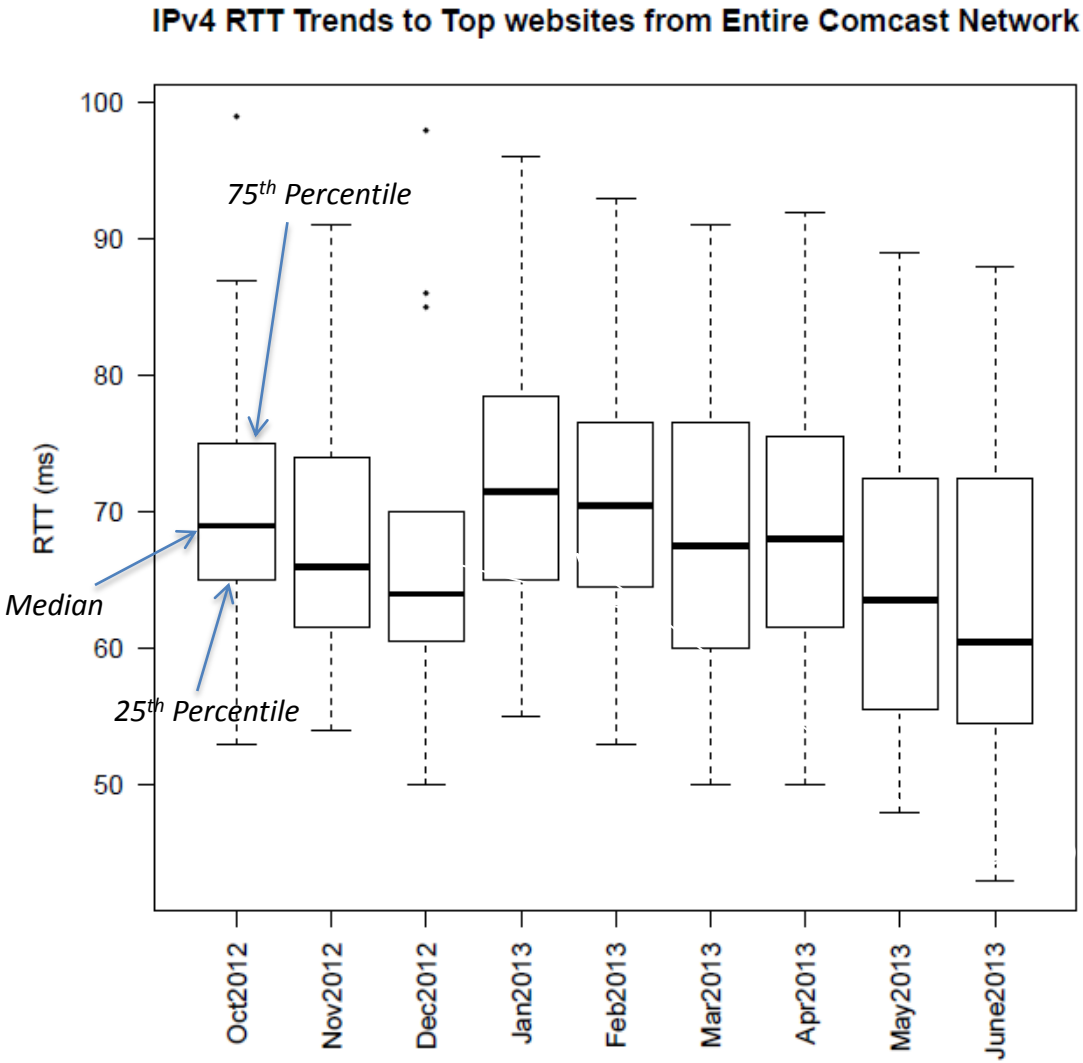

RTTs chart based on measurements for about top 25 websites

Each box represents spread of measurements from our probes at different locations.

