Special Report: Network Latency and Bahrain's Content Ecosystem

Renesys Corporation
December 2012

Table of Contents

Executive Summary

Bahrain's consumers access Web content from a wide variety of sources, and are generally unaware of where in the world that content is actually hosted. Web content may be hosted domestically, within a facility operated by a Bahraini company. More often, however, the Web content viewed in Bahrain is served from a host that resides outside the Kingdom, or in fact, outside the Region entirely.

Content providers and content delivery networks are therefore important non-telecom actors within Bahrain's Internet ecosystem, and the paths used to interconnect to them are vitally important in determining the overall consumer experience.

This report describes the Internet paths along which Bahraini consumers retrieve popular content, using both fixed and mobile Internet devices. We examine the **latency**, or round trip time, of each path, which is primarily determined by the speed of light (path length) and secondarily by factors such as congestion.

Finally, we identify some possible ways in which content latencies could be reduced for Bahrain's consumers, and suggest future strategies for measuring and managing this progress. As Bahrain's access providers learn more about the latencies of their alternative paths to content, they may be able to guide their upstream providers in improving their customers' online experience.

Methodology

Renesys operates a global network of 90 Internet measurement points, from which latency probes are sent to over a million responding hosts each day, worldwide. Using the latencies measured to each responding host and router, we build up a global picture of the fast and slow paths that Internet traffic can take between clients and servers.

It's important to note that Renesys does not currently have a measurement point within Bahrain. To examine the characteristic latencies between Bahrain's consumers and the world, Renesys performed more than 30 million traces from worldwide cities to IP addresses located in Bahrain, and used the paths and round trip times to build a map of each provider's inbound connectivity. We used the primary hosting locations for a set of popular Web destinations (identified by Alexa.com) as a guideline for identifying performance-sensitive routes.

We plotted the main trends in Internet delay for ten key Bahrain ISPs, and identified the primary geographic paths they take to reach popular content locations (see Appendix A, "Provider Latency Measurements").

Geolocation of Key Domains (July 2012)

Most of the websites identified by Alexa.com as being popular in Bahrain are actually hosted outside the Kingdom. In fact, none of the top 25 sites in Alexa's ranking are even hosted within the Gulf region. Some content providers (such as Wikipedia, or LinkedIn, or the BBC) serve their primary pages from a consistent location, regardless of where the customer lives. Other providers (such as Yahoo/Flickr, or Google/Youtube) maintain sophisticated mappings between customers in particular locations, and the datacenter cities from which they are served.

This table shows where the "front pages" of seven popular sites were served from in July 2012, for customers of various Bahraini ISPs. To compose this table, each domain name was resolved locally to an IP address, and that IP address was geolocated to identify the nearest major city.

ISP	Yahoo	Youtube	Twitter	Wikipedia	LinkedIn	Flickr	BBC
2Connect	Dublin	Frankfurt	San Francisco	Florida	Los Angeles	London	London
Batelco	Geneva	Frankfurt	San Francisco	Florida	Los Angeles	San Antonio	London
Etisalcom	San Francisco	Marseilles	San Francisco	Florida	Los Angeles	San Antonio	London
Kalaam	Geneva	Frankfurt	San Francisco	Florida	Los Angeles	London	London
Lightspeed	Dublin	Mumbai	San Francisco	Florida	Los Angeles	London	London
Menatelecom	San Francisco	Marseilles	San Francisco	Florida	Los Angeles	San Antonio	London
Zain	San Francisco	Marseilles	San Francisco	Florida	Los Angeles	San Antonio	London
Batelco Mobile	San Francisco	Frankfurt	San Francisco	Florida	Los Angeles	San Antonio	London
Viva Mobile	Geneva	Marseilles	San Francisco	Florida	Los Angeles	London	London
Zain Mobile	San Antonio	Frankfurt	San Francisco	Florida	Los Angeles	San Antonio	London

Key:
European hosting; latencies 150+ ms
South Asian hosting; latencies 250+ ms
Eastern US hosting; latencies 250+ ms
Central US hosting; latencies 250-280ms
Western US hosting; latencies 290+ ms

Challenges of Bahraini consumer content mapping

Several patterns emerged from the geolocation of these content sites' primary IP addresses.

- Broadly speaking, the **most important geographic regions** serving content to Bahrain's consumers are Western Europe (Geneva, Dublin, London, Frankfurt, Marseilles), Central USA (Texas), and Western USA (Los Angeles, San Francisco).
- The **worst case scenario** in terms of latency for Bahrain's consumers is a website, such as LinkedIn, which serves content from the US West Coast. Whether Internet routes carry the traffic west or east around the world to Bahrain, typical latencies will be 280-320ms depending on the path. Paths from the US through Europe tend to be faster than equivalent paths from the US through Asia to reach Bahrain, but are not always used.
- Bahraini ISPs who connect to the Internet through Qatar Telecom, such as Batelco, Etisalcom Bahrain, and Zain, tended to be mapped to San Antonio, Texas to reach Flickr.com, while Saudi Telecom customers like Kalaam, Lightspeed, and Viva Mobile were more likely to be mapped to London.
- The only Bahraini provider mapped to a non-European caching host for YouTube was Lightspeed, whose FLAG transit presumably influenced the selection of Mumbai. All others were mapped by Google to Frankfurt (a major European exchange point city) or Marseilles (the landing point for several of the largest east-west submarine cables serving the Gulf region).
- Yahoo connectivity displayed the greatest variability. 2Connect and Lightspeed, both downstream of the Bahrain Internet Exchange, were mapped to Ireland. Batelco's fixed broadband customers were mapped to Geneva, as were Viva Mobile and Kalaam, both downstream of Saudi Telecom. Zain Mobile Broadband was the lone service to receive a Texas address for Yahoo. Zain's fixed broadband, Etisalcom, Menatelecom, and Batelco mobile broadband were all mapped to Yahoo in San Francisco.
- As a result of this wide variance, it's possible that customers of these different providers may experience different performance when visiting Yahoo's web properties.

It's important to note that these observations were based on mappings observed on a single day in July 2012. The mappings are certain to change over time, as content providers react to network loads by shifting traffic and remapping clients. However this highlights the existence of large path-based variances between ISPs.

Even more importantly, these are only the “front page mappings:” the geolocation of the IP addresses assigned to the home domain. Much of the heavier content delivered by these sites will arrive from partner CDNs such as Akamai, whose servers have not been geolocated in this initial study. This additional level of CDN analysis could be conducted in order to complete the picture of consumer-visible latency. For example, many popular websites in Bahrain (including Twitter, Facebook, Apple, and Microsoft’s Bing) utilize Akamai caches in Singapore, the UK, the Netherlands, and the US to serve content to Bahrain’s consumers.

Strategies for Improvement

Different ISPs naturally have different latencies to popular content. Much of this is the result of their specific Internet transit and peering arrangements. Globally distributed content providers may choose to map Bahraini users to different content service locations, and some even change those mappings frequently.

Some of the latency variance, therefore, may be under providers’ control, if they carefully monitor the paths between their customers and international content. When a Bahraini provider’s best path to the US West Coast goes via Mumbai, Singapore, Hong Kong, or Tokyo, for example, they may wish to coordinate with the corresponding global carrier (e.g., FLAG or Tata) to move traffic back onto a lower-latency European path.

Similarly when traffic is routed along high-latency paths unnecessarily (for example, inbound traffic from London to Batelco travelling to New York before being returned to Bahrain in an MPLS tunnel)[1], consumers may experience significant performance degradation. New measurement tools (like those offered by Renesys) can make these suboptimal paths visible to NSPs so that they can improve customer experience.

Another, complementary strategy to improve consumer Internet performance would be to improve the placement of content caches, thus bringing popular content closer to the Gulf region. In the absence of better local caches in nearby cities such as Riyadh, Jeddah, and Dubai, users in Bahrain may experience significant delays fetching content from more remote locations. Simple speed of light delays of hundreds of milliseconds are unavoidable over such distances, compounded by the fragility and potential congestion of providers’ routes to international locations.

