3561The ISP Column
A monthly column on things Internet

November 2012

Geoff Huston
George Michaelson

Superstorm Sandy and the Global Internet

The Internet has managed to collect its fair share of mythology, and one of the more persistent myths is that from its genesis in a cold war US think tank in the 1960's the Internet was designed with remarkable ability to "route around damage." Whether the story of this cold war think tank is true or not, the adoption of a stateless forwarding architecture, coupled with a dynamic routing system, does allow the network to "self-heal" under certain circumstances. Can we see this self-healing in today's network? How true is this reputation of network robustness in the face of all kinds of adversity? While the Internet is almost everywhere these days, there are still a small number of locations that host a remarkable amount of Internet connectivity. One of these critical points of global connectivity is New York, an area that hosts a significant number of submarine cable landing points as it it the major termination point of the North Atlantic submarine system so it is a major connection point in linking the trunk cable transit systems in Europe, America and Asia (Figure 1).

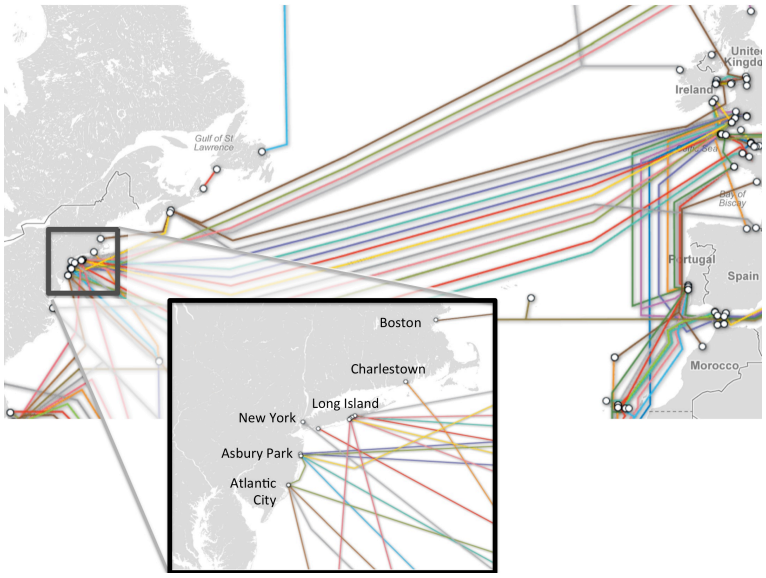

Figure 1: Submarine Cable System US landfall, Telegeography
[http://www.submarinecablemap.com]

On the evening of the 29th October this year a large storm, post-tropical cyclone Sandy, made landfall near New York at 8pm ET, causing an unprecedented tidal surge, flooding subway tunnels in New York City, and causing widespread power blackouts across much of the north eastern seaboard of the United States. The centre of the storm made landfall about 5 miles southwest of Atlantic City, New Jersey, as seen in this NOAA GOES013 infrared satellite image (Figure 2). What this image does not show is that because of the anti-clockwise around the storm, there were heavy southerly winds blowing directly onto Long Island, causing a significant storm surge in Long Island and New York City even though the centre of the storm was further south in New Jersey at this time.

