IPv4 & IPv6 Latency Analysis for Faster Web Browsing and Network Tuning

By:

Rich Andreas, Salil Banerjee, John Brzozowski, James Ehlinger, Eric Frishman and Tony Tauber

Comcast Cable

Introduction

This paper discusses the methods and results of Comcast's efforts in monitoring latency metrics that impact customers' web browsing experience. Customer experience with web browsing is dependent on a variety of factors that include not only performance of the Comcast network but also the number of HTML objects, number of servers, number of TCP connections, characteristics of web browsers and mechanisms used for content distribution. The primary focus in this paper is TCP connection setup time or TCP RTT. Every TCP connection or every HTTP GET operation may involve the equivalent of one or more TCP RTT delays. Thus, TCP RTT has a multiplying effect on the overall delay in accessing websites. Another benefit of focusing on TCP RTT is that the process of mitigating high TCP RTTs oftentimes improves network and routing performance and efficiencies through network and route "tuning".

Comcast Measurement and Analysis Infrastructure

Comcast has deployed a large number of probes across its network which use CURL and CURL libraries (CURL is a command line tool for getting or sending files using the URL syntax) to access a large number of websites at regular intervals and provide measurements of various parameters including TCP RTT between the Comcast probe and the server hosting the website. The same Comcast probes also initiate traceroutes to all the websites and record the traceroute data for each probe-website pair. Similar probes that run within embedded and standalone computing systems have also been developed and will too be used to diversely measure various network attributes and properties.

The Round Trip Time (RTT) data is stored in a database and analyzed by a system called the Website Latency Analysis (WLA) system. The WLA system provides capabilities to view that data in various different ways: by simply probe-website pair over different intervals, by CDN hosting the website, and, by deviation as defined by Comcast. Using these various means, Comcast is able to identify the probe-website pairs, which are having persistent high RTT values, which in turn may result in an impact to customer experience.

Once the high RTT cases have been identified, by using the WLA system, Comcast uses its Traceroute Analysis system to diagnose the issue and identify the root cause of the problem. The Traceroute Analysis system database stores details of each probe-website traceroute (including all the traceroute hops), and the Traceroute Analysis system is able to review the past data and identify whether the increased RTT is within the Comcast network or outside, exactly at what hop in the traceroute does the RTT go up, the different servers which served the website and whether high RTT was contributed by only one of those servers.

Characterization of Comcast TCP RTT Data

Before starting our discussion of specific examples of how we used our RTT measurement systems to solve real life latency problems, this section provides an overall statistical characterization of the RTT data for IPv4 and IPv6. Figure below shows the cumulative distribution function (CDF) of the TCP RTT data for accessing several hundred IPv4 and IPv6 websites from Comcast probes distributed across the Comcast network.

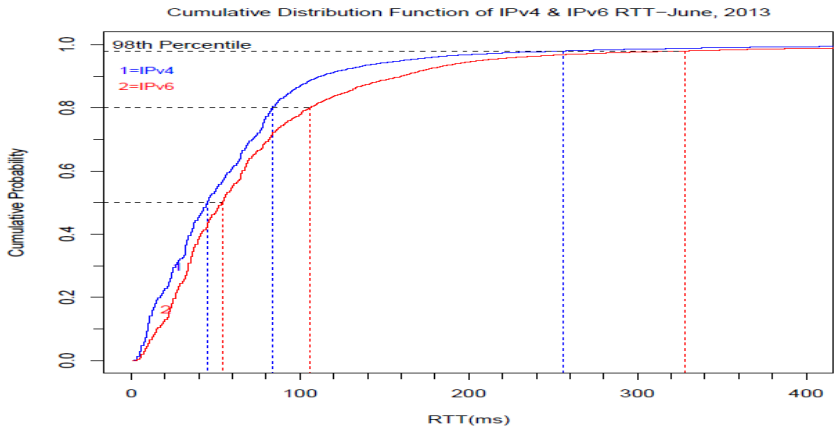

80th percentile RTT is just around 100ms, slightly less for IPv4 and slightly more for IPv6. We routinely analyze RTT data from probe-website pairs which were exhibiting higher RTTs and use our Traceroute Analysis system to identify the root cause and follow up with stakeholders as needed to reduce the RTTs.

