Internet Measurements

Data Center interconnection analysis

By:
Rezfan Pawirotaroeno (University of Twente)
Ferdinand Kwarteng (Brandenburg University of Technology)
Philip Doku (Brandenburg University of Technology)

MAY 28TH 2021

Group Codename - India

Delay and hop count between hosts

Enschede and Bangalore - IPv4

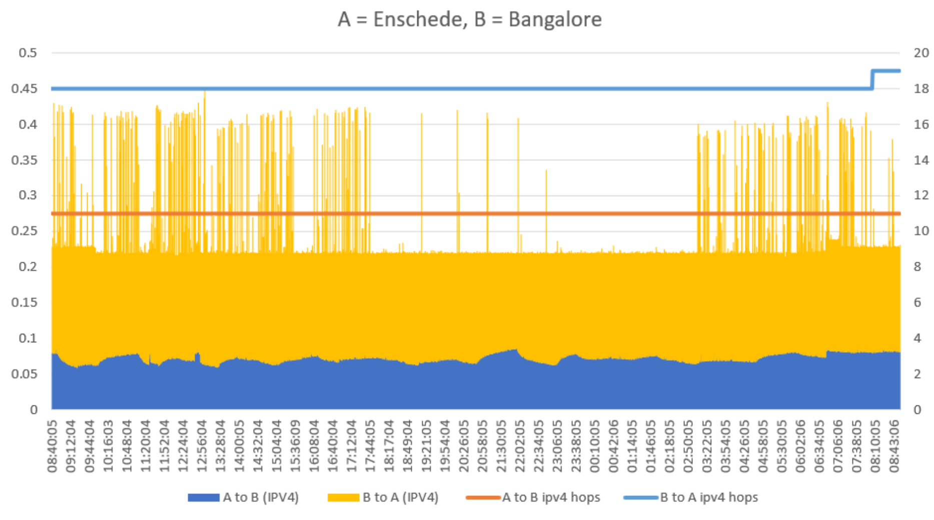

Enschede and Bangalore - IPv6

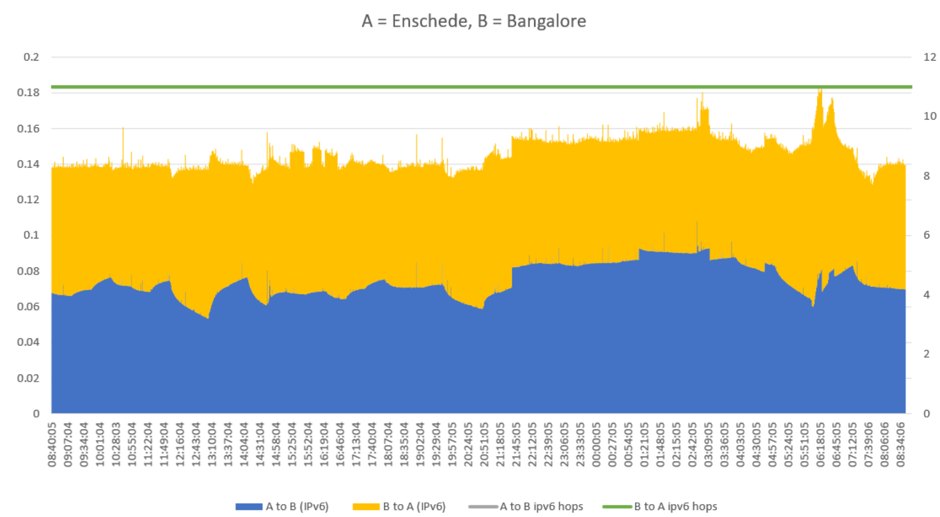

Enschede Bangalore Discussion

For the IPv4 interconnection between the server from Enschede and Bangalore, we generally observe significant increases at one of the routes at certain time intervals. Regarding the path symmetry, it is clear that the paths from Enschede and Bangalore are not very symmetrical regarding the delays between the two servers. The path from Bangalore to Enschede takes significantly longer than the path from Enschede to Bangalore. Additionally, regarding the delays it seems to be much more volatile. Even though the path from Enschede to Bangalore is much shorter than the other way around, it is just as volatile.
Even though there is a lot of change regarding the delays between the two hosts, there is for a significantly large majority no change in the number of hops between the two hosts. This indicates the possibility that one or more intermediate routers in the path are causing an increased delay.
Furthermore, even though the delays between both hosts are highly volatile, we see that most delays are around 0.23 seconds. This indicates a mechanism that ensures a consistent RTT between the two hosts.