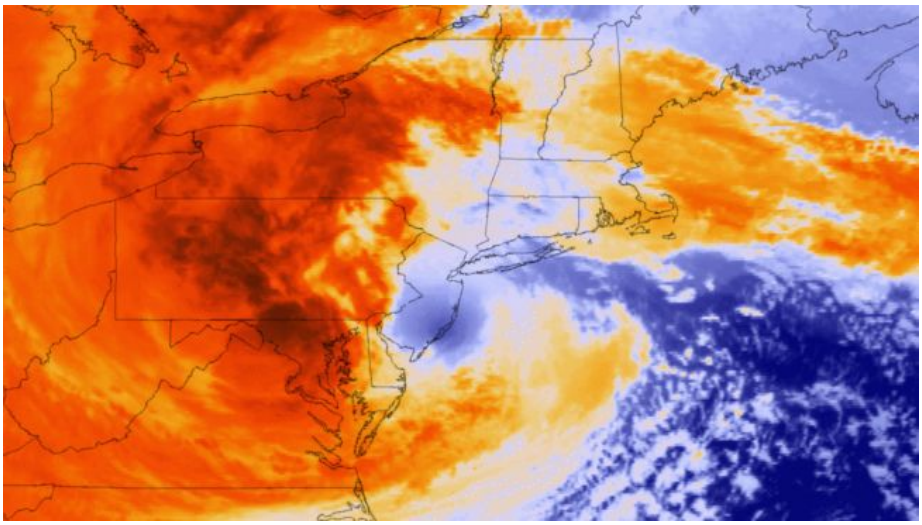

Figure 2. Satellite image of Sandy's Landfall on October 29
[http://www.weather.com/news/weather-hurricanes/superstorm-sandy-by-the-numbers-20121030?pageno=5]

If you compare these two figures, looking at the storm's landfall and the map of submarine cables and their landing points, then it's evident that Sandy's impact was directly aligned with the landfall points of some 25 submarine cable systems on the New Jersey and New York shores.

How well did the Internet fare in the face of this rather severe climatic event? Did this storm generate outages that were visible across the entire Internet?

It's certainly the case that much of the storm's impact was local to the north east of the United States, and New York in particular. Renesys has produced an animation of the localized outage consequences for this region, using fine-grained geo-location information to pinpoint the network outages (Figure 3).

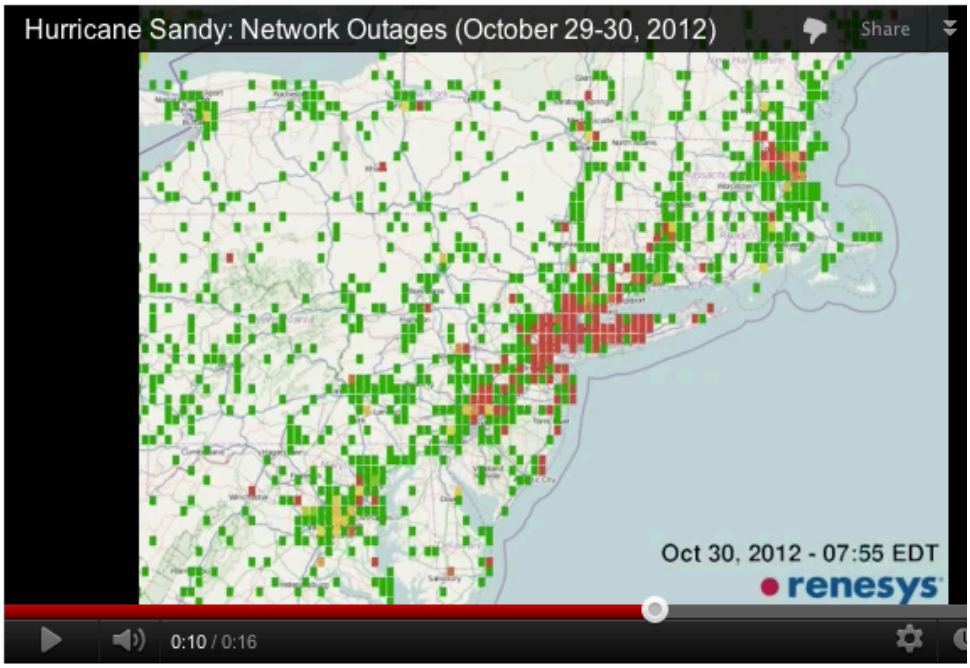

Figure 3: Renesys animation of localized outages from Superstorm Sandy
[http://www.renesys.com/blog/2012/10/hurricane-sandy-outage-animati.shtml]

It is clear from this Renesys study that a significant number of Internet services in metropolitan New York, across the state of New York, and in neighbouring states were affected, especially between the

critical period of landfall and the following twelve hours, as the storm surge took out electricity distribution and supply, and cascading failures of systems came into effect.