The complexity of the Internet performance challenge calls for careful measurement and mapping of the nation’s entire content delivery ecosystem. Because paths and performance to content can vary so widely among NSPs, measurements would

[1] In the first days of November 2012, some measured path latencies between Jeddah and Batelco increased significantly, an effect that turned out to be due to the impact of Hurricane Sandy on the New York region.

optimally be performed within each fixed-line and mobile broadband Internet provider's infrastructure. At a minimum, the challenge requires identification of the primarily domain names for popular content (including CDN resources used to host content), hourly or daily DNS resolution of those names to IP addresses, and recurring active measurement of latencies to those IP addresses. Combined with global analysis about the paths traffic follows, such a measurement effort would provide valuable information about the delays consumers are likely to encounter when accessing popular content. It would also give daily feedback to NSPs about the causes of high latency to content, and suggest alternative strategies.

The basic shape of the challenge is clear: Internet consumer experience within Bahrain is currently strongly shaped by Bahrain's ISP performance to a diverse array of remote locations around the world, from Europe to Southeast Asia and across the United States. Until content comes closer (through improved regional hosting), Bahrain's consumers will continue to experience high latencies and significant variance, depending on which provider they use. Those providers may be able to influence international path selection to popular destinations, thereby improving Internet performance for Bahrain's consumers.

Provider latency measurements

Each of the following subsections shows the 2012 history of primary latency modes measured from a major content city into a given Bahraini provider. These plots are based on direct measurements from the cities in question, identifying up to 4 primary latency modes per day over the course of a year from London (Western Europe), Dallas Texas (Central USA) and Palo Alto California (Western USA).

Intuitively, one might expect to see a single, relatively constant round trip delay for traffic between a given world city and an ISP in Bahrain. The actual picture is significantly more complicated. Most Bahraini ISPs have multiple Internet transit providers, and most of those transit providers have multiple strategies (physical paths, such as submarine cables or terrestrial connections) for reaching far-away providers.

As a result, the typical latencies between a given Bahraini ISP and a remote city will tend to have multiple "modes" -- characteristic levels of delay – corresponding to the various paths traffic can take to reach the destination. This "delay floor" is based on the speed of light, but it can also vary continuously, as routing fluctuates and the length of the paths between Bahrain consumers and Europe/US web hosting changes in response.

Compounding this measurement challenge are the add-on effects of link congestion, which can cause latencies to further rise and fall with the time of day. When some long-distance links become congested, global NSPs may choose to reroute traffic to less-congested, lower-cost routes without notice. In many cases, this explains the emergence of higher-latency "backup paths" from the USA to Bahrain through Hong Kong, Singapore, and Mumbai. Without careful measurement, local providers may not be aware of the selection of these long paths and their effects on customer experience.

It's also possible to see the effects of latency in the generation of "outlier" points: occasional latency measurements that are significantly higher than the simple speed-of-light path length would suggest. When these outlier points tend to occur at particular times of the day and night, it can be a sign that congested links are introducing additional delay.

When measured latency values are plotted per provider over a period of months, certain stable patterns emerge, as shown in the following plots. These patterns can then be interpreted in terms of the physical paths Bahrain's traffic takes as it crosses the globe.

Two factors are important to consider when interpreting these plots. In addition to the latency measurement itself (lower is better), the variance of the latency is an important contribution to a stable perception of connection quality. An Internet path whose latency varies by hundreds of milliseconds over the course of a few days is more likely to offer poor subjective Web performance due to some combination of route instability, packet reordering, and congestion.

2Connect

For most of 2012, 2Connect's latencies from London have been between 90ms and 110ms [A].

In September, wrong-way routes on NTT to Etisalat, going London-New York-San Jose-Tokyo-Singapore at a higher latency mode (225ms) became more common [B].

In mid-October, latencies to 2Connect's UK subsidiary appear to have finally stabilized at 100ms [C].

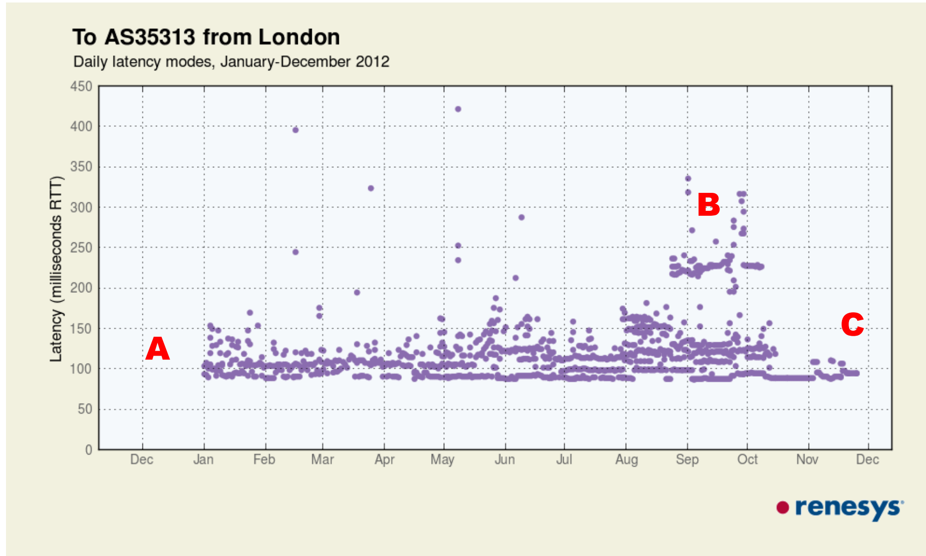

2Connect sees 200-300ms latencies to Dallas [D], and 225-325ms to Palo Alto [E].

Most of the variance is accounted for by different paths across the US before the hop to London.

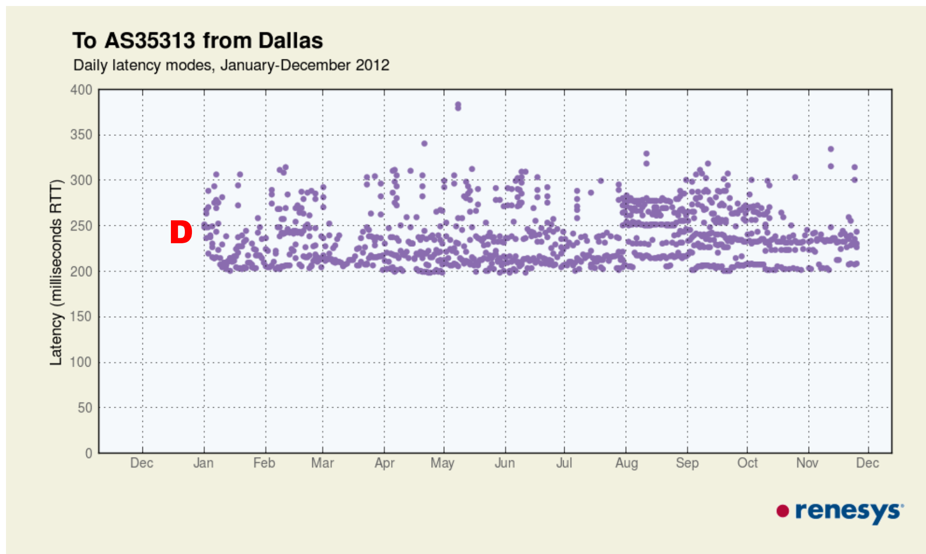

The highest Palo Alto modes [F] are westbound Tata transit through Los Angeles, Tokyo and Singapore. These, too, appear to have become less frequent since mid-October.

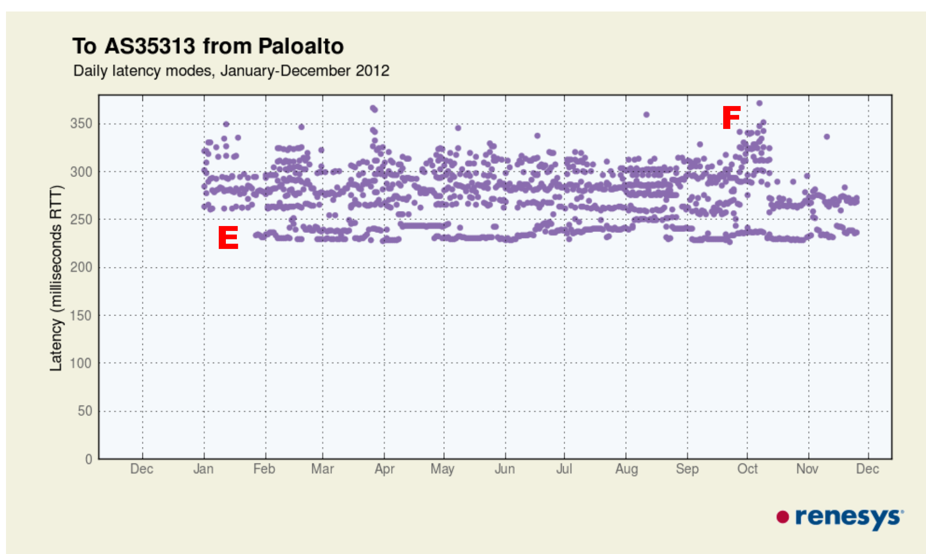

Batelco

Batelco's London connectivity has three primary latency modes.