For the IPv6 interconnection between the server from Enschede and Bangalore, we also observe that these delays are highly non-symmetrical. And there is no mechanism that ensures RTT consistency. Whenever one delay (from host A to B) increases significantly, the other one (host B to A) does not decrease to ensure around the same RTT.

Enschede and New York - IPv4

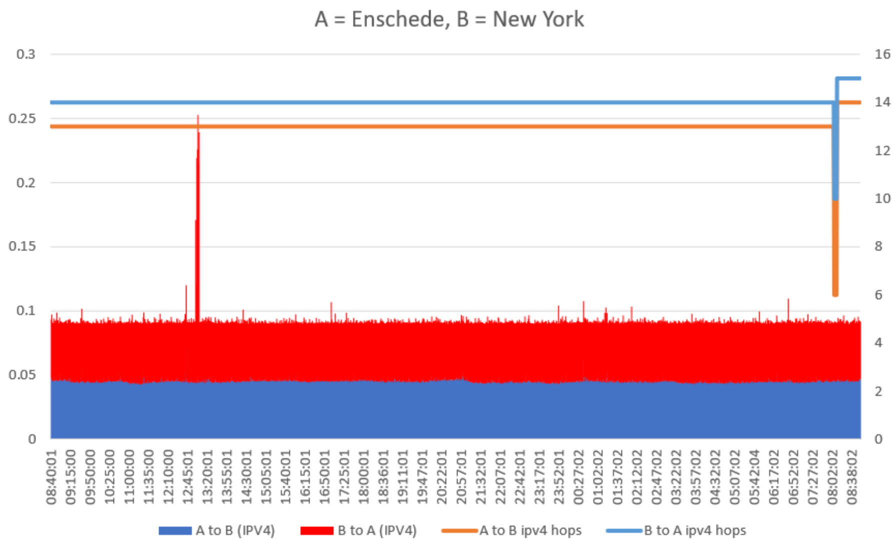

Enschede and New York - IPv6

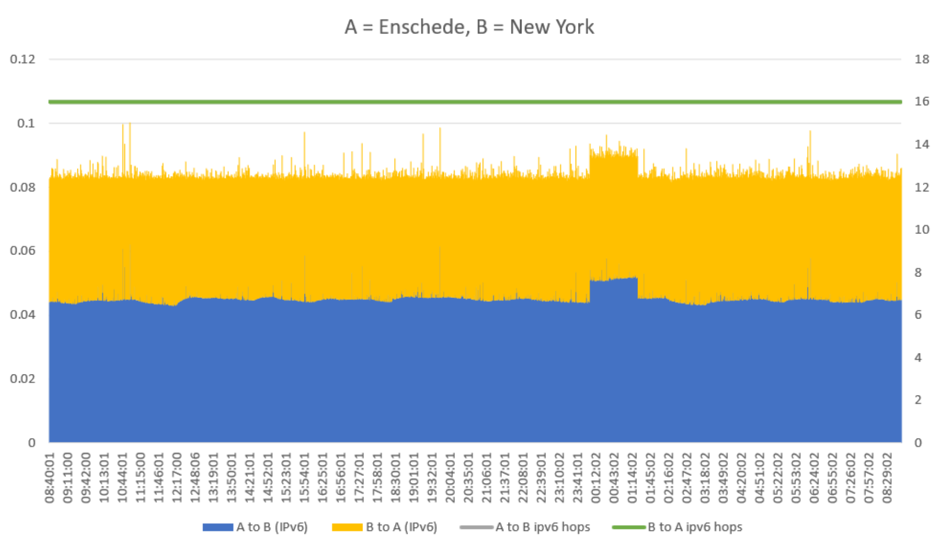

Enschede - New York Discussion

For the IPv4 case, we observe that for most of the time the delay between Enschede and New York is highly the same. For this case you can say that the delay in path is relatively symmetrical. However, we do notice a significant spike between 12:10 and 12:45. Additionally, we observe that there is no decrease in the delay from B to A to compensate for this. This could mean that these two hosts do not have a mechanism in place that ensures consistent RTT like the other ones. Between 08:02 and 08:30 it is observed that there is a large decrease in the hop count, however it seems that there is no decrease or increase in the delay between the two hosts.

Regarding the delay in the IPv6 space, it is observed to be very similar to the IPv4. Regarding the height of the graph, one can observe that the two heights (A to B and B to A) are relatively similar in size, meaning that the delay from A to B and B to A are very similar. In other words, the delays are very symmetrical. However, from around 00:12 and 01:45 you see there is an increase in delay for the Enschede and New York interconnection. And you see that the total RTT also because the delay from New York to Enschede does not decrease. This indicates that the IPv6 space also does not have a mechanism in place to ensure RTT consistency.