But to what extent were these localized events, where networks located in New York around the point of landfall were affected? And to what extent did we observe non-localized effects, where other networks that use facilities in New York for transit services were affected? Can we assess the impact of Sandy's landfall in New York on the global internet?

One way to answer this question is to examine the Internet's routing table during this time. Figure 4 shows a second-by-second plot of the size of the BGP forwarding from the start of the 29th October (0000UTC) to the end of the 3rd November (2359 UTC) based on scanning the logged BGP updates from a router located in Japan, in AS 4777 and a second router located in Australia, in AS 131072.

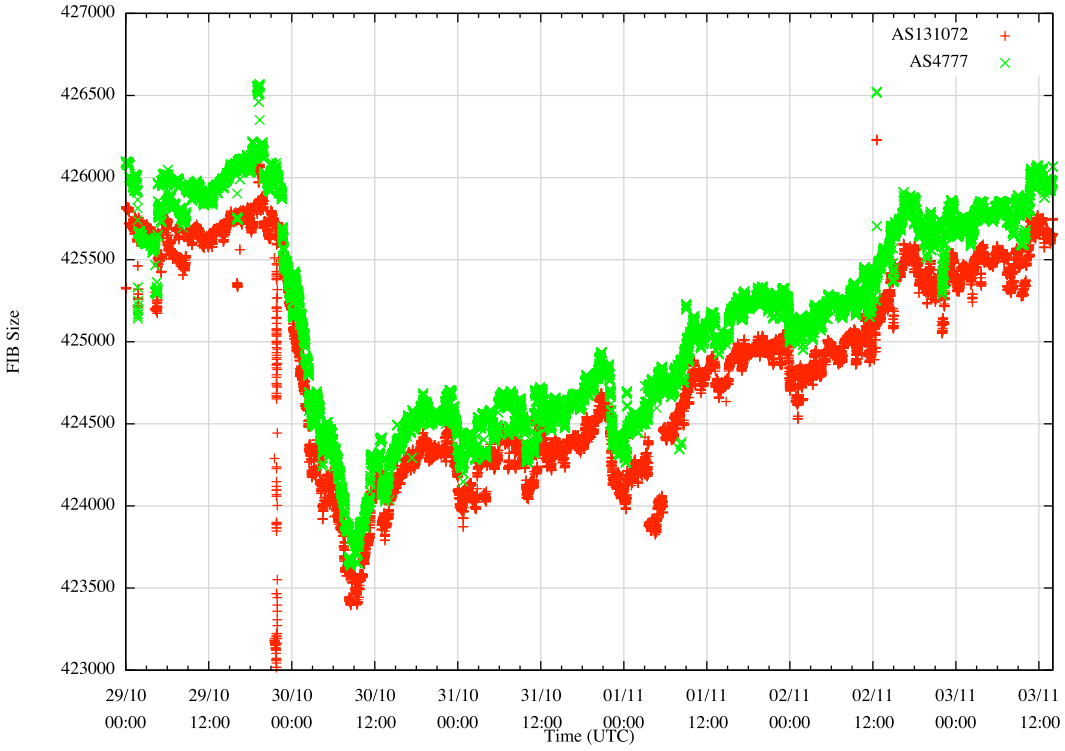

Figure 4: BGP FIB Size for AS4777 and AS131072

It appears that the time of interest with respect to Sandy's disruption on the Internet is the period from Mon Oct 29 1900 UTC (3pm US East Coast time) to Tuesday Oct 30 0900 UTC (5am US East Coast Time), where a net loss of 2,500 routes was observed over this 14 hour period. What this shows is a sharp drop in the number of routes, followed by a rapid restoration of 1,000 of these routes, and a somewhat longer timeframe to restore the remaining routes. There is also the anomaly recorded just before midnight UTC in AS131072, which is the result of a local peer session reset in BGP at the route collector, and is not related to Sandy at all!