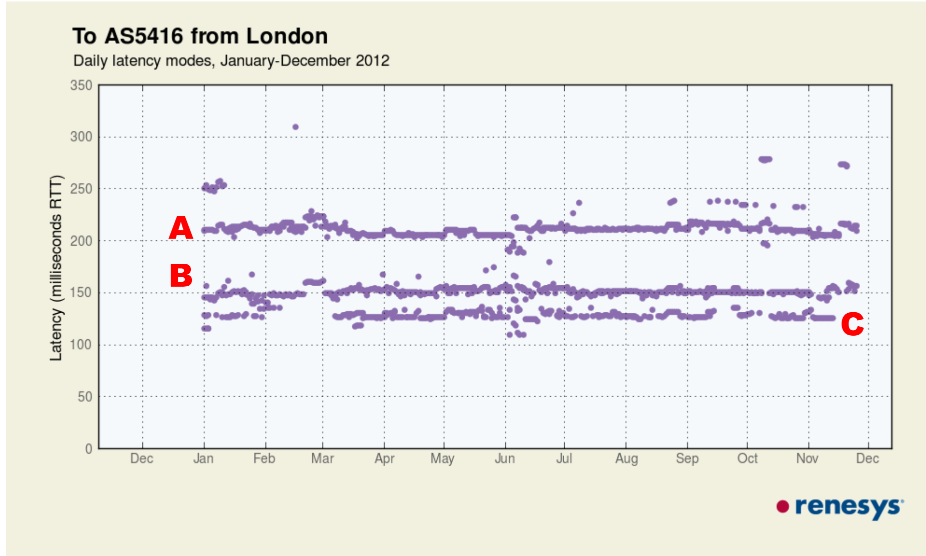

The high mode (A, 212ms) actually appears to include traffic routed by Tata from London to New York, and then directly back to Batelco (i.e., via MPLS) in Manama.

The two lower London modes represent paths routed directly from London, either around the Arabian peninsula on subsea cable (B, 151ms) or directly across it on a terrestrial cable (C, 127ms).

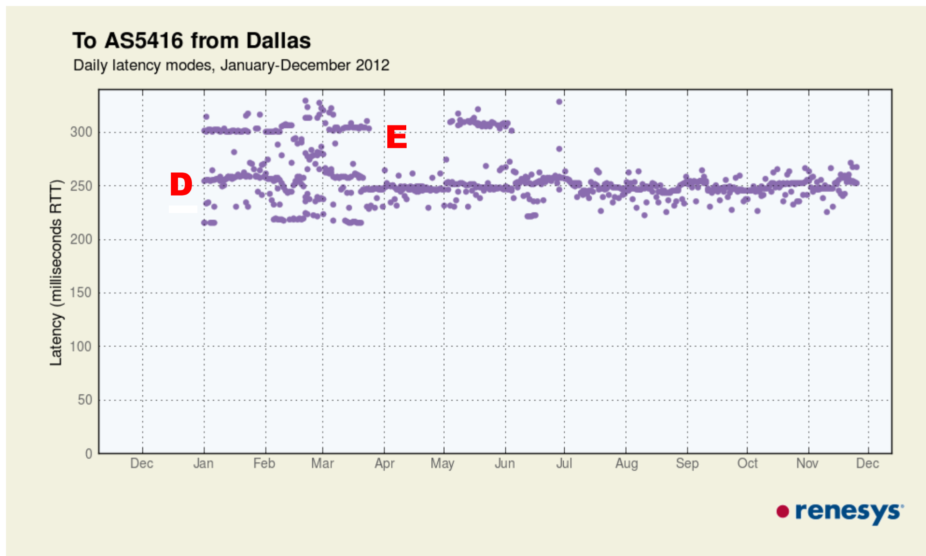

Traffic from Dallas (D, 247ms) passes east through New York to Bahrain; at the start of the year, it could also go west (E, 300ms) through Singapore.

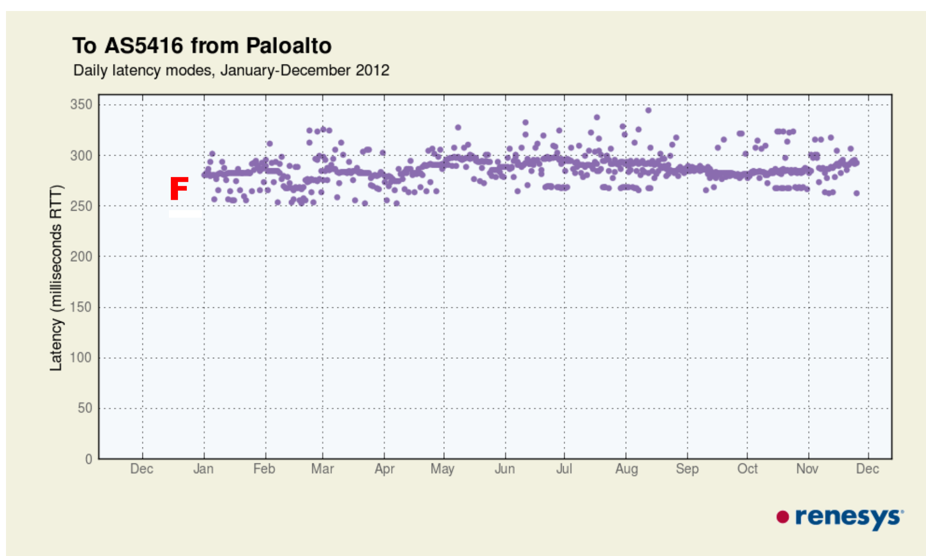

Traffic from Palo Alto (F, 270-290ms) typically passes west through Singapore to Bahrain.

BIX (Bahrain Internet Exchange)

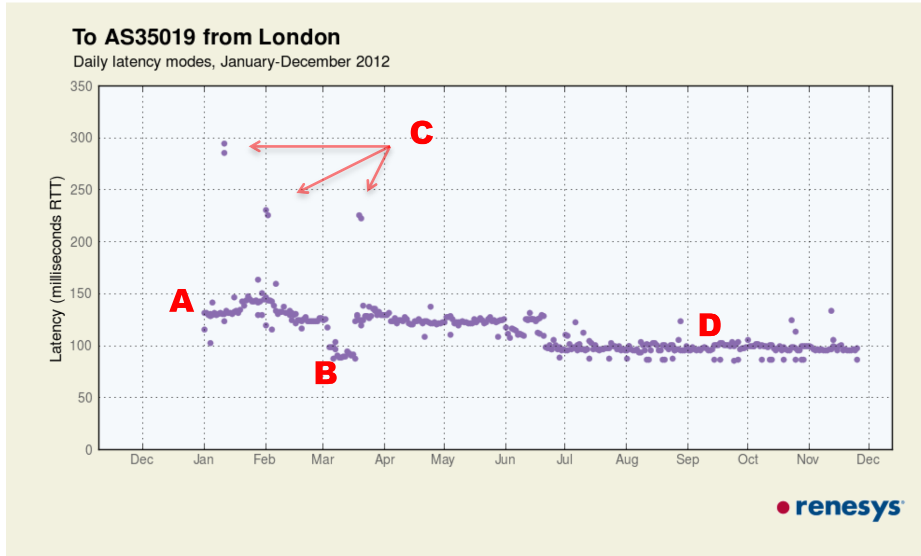

With the FOG cut in March, typical latencies of 124ms to London via Tata [A] improved temporarily to 87ms [B], as Tata traffic shifted west to traverse Saudi Arabia to the landing at Jeddah.

Etisalat provided a rarely used backup route through Singapore [C] at 250-300ms.

In July BIX shifted transit west again, and eventually turned off Etisalat entirely in favor of single-homed reliance on Tata, potentially through the new submarine cable [D] at 100ms.