Enschede and Nurnberg - IPv4

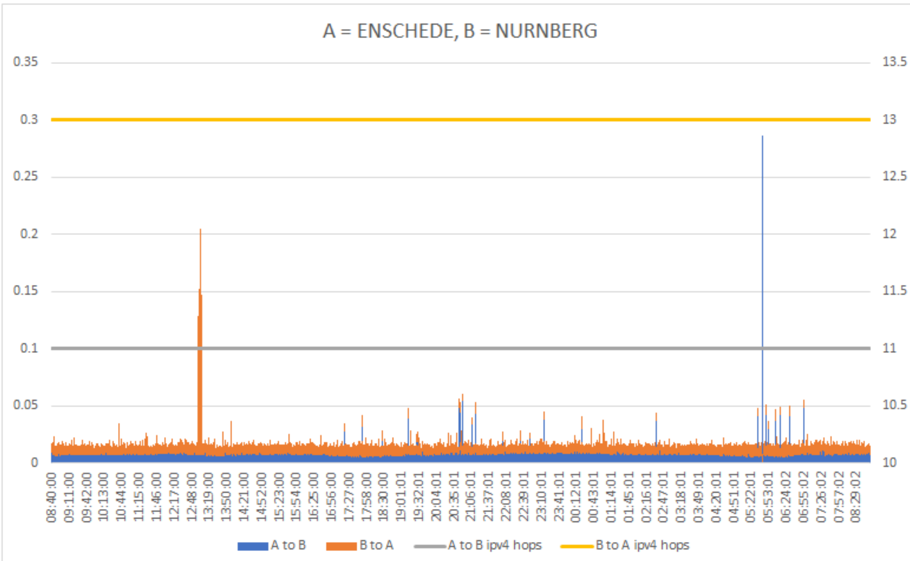

Enschede and Nurnberg - IPv6

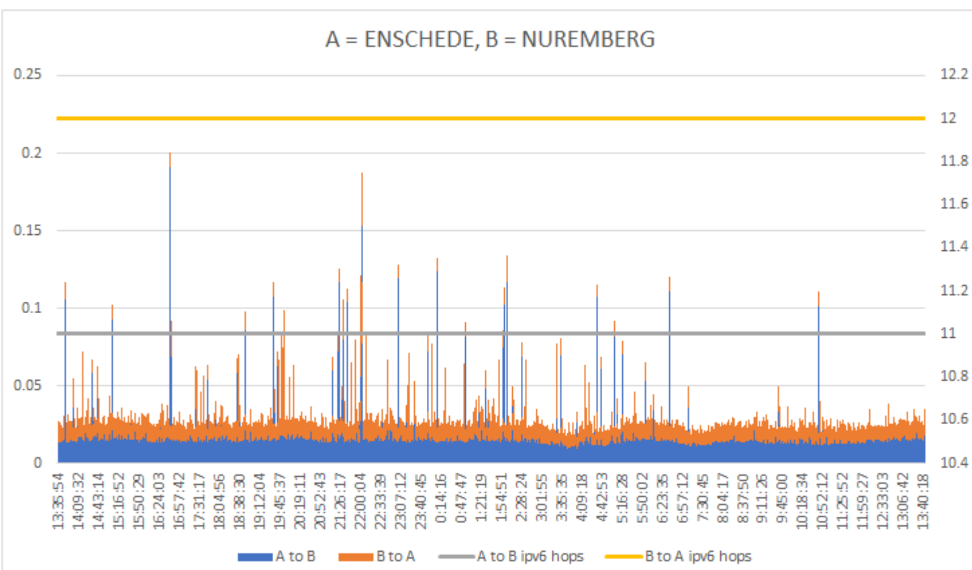

Enschede and Nurnberg Discussion

For the IPv4 interconnection between Enschede and Nurnberg, we observe that there is clear no-path symmetry between the two hosts. One delay seems to be significantly larger than the other. Furthermore, we also observe no mechanism for RTT consistency. And similar to one of the previous graphs, we also observe a sudden spike in delay around 12:48 in time. The number of hops from A to B is smaller, and the average delay from A to B also seems to be shorter. The same goes for the B to A case, there the number of hops seems to be larger and the delay is larger as well.

Regarding the IPv6 interconnection delay between the two hosts, we see similar results. Compared to its IPv4 version, the delays between the two hosts seem to be a lot more volatile. Lots of larger spikes seem to come mostly from the connection from Enschede to Nurnberg. Do note that for the IPv6 case the measurement was taken at a different time than the others due to the first collected set having issues regarding the way the code parsed the csv’s. The code was then run several hours later to properly collect the data from the two hosts again.