Typical Cases of High RTTs

We have noticed a few different types of cases which exhibited high RTTs.

1. Website provider has multiple serving locations but has not optimized assignment of customers to servers. RTTs could be reduced in these cases by improving their server assignment algorithms, or by explicitly including Comcast resolving DNS server IPs and geographical locations in their algorithms.
2. Website provider has multiple serving locations but assigns customers to servers based on criteria (like, load balancing at high-utilization times) other than RTTs. RTTs could be reduced in these cases by improving the assignment algorithms used by the website provider. An example of this is shown below where the website is served out of two different locations and the two bands of RTTs correspond to the two different server locations that the customer gets served out of, typically during different times of the day.

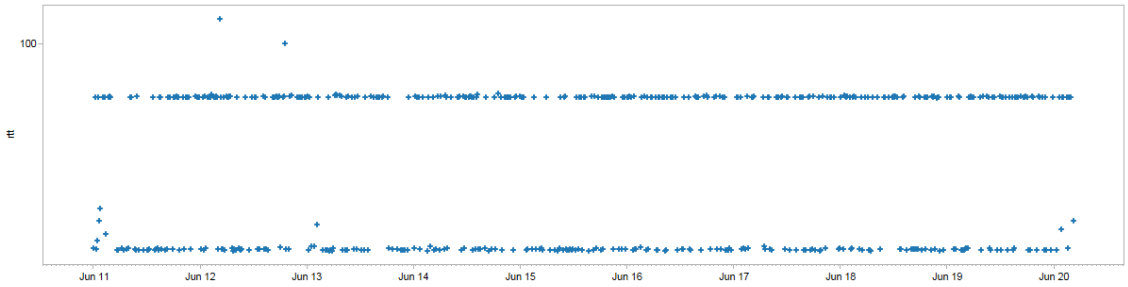

3. The components involved in serving the website have either inadequate resources or non-optimized configurations. This may include inadequate server capacity, inadequate bandwidth somewhere in the connecting path or other mis-configured components.
4. The website has only a single serving location (mostly for relatively less popular websites); in these cases reduction in RTTs is dependent on business decisions by the website provider.

Examples of Using RTT Data to Reduce Latencies for Comcast Customers

Example 1

The WLA system detected persistently high TCP RTT for a certain highly popular website (say, www.example.com) from a number of probes located in the eastern or the central parts of the country. Traceroute analyses revealed that Comcast customers from the eastern part of the country were accessing example.com servers in the western part of the country resulting in high latency.

Comcast worked with the engineers at example.com and they attributed the problem to incorrect user-server mapping used by example.com's authoritative DNS servers. Example.com updated their mappings using information Comcast provided to them containing a list of IP addresses of its resolving DNS servers and their geographic locations. This resulted in about 80% reduction in RTT from a number of our probes.

Example 2

The WLA system detected a number of websites hosted by a commercial CDN experiencing rising TCP RTTs from 4 probes in the southwestern part of the country, particularly during busy hours. Comcast worked with the CDN provider and the root cause was attributed to inadequate disk capacity and server CPU for the resources serving the specific locations. When additional resources were deployed by the CDN provider, we noticed almost over 50% reduction of the RTTs.

Conclusion

Comcast has developed and deployed a measurement and analysis infrastructure to monitor RTTs to hundreds of popular websites from many parts of its network. Comcast identifies cases of high RTTs, analyzes its root causes working with its own tools as well as working with other stakeholders, and tries to mitigate the high RTT issues. Comcast welcomes opportunities to work with all parties to reduce Internet RTTs in order to provide a better experience for its customers.