RTTs have not changed significantly since 4Q2012 although median is down by ~15%

Characterization of Comcast RTTs – IPv6

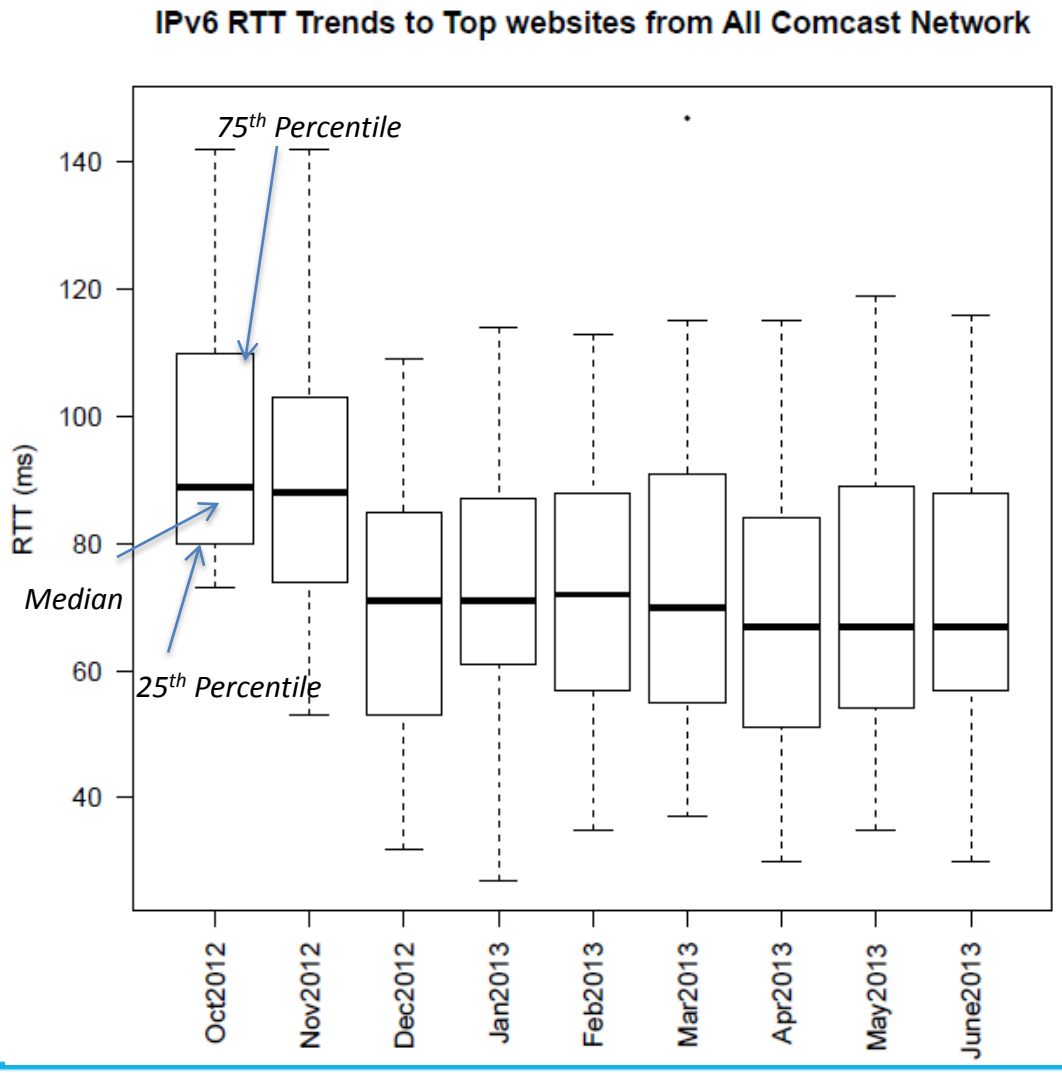

RTT data based on measurements of about 25 IPv6 websites

RTTs have come down by over 20% since last October

We attribute this reduction to Comcast's work with high RTT cases based on WLA analysis.

Comcast Drivers for Reducing RTTs

Comcast drivers for this project are:

- ❑ RTT is an important component of overall page download time; so reducing RTT would result in improved QoE for Comcast's internet service customers
- ❑ High RTT often indicates inefficient use of resources inside of Comcast or outside; reducing RTT would help fine-tune such inefficient configurations
- ❑ Website RTTs helped Comcast understand/compare CDN RTT performance changes

High RTT cases typically fell in the following categories:

- ❖ Website Provider's algorithm for assigning customers to servers not adequate
- ❖ Website providers use criteria (e.g., load) other than RTT in assigning customers
- ❖ Components involved in serving web content have inadequate resources or non-optimized configurations
- ❖ Smaller websites may not have multiple server locations
- ❖ Occasional mis-configurations