Frankfurt and New York - IPv4

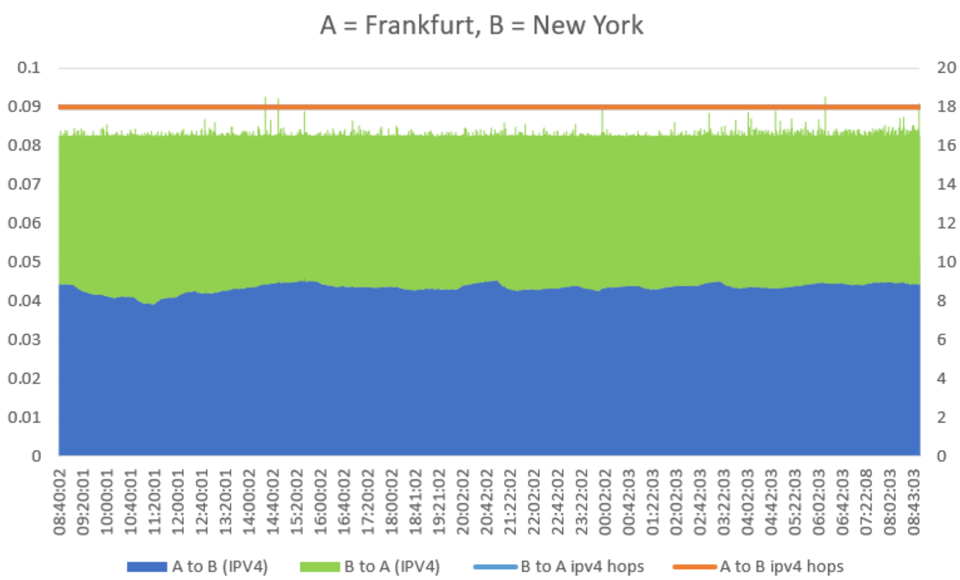

Frankfurt and New York - IPv6

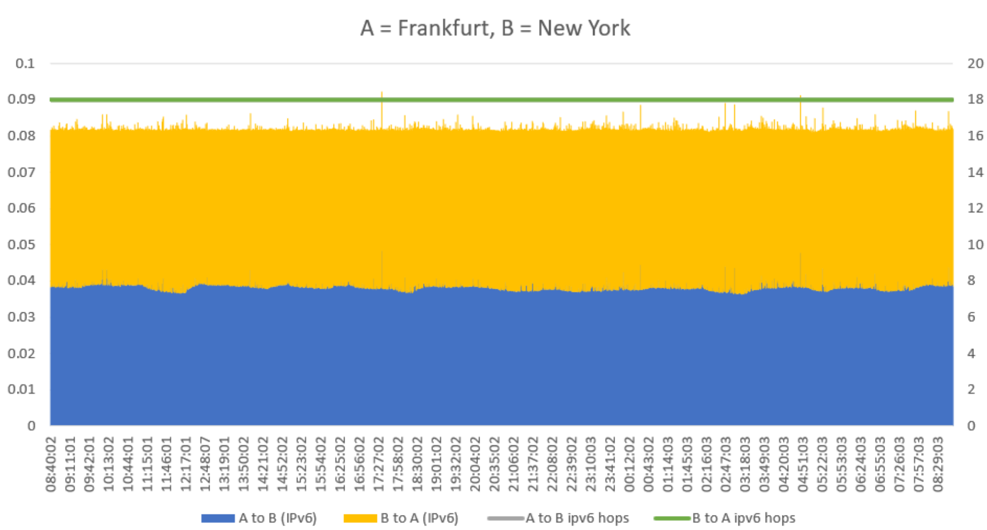

Frankfurt and New York Discussion

From the graph, it is noticed that there is a relatively even latency between the servers from Frankfurt and New York for IPv4 connections at the different time intervals. It is then possible to conclude that there is a symmetry between the connections. However, there are slight spikes from New York to Frankfurt at some irregular times. In addition, it can also be observed that there is consistency in the RTT delays in both directions. A change in RTT from one direction does not affect the other. Regarding the hop counts and intervals, it is also symmetrical. There is an even distribution in the hops.

Similar to the IPv4 graph, the IPv6 graph has symmetry in the latency in both directions because the delays are averagely in the same time limit. Hop counts are also the same. There are some slight spikes in the from New York to Frankfurt as observed in the IPv4 graph. Regarding the RTTs, there is consistency as well.