As this is a second-by second view based on the recorded BGP updates, its possible to bring into sharper focus the time period around Sandy's landfall at midnight UTC (Figure 5). Interestingly, apart from the route reset, there is no immediately obvious single timeframe where many hundreds of networks were withdrawn in a single event, which tends to suggest that there was no large scale loss of a set of networks that exclusively use the facilities in New York for transit.

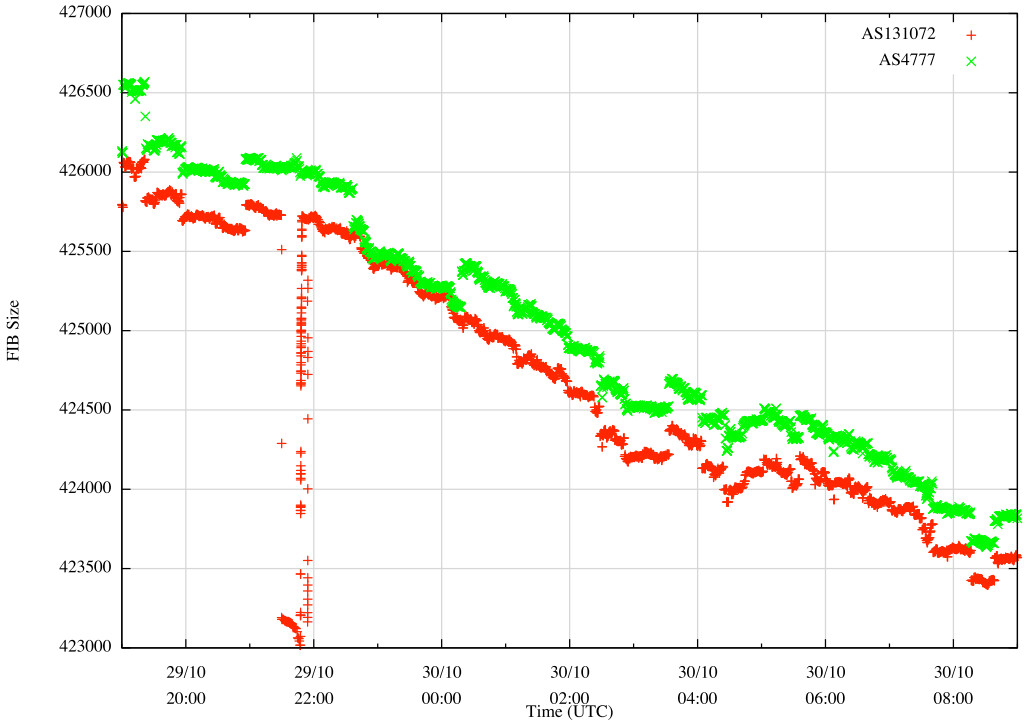

Figure 5: BGP FIB size as seen at AS4777, using BGP Updates

In this 14 hour period we observed some 159,385 "prefix updates".

> Here a "prefix update" is defined as a withdrawal of a prefix, an announcement of a prefix, or an update to an existing announcement of a prefix. As BGP will endeavour to pack multiple withdrawals into a single BGP protocol message, and pack multiple announcements into a single protocol message if they share a common set of attributes, this count of "prefix updates" is not the same as a count of the number of BGP update messages. This count is the number of prefixes that are updated by these BGP protocol interactions.

The following figure (Figure 6) shows the amount of prefix update activity superimposed upon the log of the FIB size using the view of AS 4777. Notable in this figure is the peak of activity just prior to midnight UTC on the evening of the 29th October (8pm ET), which saw a peak of some 4,000 prefix updates received in a single BGP update interval.