Example of Multiple Server Locations

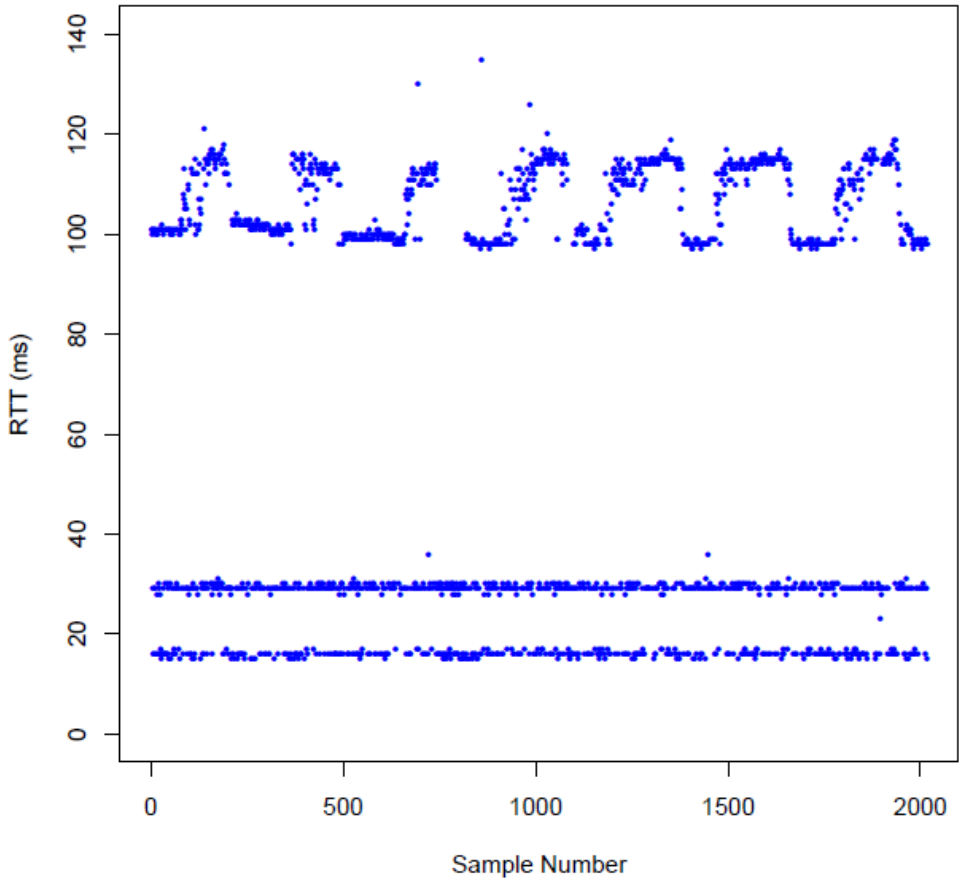

This website is served out of 3 different locations

RTTs from 2 locations quite low

But third location has high RTTs (>100ms) and a clear diurnal pattern indicating congestion during busy hours.