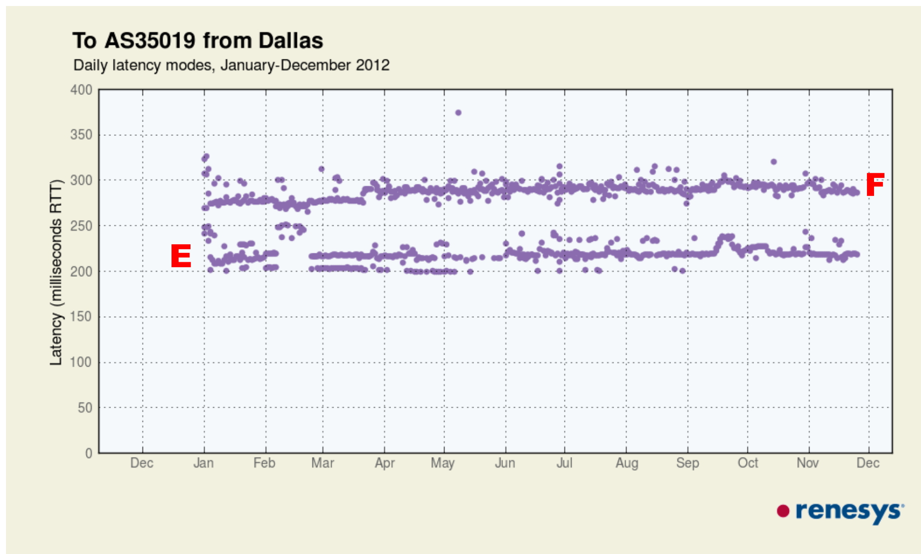

From Dallas, the two major modes are eastbound [E] 200-220ms through/around Saudi Arabia, and westbound [F] at 300ms through Singapore.

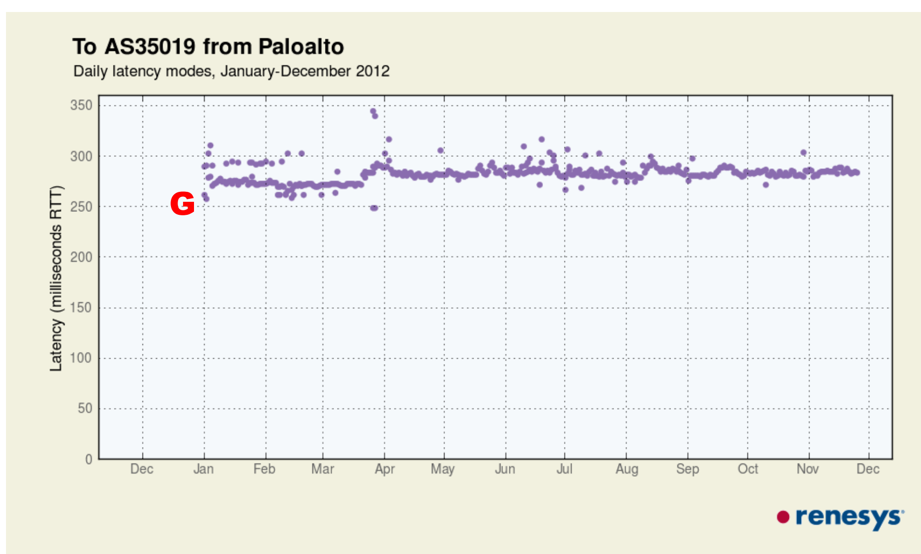

From Palo Alto, the BIX sees only the westbound route [G, 290ms] through Singapore.

Etisalcom

With transit from STC and QTel, Etisalcom experiences two basic modes from London: 103ms [A] and 122ms [B].

Higher modes at 217 and 235ms [C] tend to be longer paths through QTel.

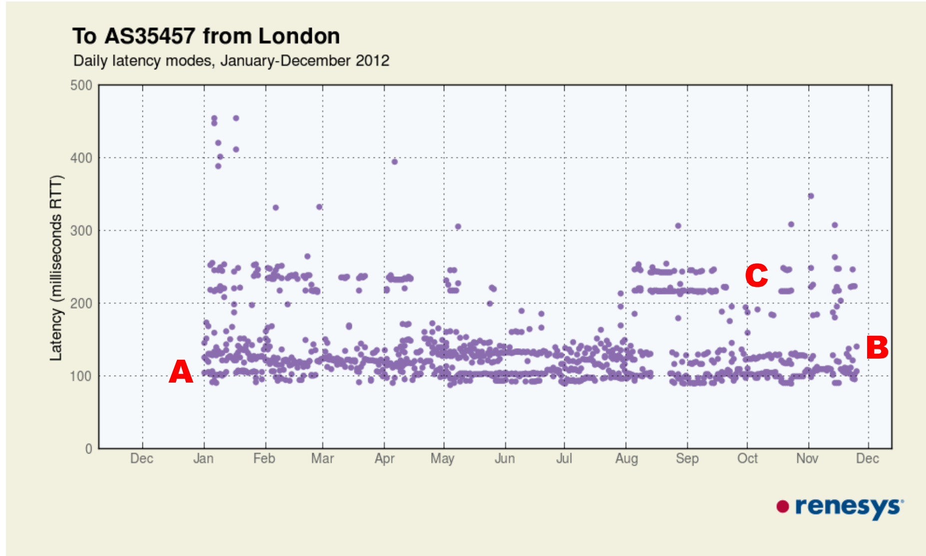

From Dallas, the same two basic modes manifest at 199ms and 221ms, along European routes [D].

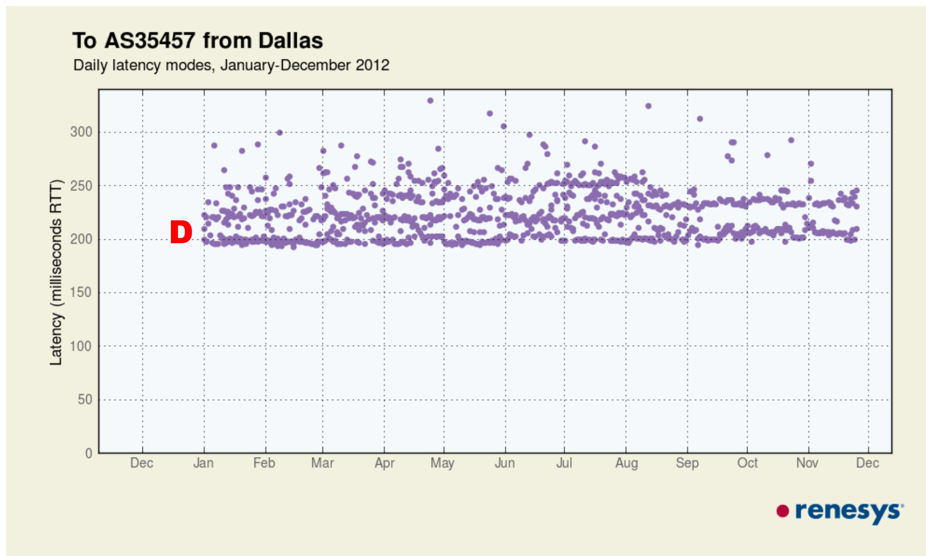

From Palo Alto, paths to Etisalcom head east through New York at 251ms [E] or west at 361ms [F].

The longer western path, which has been seen less frequently since July, uses Verio transit through Hong Kong to reach QTel.