Of these 159,385 prefix updates, how many of these routing updates were associated with some form of extended reachability impairment? What we are trying to do is to filter out the churn of updates associated with a dynamic change in the routing state of a prefix due to backup routes being invoked, and try to identify those prefixes which did not fail over to a backup route, but failed completely. Also we would like to filter out transient route failure as the protocol searches for a new converged routing state, and identify those prefixes that were seen to fail for an extended period. Here we will use a filter that identifies those prefixes that became unreachable during the night of the landfall of Sandy on the evening and night of the 29th October any time on or after 3pm US East Coast time, and remained unreachable until at least 5am US East Coast time on the morning of the 30th October.

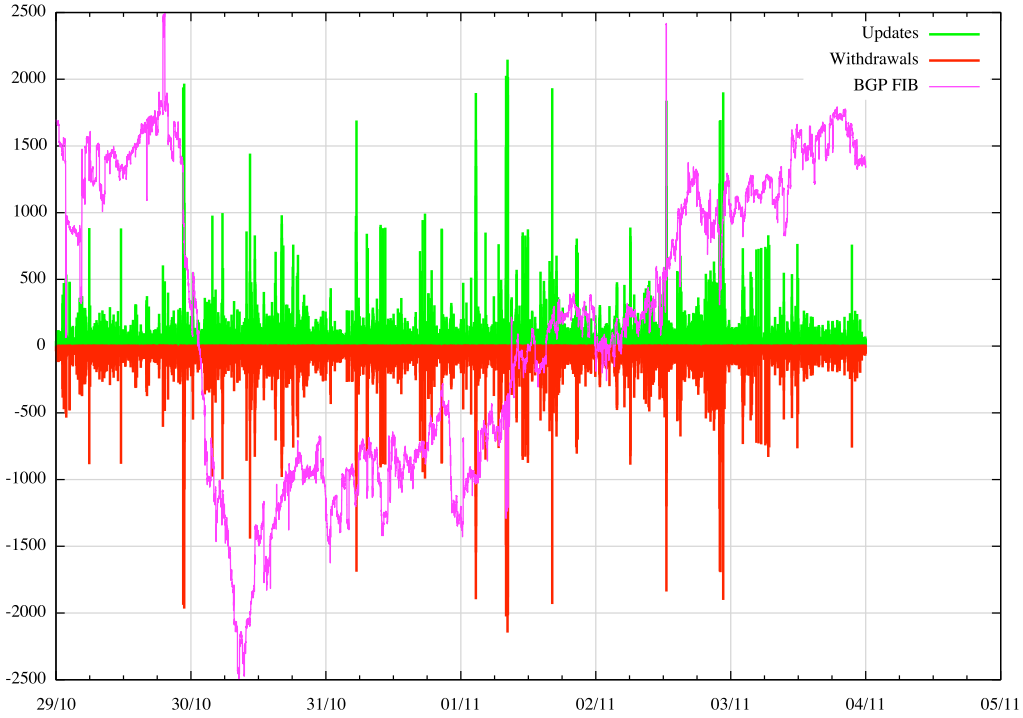

Figure 6: BGP FIB size and Prefix Update and Withdrawal rates, as seen at AS4777, using BGP Updates

In this 14 hour period the BGP observation point at AS 4777 saw some 3,721 routes withdrawn from the global routing table that were not restored by the end of this period. As the Renesys study has shown, we are aware that a large number of these prefixes were local to the north east of the United States. What we are interested in is to see if the effects of the storm included effects on the transit arrangements that connect the continents together, as this area of the US contains the landfalls of most of the trans-Atlantic submarine cable systems. Any outage on these cable systems would be expected to cause some level of impact on the overall picture of global reachability within the Internet.

One way to look at this is to use geo-location to map each withdrawn prefix to the country where the prefix has been assigned or allocated. The withdrawn prefixes were from 77 countries. The top 10 withdrawals on a country-by-country basis is shown in the table below.