Enschede and San Francisco - IPv4

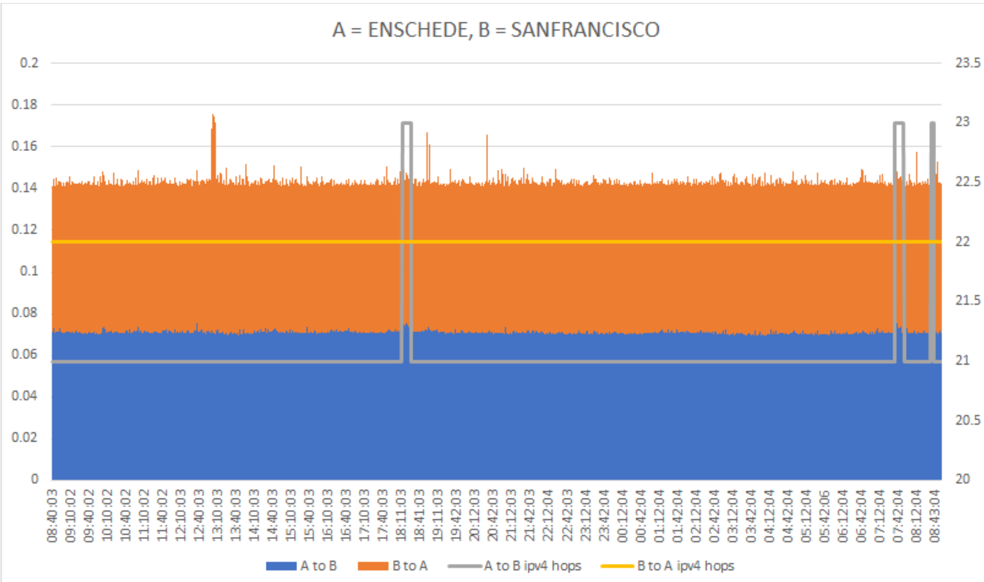

Enschede and San Francisco - IPv6

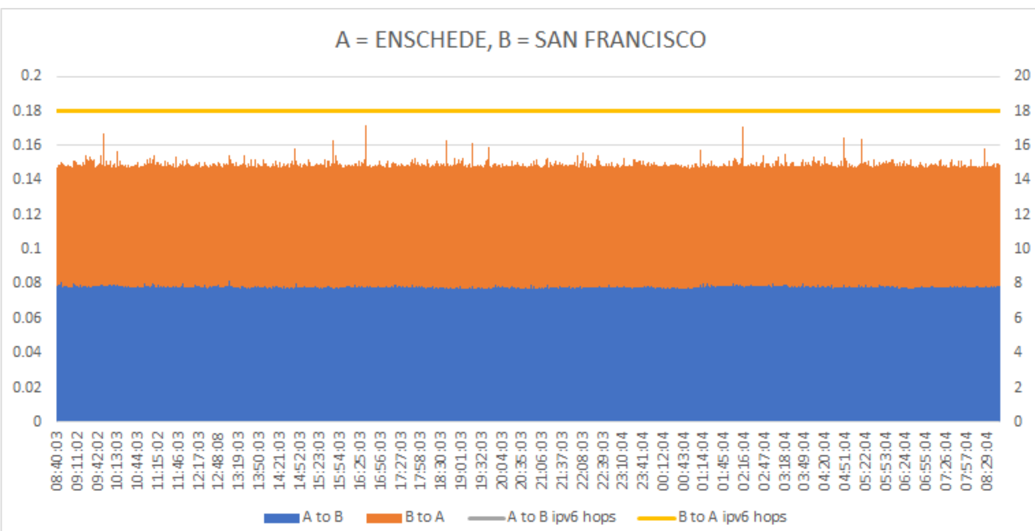

Enschede and San Francisco Discussion

For the case of Enschede and San Francisco, we see in both IPv6 and IPv4 that the paths are highly symmetrical, because both the delays are around the same time. For the IPv6 the hop count is also the same for both cases. However, for the IPv4 case, we observe several significant increases in the hop counts from Enschede to San Francisco. However, no further change is observed regarding the delays as a result of this. It is difficult to determine whether these two hosts have an RTT consistency mechanism in place, due to both delays being very similar and non-volatile/consistent.