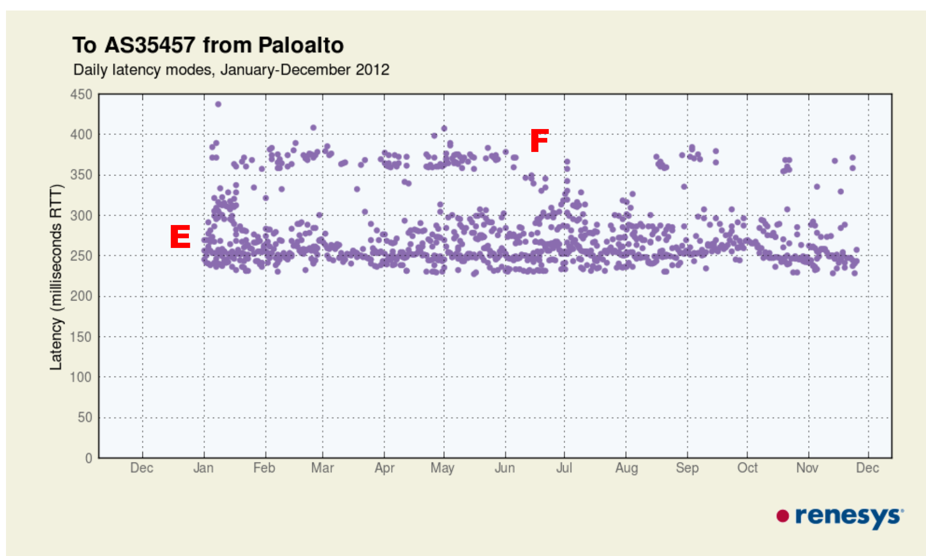

Kalaam

Kalaam receives all its Internet transit from Saudi Telecom. It has a correspondingly simple primary mode of 97ms from London [A].

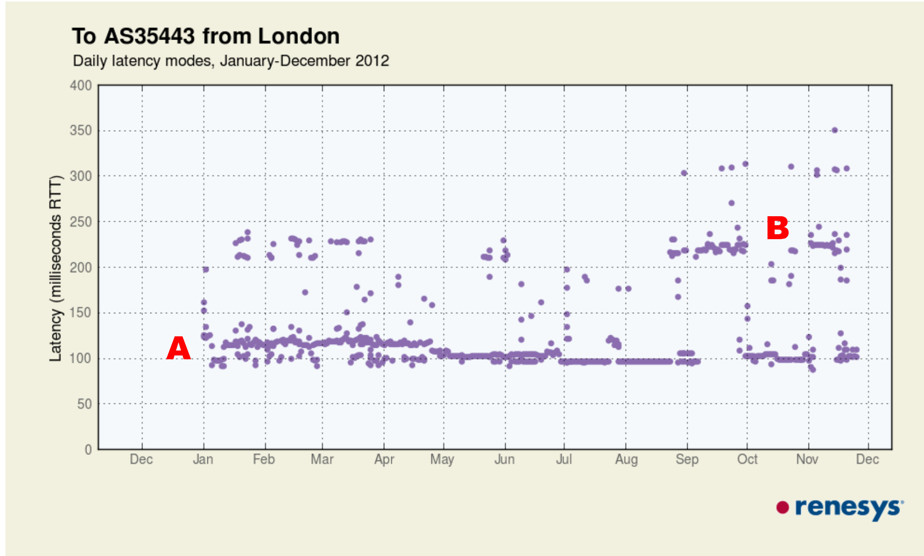

When traces to STC from London run through other carriers, traffic can take circuitous side trips through other European exchange cities such as Amsterdam and Frankfurt. This leads to the higher modes at 116ms and 225ms. [B]

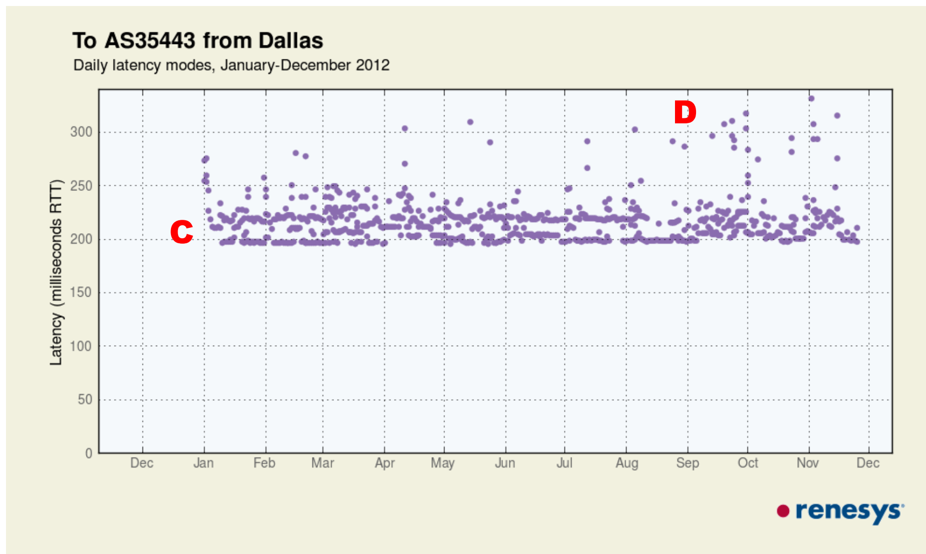

From Dallas, Kalaam experiences delays of 197ms [C] through London to STC. More rarely, Dallas traffic travels the long way to Mumbai, where STC has peering with Bharti Airtel, resulting in latencies of 300ms+ [D].

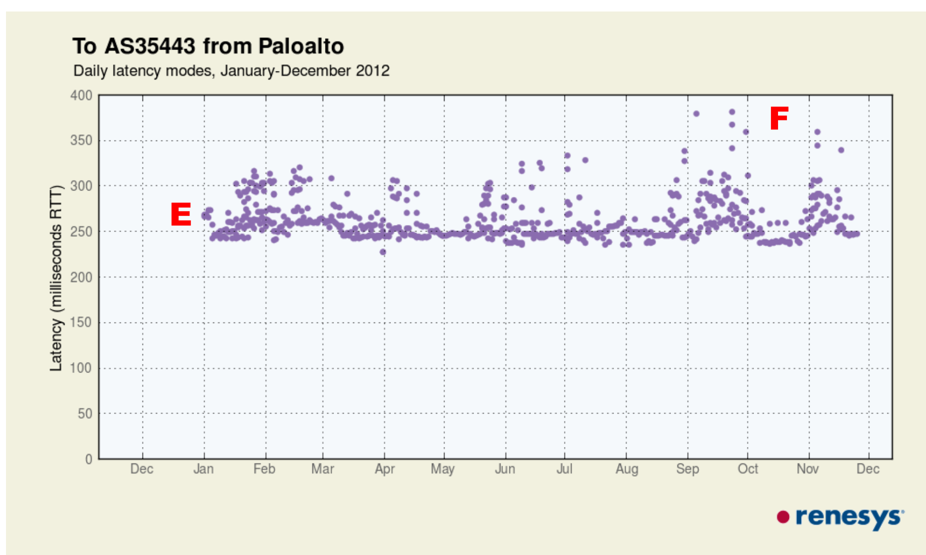

Palo Alto latencies to Kalaam tend to be 248ms through NYC [E], again except when resorting to STC's Mumbai paths, which can exceed 325ms [F].

Lightspeed

Lightspeed switched from FLAG transit to QTel transit in August, with a small amount of BIX transit year-round.

This creates a somewhat complex latency picture from London, with modes at 115ms, 129ms, and 143ms. The highest latencies are 300ms+ paths from London that cross the USA, head to Hong Kong and Tokyo, and return to Bahrain through QTel.

Lightspeed is 229ms from Dallas along UK routes to the BIX through Tata,[C] or 249ms through FLAG [D]. Modes above 300ms include Tata paths through Singapore [E].

Palo Alto latencies are 293ms through London [F]. The very longest paths, on Tata and FLAG through Singapore and Hong Kong, can be more than 400ms! [G]

Since dropping FLAG in August, these longer modes from the US for Lightspeed have largely vanished [H].