Withdrawn Routes	CC	Country Name
2291	US	United States of America
476	AU	Australia
224	RU	Russian Federation
172	BR	Brazil
45	GB	United Kingdom of Great Britain and Northern Ireland
39	IR	Iran (Islamic Republic of)
33	LB	Lebanon
31	GT	Guatemala
31	BF	Burkina Faso
28	MX	Mexico

Table 1: Country of Origin of Withdrawn Routes seen on the evening of the 29th October, as observed by AS4777

This table indicates that if there was a problem that encompassed more than the locality of immediate storm damage then it was only evident in Australia, Russia and possibly Brazil, After that the numbers of route withdrawals are sufficiently small that it is not clear that this was the side effect of Sandy or part of the normal background activity of the Internet's routing system.

In the case of the withdrawals of Australian routes, 421 of these withdrawals occurred at 19:22 on the 29th of October, UTC. As seen in Japan, these prefixes all had a path of:

AS4777 (APNIC, Japan)
AS2516 (KDDI, Japan)
AS3257 (TINET, transit)
AS7473 (Singtel, Singapore)
AS4804 (Optus, Australia)

So, could this outage of 421 routes be attributed to Sandy or not? An Australian view of the same set of outages saw the withdrawal of a path of:

AS131072 (APNIC R&D, Australia)
AS4608 (APNIC, Australia)
AS1221 (Telstra, Australia)
AS4637 (Telstra Global, transit)
AS3561 (Savvis, US)
AS3257 (TINET, transit)
AS7473 (Singtel, Singapore)
AS4804 (Optus, Australia)

It is possible that the one common factor here, the transit through AS3257, TINET, was impacted by Sandy in this period, but the connection is not immediately obvious. Indeed during the withdrawal of these AS4804 routes we saw from AS 4777 the following path appear for just 30 seconds, and then the entire route was withdrawn:

AS4777 (APNIC, Japan)
AS2497 (IIJ,Japan)
AS7473 (Singtel, Singapore)
AS4804 (Optus, Australia)

If this was just an outage in TINET's networks then this backup path would've probably stayed up, and not be withdrawn, so it does appear that what we see for these AS 4804 routes the outage we've noted is the result of some form of outage in or close to AS4804 in Australia.

I'm not sure that I could make the case that the withdrawal of these Australian routes were readily attributed to Sandy. Given that many routes are exchanged directly between AS1221 and AS7474, and AS7474 is the upstream AS for AS4608, then perhaps the observation that instead of the path <1221 7474 4804>, which is contained with Australian network infrastructure, prior to the evening of the 29th of October we were seeing a path of the form <1221 4637 3561 3257 7473 4804> was itself a signal of a routing problem with AS4804, and perhaps Sandy was not the root cause here.

What about the other large set of routes, those 224 withdrawn routes from Russia? This presents differently, in so far as there is no single event that withdraws the majority of these routes. There is a steady trickle of single withdrawals over the 14 hour period, with just one event withdrawing 18 routes (Figure 7). Many of these routes include AS 12389 (ROSTELECOM), and AS12389 has connections to a number of providers for transit towards Asia/Pacific, including AS1273 (Cable and Wireless), AS2914 (NTT America), AS25165 (KDDI), AS3257 (TINET) and AS6762 (SEABONE, Telecom Italia). Over this 14 hour period we saw some 20,453 updates relating to these Russian routes, but the level of service disruption in terms of extended period of unreachability was much lower.