Singapore and San Francisco - IPv4

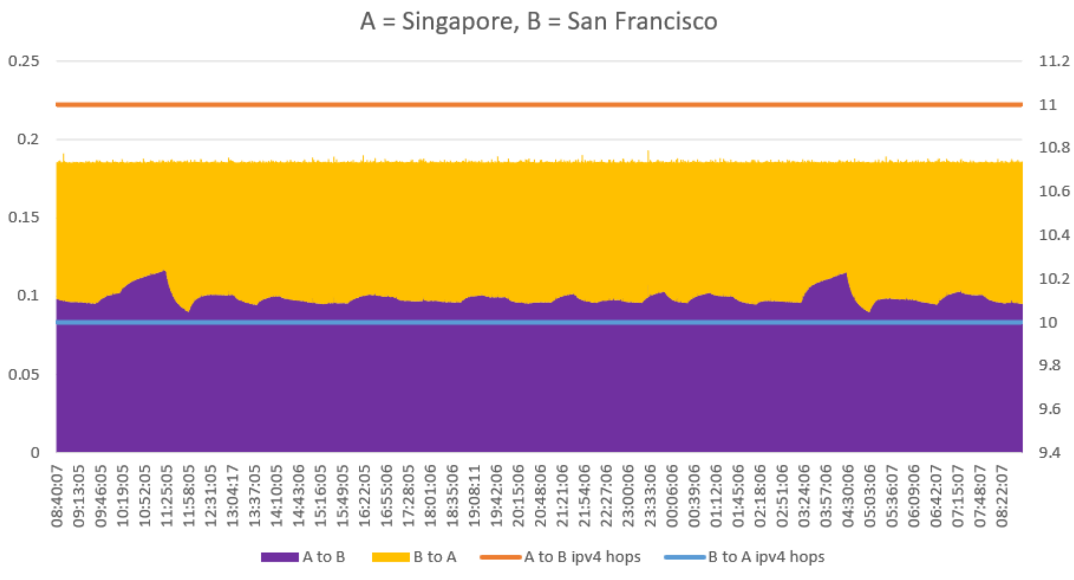

Singapore and San Francisco - IPv6

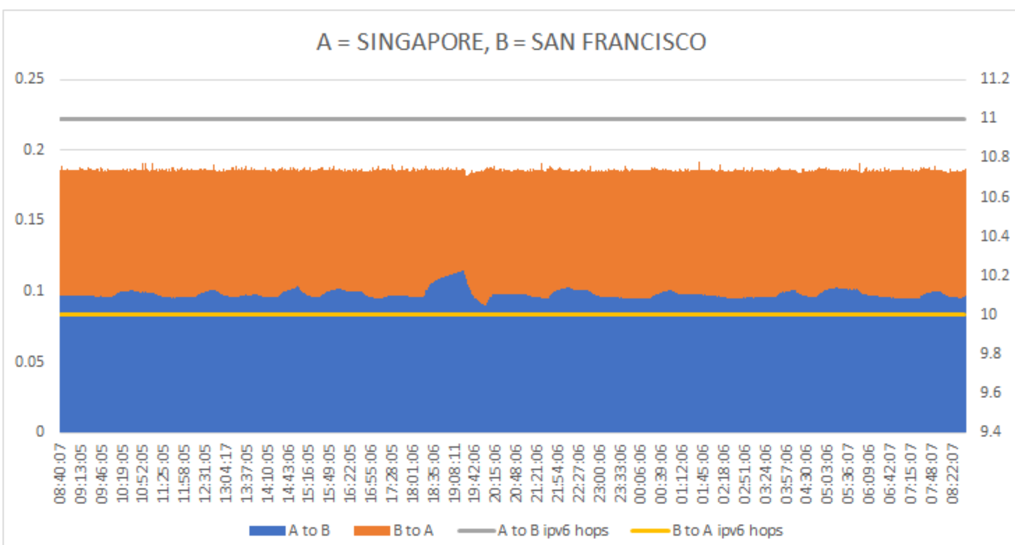

Singapore and San Francisco Discussion

From the interconnections from the hosts Singapore and San Francisco, it is clear that both IPv4 and IPv6 have mechanisms in place to ensure RTT (summed delay) consistency. An increase in one delay results in a decrease in the other delay to ensure that the summed delay is around the same time. Furthermore, it can be seen that the two delays are around the same size, even though the hop counts are different. There seems to be no changes regarding the hop count and no correlation between the slight changes in host to host delay times.