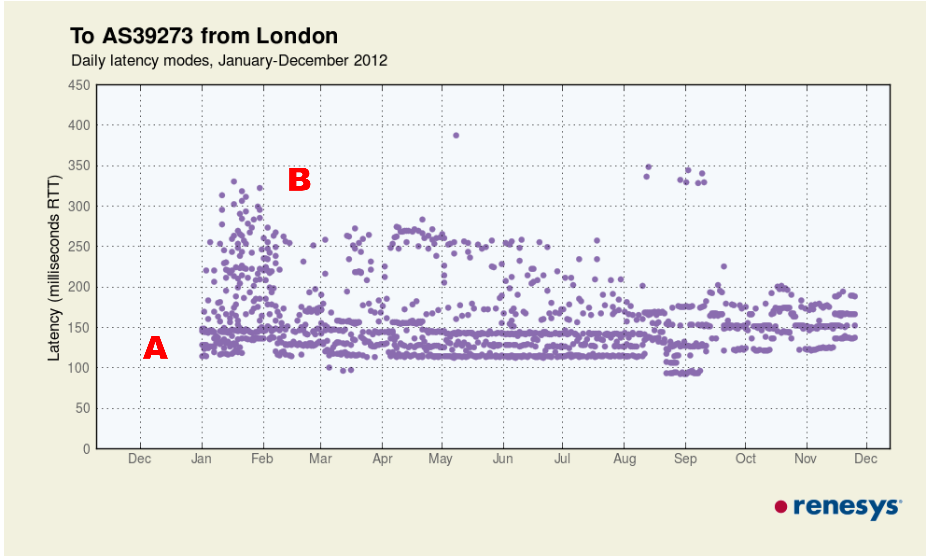

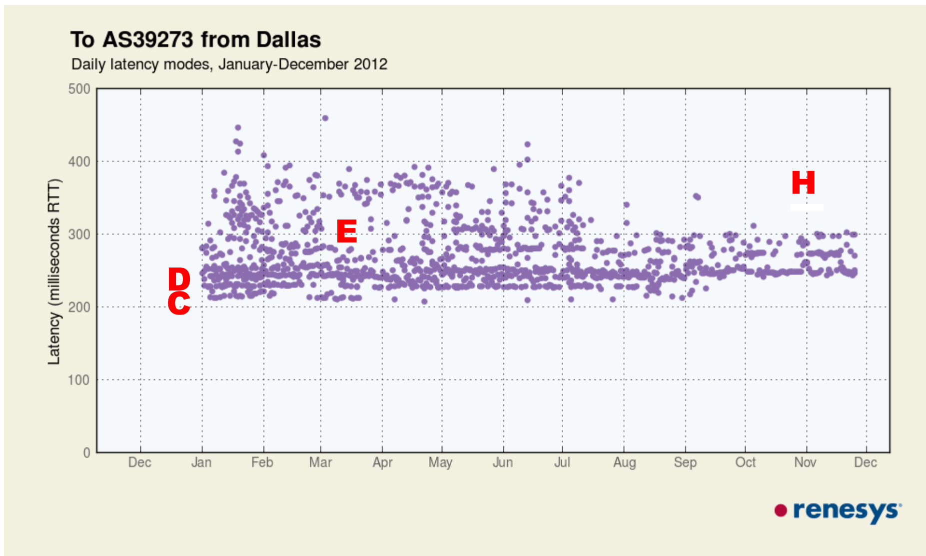

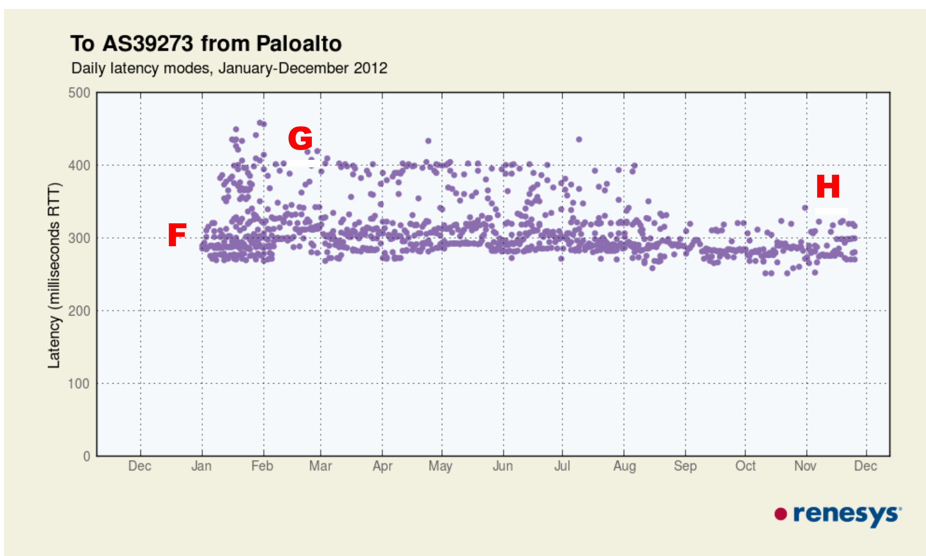

Menatelecom

Menatelecom dropped FLAG transit in May, leaving them single-homed to STC.

From London, this improved their floor latency from 112ms [A] to 94ms [B].

Higher modes are tradeoffs between alternative Saudi infrastructure: AS39386 [C] versus AS51375 [D].

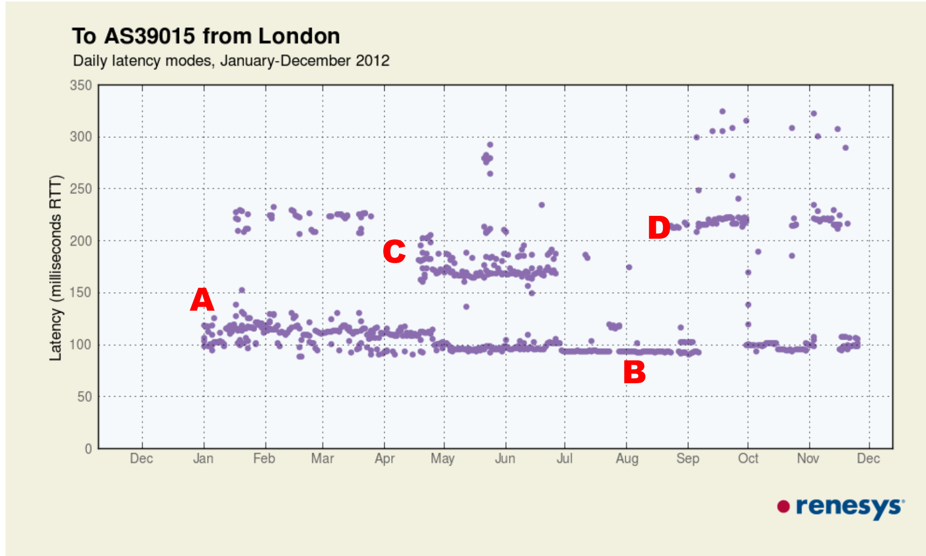

Latencies from Dallas also improved slightly with the loss of FLAG transit in May, from 221ms to 196ms [E]. Paths go via NYC through Europe.

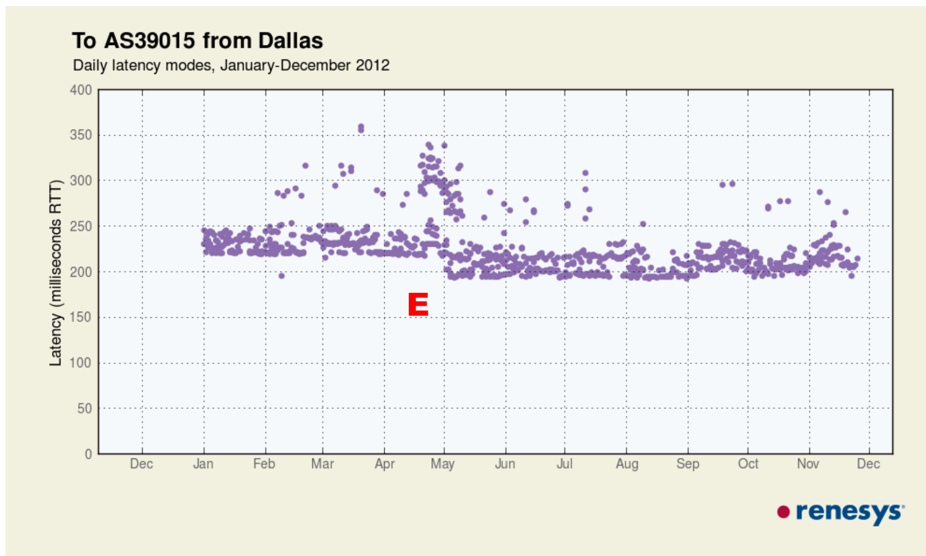

From Palo Alto the effects of the transit shift are subtler [G]; the same substitution of different Saudi infrastructure paths seems to explain the bursts of higher-latency modes in May and July [H]. There is no significant Asian mode.