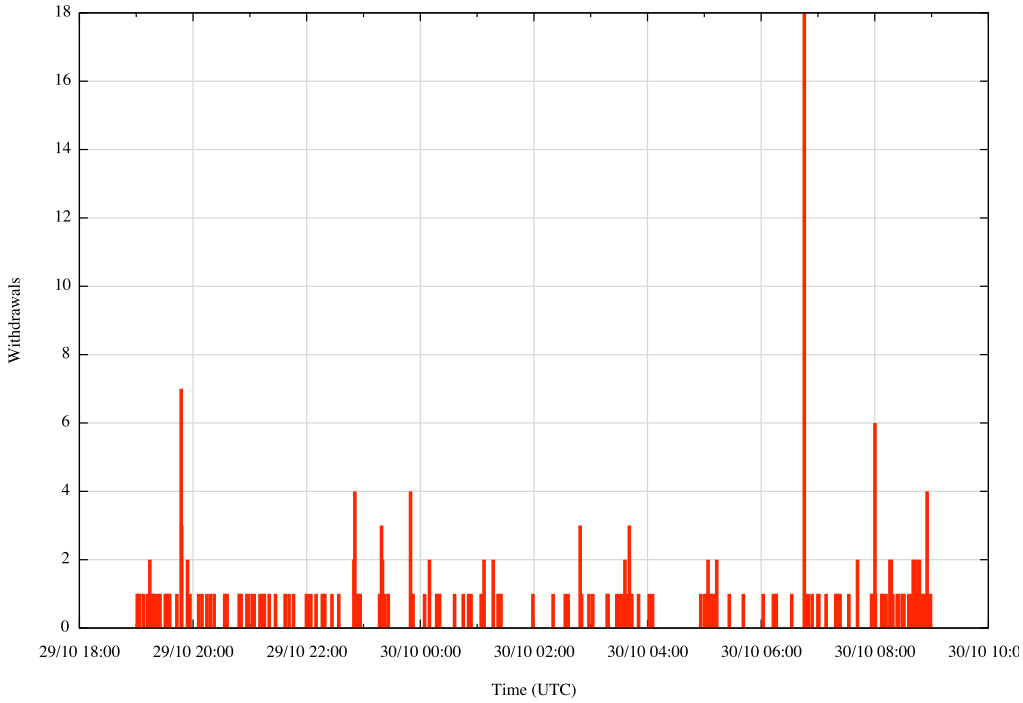

Figure 7: Time Distribution of withdrawal of RU origin routes

A similar story can be see with Brazilian routes, with 17,977 updates in this period with 172 extended withdrawals.

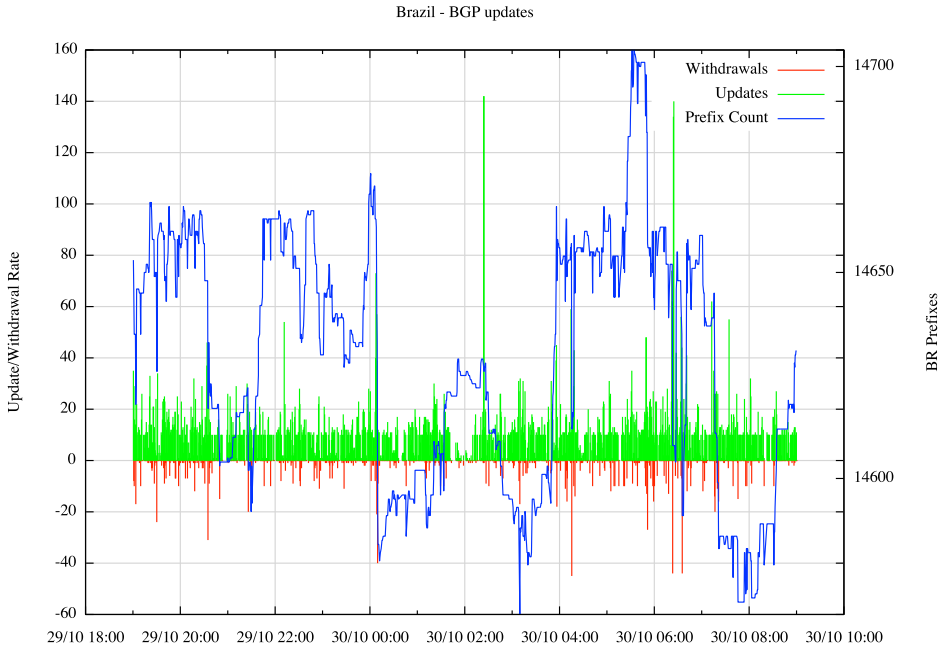

Figure 8: Time Distribution of BGP activity of BR origin routes

It is evident that a small scale disruption of some 80 routes occurred from 8pm ET that were restored by midnight, and a smaller scale disruption of some 40 routes 3 hours later at 3am ET that were restored one hour later, but, as with the Russian origin routes there was no visible extended outage of a large set of Brazilian routes.