Enschede and Cottbus - IPv4

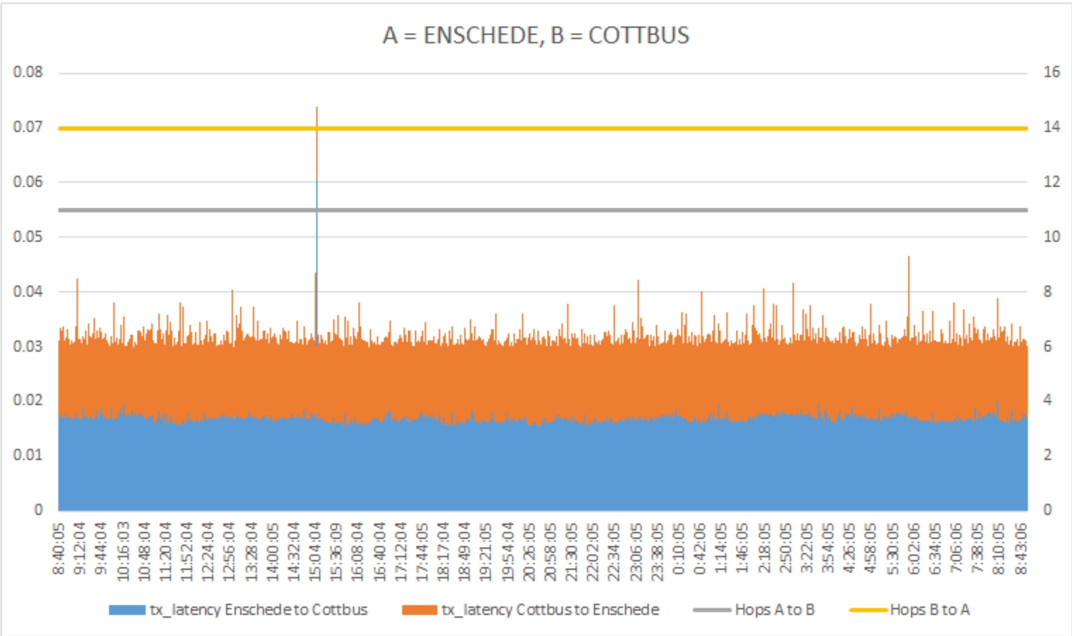

Enschede Cottbus Discussion

From the graph it is obvious to say there is consistency in the hops from both directions. In addition, there are evenly slight spikes in the latencies and at least one sharp spike around 15:04 and this was for a short period of time. This may be due server latency times. There seems to be no implementation of a consistent RTT because there is an irregular change in RTT from both directions.

General discussion

From the different interconnection results, it is observable that there is a clear distinction between servers that are symmetric in their delays and those that are not. Additionally, it is observed that some of the host pairs show mechanisms that try to keep the RTT consistent between the two hosts.
It seems that for a lot of the plots of the pairs from which one is Enschede, there is a sudden increase in delay (as observed by Enschede - San Francisco IPv4, Enschede Nurnberg IPv4, Enschede - New York IPv4), at a specific time. Because Enschede is in all of these, it would be a reasonable assumption to conclude that the server from Enschede either has problems or the routing around Enschede might have some issues that are causing these delays. Furthermore, we observed that the hop count stayed consistent for the majority of the measurements, and observed no correlation between hop count and delay. An increase in hop count and increase in delay would seem trivial, however there was an increase in delay but no increase in hop count, meaning that this spike in delay has to be a result of one of the intermediate routers on the path.

Furthermore, the interconnections between Enschede and Bangalore need to be looked at for the reasons behind those large spikes in delays.

For the measurement intervals (times at which data was logged), 1 minute was chosen. This interval of 1 minute was chosen to get a general yet detailed enough overview of the data. Initially, we observed that the data from sping files of the servers changed multiples times every second, whereas the traceroute files changed around every 2 minutes. To meet somewhere in the middle, we chose 1 minute. Do note that this can also function as a limitation that introduces errors into the research. When collecting the data every 1-minute intervals, one might miss out on a possible event that happened in between these minutes.

The first possible factor that could introduce measurement errors are the personal internet connections. The computers used to acquire the data might also suffer from increased delays when requesting data from the server.
The second possible factor that could introduce errors could be the time it takes for the servers to log the data. For example, if a server is slow in processing the data into the retrievable csv’s, it would give its data at a certain time, however in reality that given information is already an extra 5 seconds older than the time it claims to be logged at.
Thirdly, the tools used to generate the data could also possibly provide artifacts. These could be delays in processing that cause the acquired data to be several milliseconds behind the ground truth.