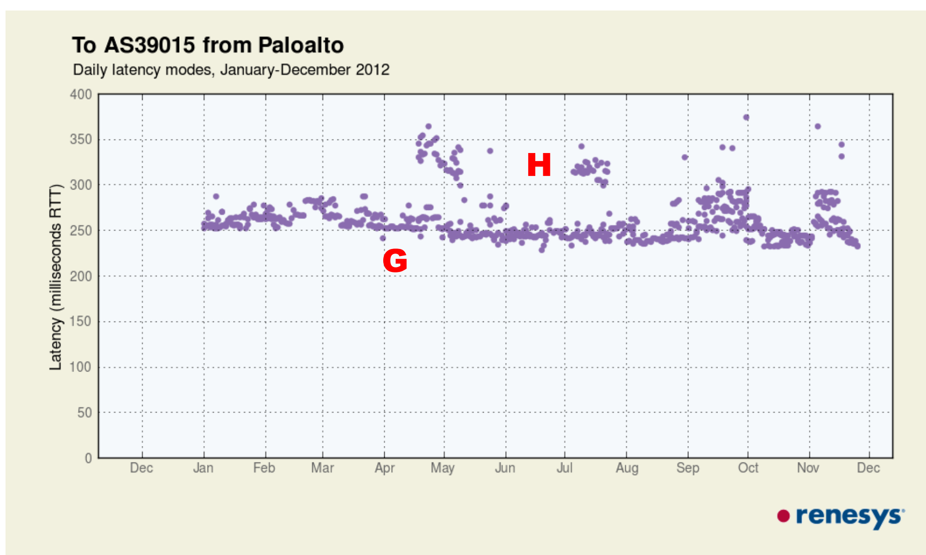

Nuetel

Until the end of August, Nuetel's classic pair of latencies from London was relatively stable, at 97ms and 112ms. [A]

At that point, however, Nuetel dropped BIX transit and added Viva as a provider. Latencies immediately increased and became more variable to all regions.

New, slower modes based on FLAG transit from London appeared at 143ms [B] and 243ms [C].

The 243ms mode [C] appears to be FLAG routing from London the "wrong way," through Mumbai and Hong Kong to Bahrain.

Similar increases appear in Dallas, going from 215ms to 249ms [D]. Some of the 270ms+ modes are Viva bringing transit via STC from Bharti in Mumbai, or FLAG from Hong Kong [E].

Palo Alto, at 283ms+, is all Asian routing through Tata and FLAG [F]; 350ms+ is FLAG through Hong Kong [G].

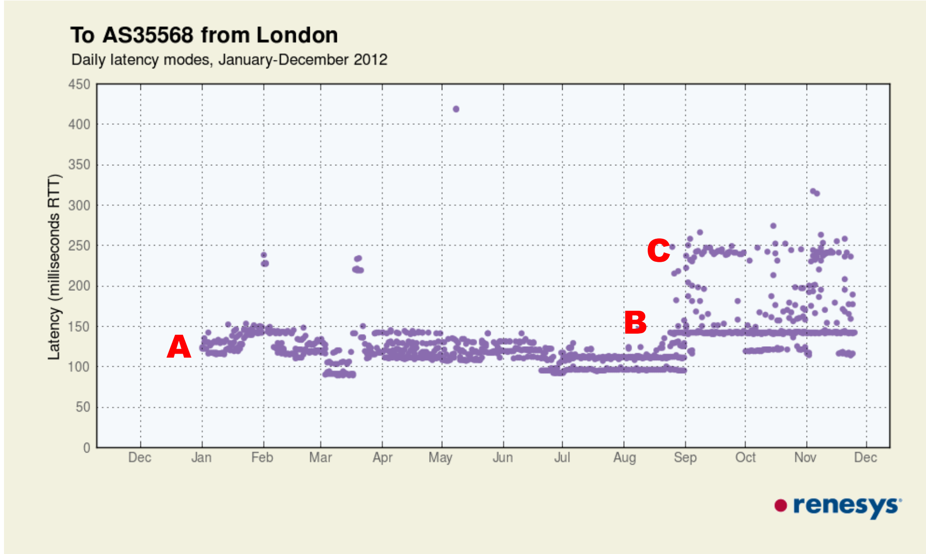

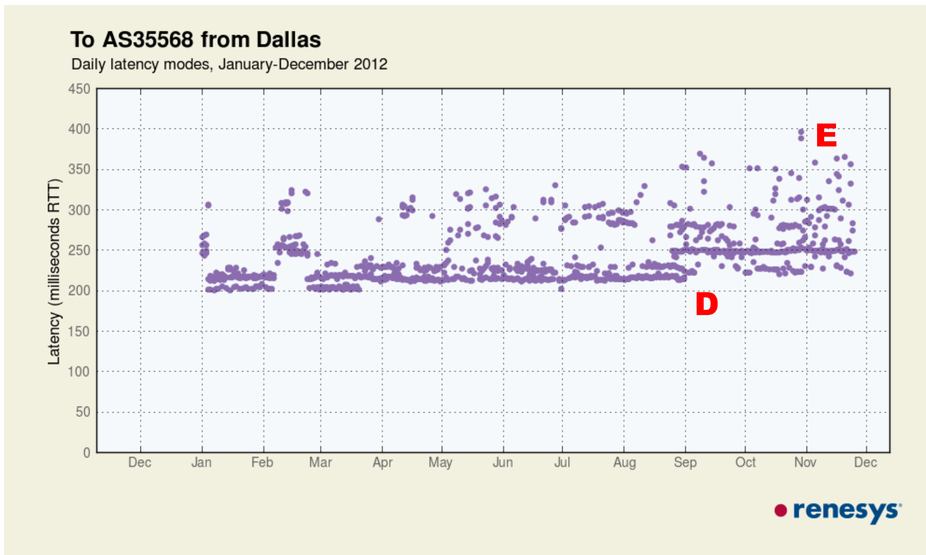

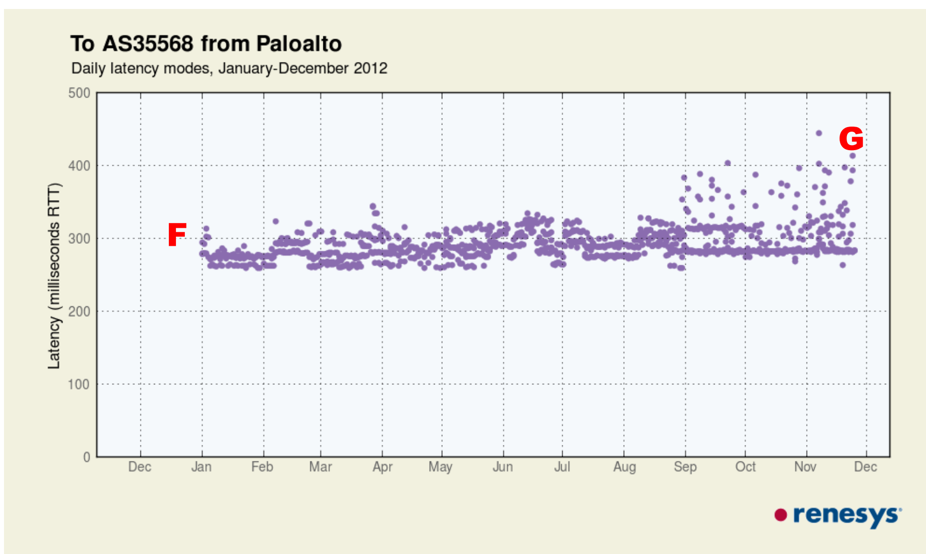

Viva

Latencies from London to Viva have a strong mode at 100ms [A].

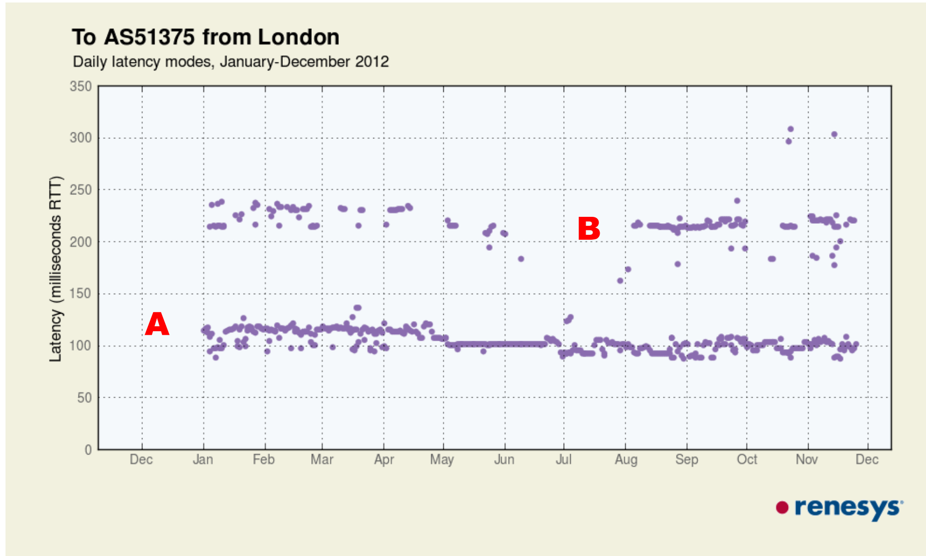

Another, at 215ms, is due to different paths through Saudi Telecom's infrastructure, either as AS39386 or AS51375 [B].