Example of RTT Reductions

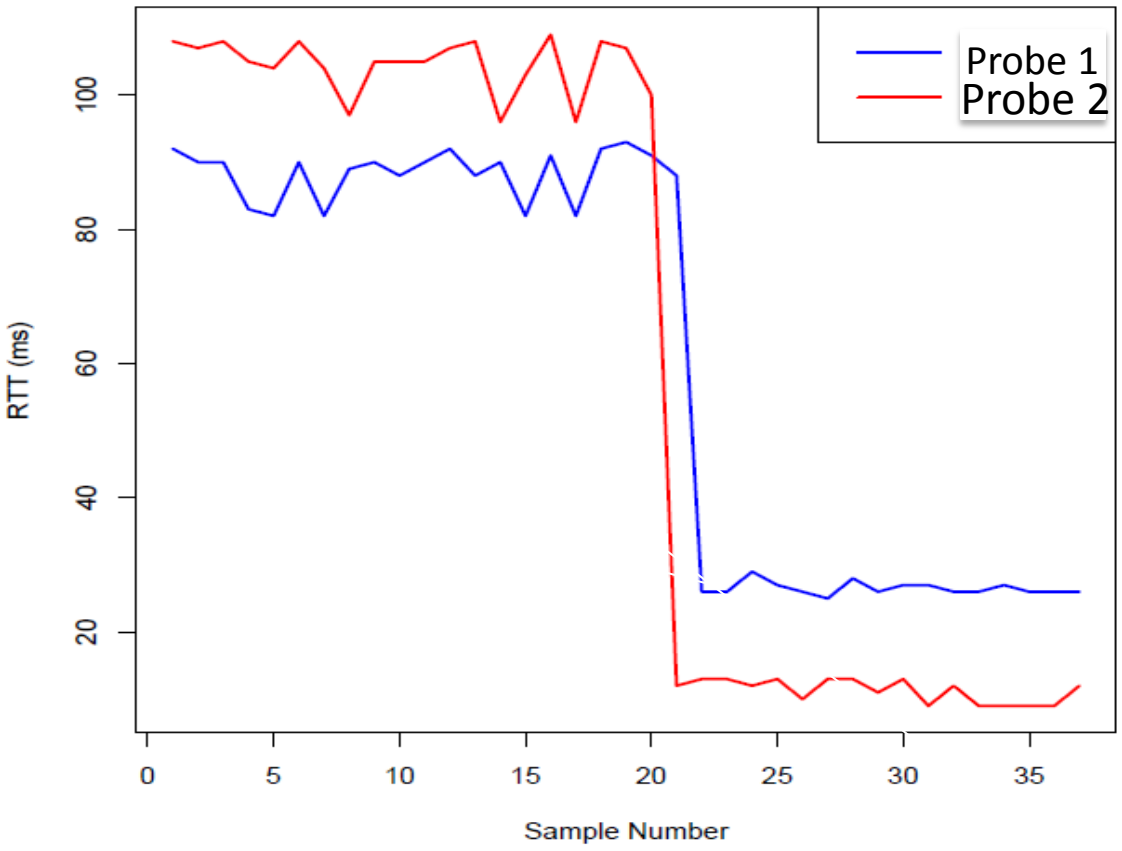

WLA system detected high RTTs to a popular website from several east coast probes.
Traceroute analyses showed that east coast customers were being served out of west coast servers.

Inadequate user-server mapping algorithm.
Website provider updated its mapping algorithm using information provided by Comcast.

Almost 80% reduction in RTTs

Example of RTT Reductions

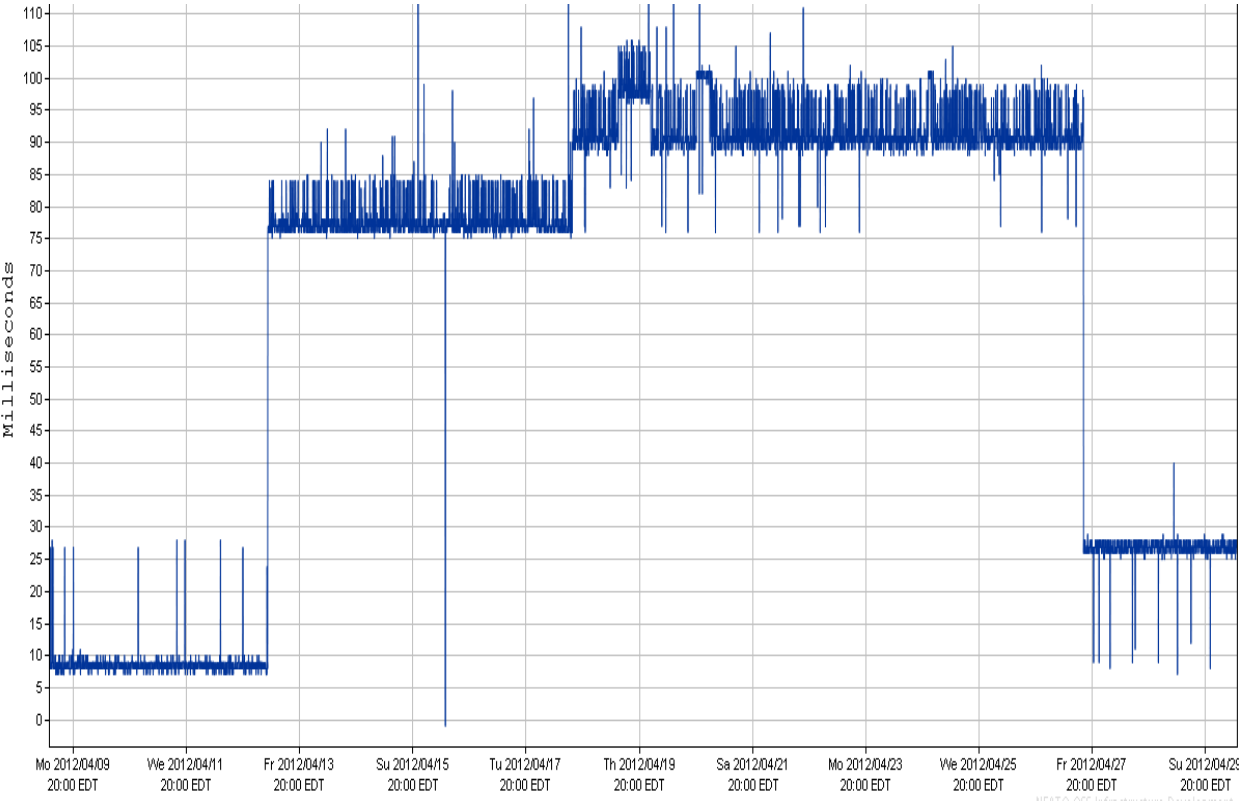

Comcast split an existing AS into two.

Website provider did not recognize new AS; so served those customers from a non-optimum default server.

WLA detected the issue

Comcast informed website provider and high RTT was quickly mitigated.