Again we can look at the time distribution of extended outages of Brazilian routes, were the withdrawn prefix was not restored during the 14 hour period of the storm's landfall. Once more there is no sharp peak of withdrawals observed for this collection of routes (Figure 9)

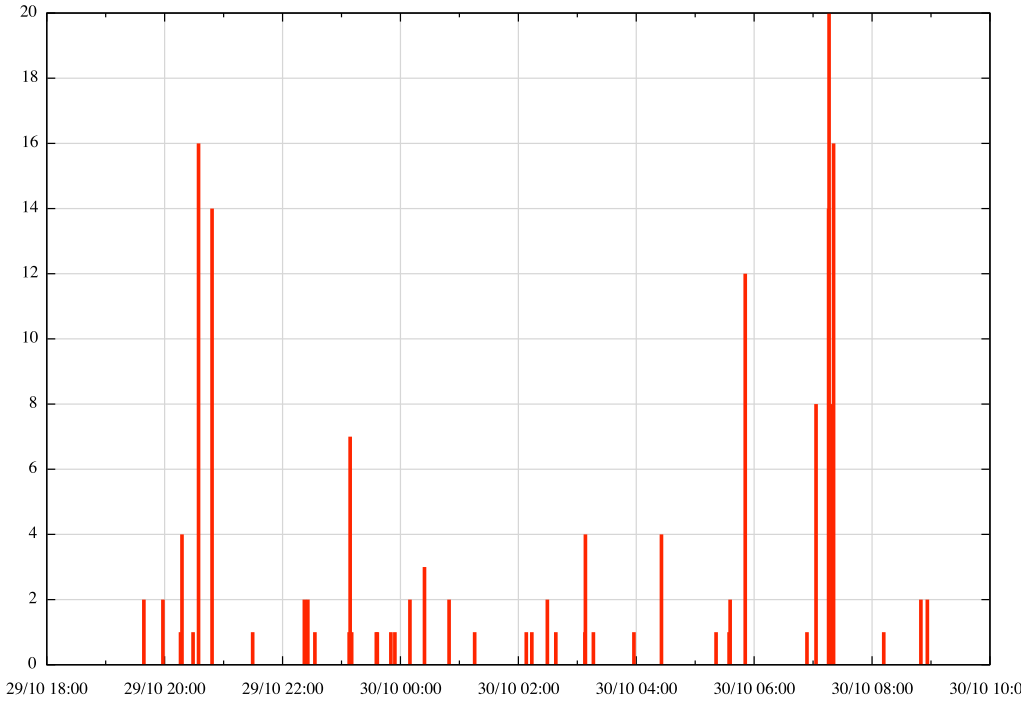

Figure 9: Time Distribution of withdrawal of BR origin routes

What can we conclude here from this routing data? While the effects on Sandy on the network infrastructure in the United States, and the North East in particular, were severe, the effects on the global Internet were far more subdued. The large collection of trans-Atlantic cable systems that surface in this area of the US were largely unscathed by this storm, as were the related trunk switching and transmission plant. We did see some disturbances in a number of routes, but the global Internet managed to live up to its name and route around much of the local damage, and the impact on global connectivity looks to have been relatively minor and short lived. As long as the underlying network connectivity mesh is sufficiently dense the Internet can indeed "route around damage"

Disclaimer

The views expressed are the author's and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.