Latencies to Viva from Dallas are 200ms, with some variance along European paths.[C]

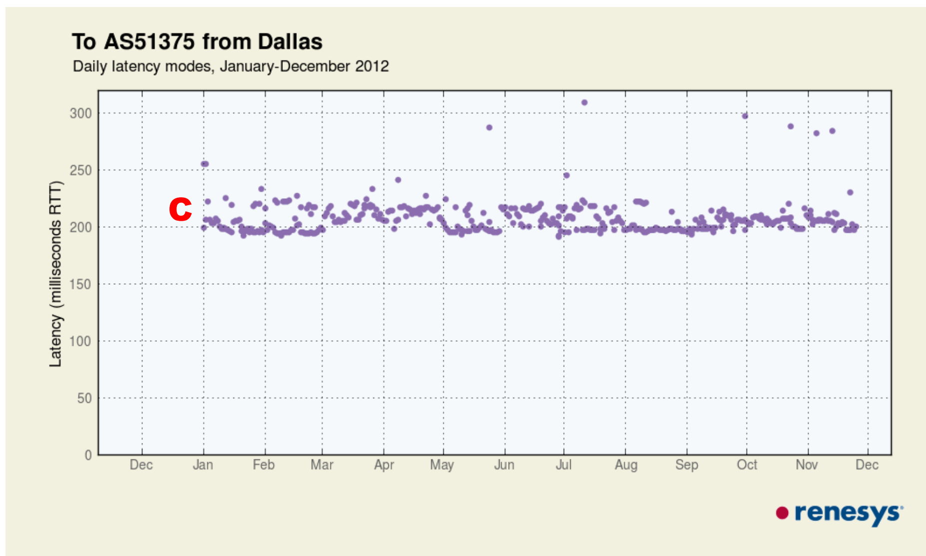

Palo Alto gives 244ms, again with European paths [D]. Above 370ms a weak second mode emerges, based on FLAG transit the other way around the world, through India [E].

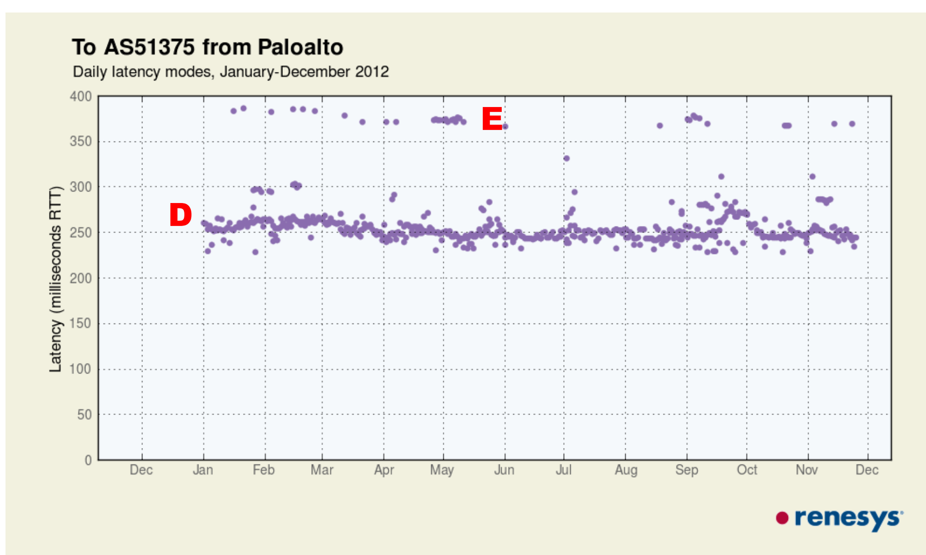

Zain

Zain made two provider switches in 2012, losing Tata in May and gaining QTel in July, while retaining STC. Latencies from London begin at 100ms through STC [A] with a second mode through Tata at 172ms [B]. They shift in August to include a new set of modes through QTel, including one at 215ms. [C].

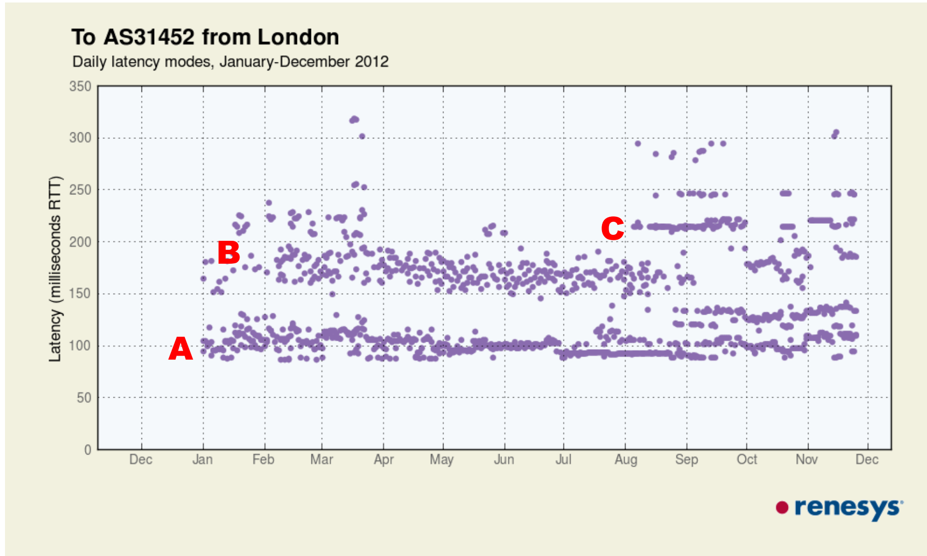

Latencies from Dallas are flat, with a floor at 197ms via London [D] and modes through various European exchanges [E]. No Asian modes are in evidence.

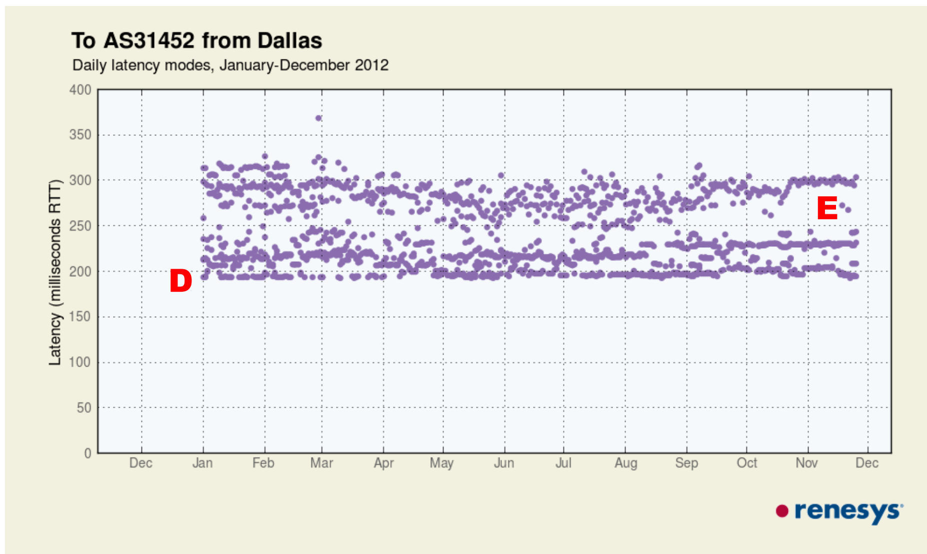

Palo Alto latencies go either way: 251ms via Europe [F], or 327ms via Hong Kong, Singapore, and Tokyo [G].