BGP Behavior through Analysis of Prefix Beacons

Matthew Guidry and Yashas Shankar

ABSTRACT

We are analyzing BGP beacons which are announced and withdrawn, usually within two hour intervals. These announcements and more particularly the withdraws have an effect on neighboring prefixes and down the line to their neighbors as well. We are conducting analysis on the number of updates that are propagated as a result of these events. We then look at the number of updates which could have been eliminated if route flap damping had been installed on the routers. We also analyze the relative convergence period associated with each beacon event and how it is correlated to the number of update messages collected by RouteViews[4].

I. WHAT IS A BGP BEACON

The Border Gateway Protocol (BGP) is the central agent used to conduct data through the internet. A large group of networks put together have a gateway router which communicates to the internet at large. All data packets which are not meant for the immediate network go through this gateway router into the global internet through this router. This router has a prefix which is associated with it and this is how other networks identify it.

Once a packet leaves your local prefix and enters the BGP it can take a number of different paths specified at each router to reach its destination. A number of previous analytical and measurement studies have shown the existence of BGP path exploration and a slow convergence in the operational Internet routing system, which can potentially lead to sever performance problems in data delivery. [1] This path exploration suggests that some BGP routers, in response to a path failure, may try a number of transient paths before settling and declaring a new best path or declaring a destination unreachable. This may cause the entire network to take a long time to settle and converge to the final decision, which causes slow routing convergence. An example of a failed path and resulting path exploration is depicted in figure 1.

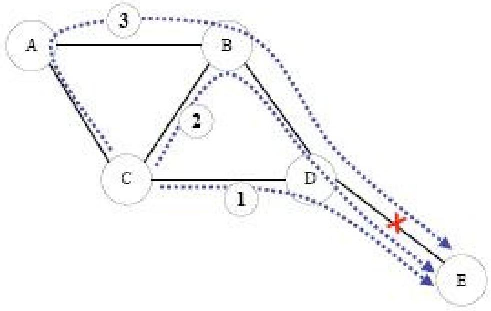

Figure 1: Failed Path and Resulting Path Exploration

To cause these events to occur in predictable and measurable time intervals

BGP beacons are used. These beacons are an active approach to be announced and withdrawn predictably, usually every two hours, to provide input to the routing system which is known. Without this predetermined active manipulation of the BGP, inferences would be very difficult or all together impossible [2]. The beacon announcements do not cause as many updates to be propagated as the withdraws, which are of more interest.

II. BGP BEACONS

There are two groups of BGP beacons, which differ slightly in implementation. The first group is called PSG beacons because they are hosted at psg.com. With PSG beacons, two attributes have been "hijacked". The aggregator IP attribute, which is an IP address, is set to have the form 10.X.Y.Z where 0.X.Y.Z (in binary) represents the number of seconds since the start of the month (GMT). The aggregator ASN attribute is incremented with each announcement and cycles through the values from 64,512 and 65,635. [2] [5]

The second group is called the RIPE beacons, they have been set up as part of the RIPE Routing Information Services (RIS). Each route monitor is associated with a different BGP prefix ranging from 195.80.232.0/24 to 195.80.244.0/24. [2] [6]

III. Monitoring Points

There are a number of monitoring points that are publically available and many which are maintained by businesses and corporations for private security use. RouteViews[4] is a popular example of one of the free monitoring stations. It peers with about 30 different networks and receives all of the BGP updates. Since the Beacon prefixes are not aggregated and should be globally visible, they can be seen by the feed available at all of the monitoring points.

IV. DATA SET AND PROCESSING

In this paper we present measurement results which were collected from RouteViews [4]. We are presenting data from the RIB files for 2004-02-06 to 2004-02-11. We then opened up the data file by unzipping it and using the provided route_btoa function to transform the files into a machine readable format and used perl to insert the data from those large files into a mySQL database. We categorized each of the announcement and withdraw periods into different events, with the withdrawal periods being the one of particular interest. See figure 2 for an example of the data contained within each update.

To do our analysis we processed the beacon updates following the methods described in [2]. We used the PSG beacons which are set up at psg.com. These beacons are announced with timestamps and sequence numbers, while RIPE beacons are not. In particular we are interested in an attribute associated with a PSG prefixes referred to as the *anchor prefix*. The anchor prefix is used for processing out unwanted announcements which are not associated with our events.

First we categorized each event to correspond with the time period that a prefix

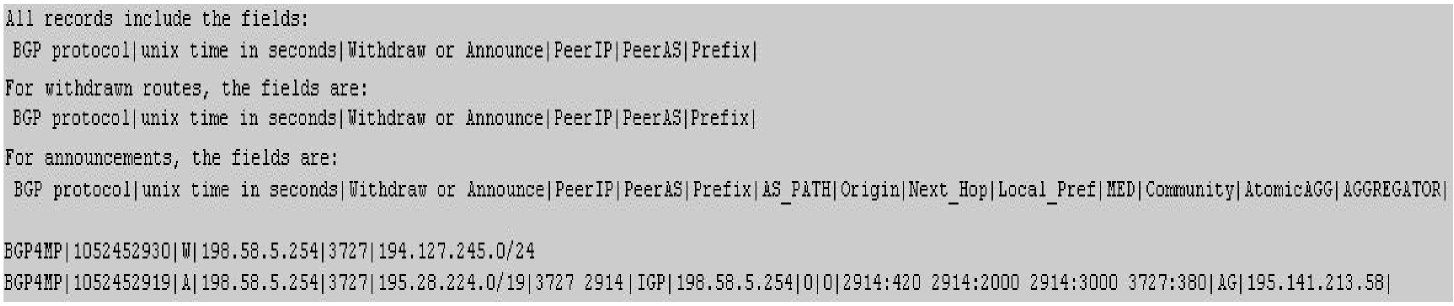

```
All records include the fields:
 BGP protocol|unix time in seconds|Withdraw or Announce|PeerIP|PeerAS|Prefix|
For withdrawn routes, the fields are:
 BGP protocol|unix time in seconds|Withdraw or Announce|PeerIP|PeerAS|Prefix|
For announcements, the fields are:
 BGP protocol|unix time in seconds|Withdraw or Announce|PeerIP|PeerAS|Prefix|AS_PATH|Origin|Next_Hop|Local_Pref|MED|Community|AtomicAGG|AGGREGATOR|

BGP4MP|1052452930|W|198.58.5.254|3727|194.127.245.0/24
BGP4MP|1052452919|A|198.58.5.254|3727|195.28.224.0/19|3727 2914|IGP|198.58.5.254|0|0|2914:420 2914:2000 2914:3000 3727:380|AG|195.141.213.58|
```

Figure 2: Example of data format found on RouteViews

was being announced and withdrawn. There are six withdrawals per day, on the PSG schedule we are analyzing the ones that announce first at 3am and have their first withdrawal at 1am [5].

For these two hour intervals we typically see a total number of updates propagated to RouteViews of around 5 million. We then begin to dissect these updates into groups associated with our beacon prefix being withdrawn and those that are being sent because of other effects in the BGP. We look for the source AS associated with the beacon prefix. However, this subset may still contain announcements due to a small set of updates not related to our beacon being withdrawn.

To differentiate these updates we use the *anchor prefix* to detect such unexpected routing changes. An anchor prefix is a statically nailed down prefix belonging to the Beacon AS or the Host AS. Such a prefix can contain live hosts and does not experience any of the routing changed due to our beacon prefixes. Anchor prefixes serve as collaboration points to be used with our beacon prefix to observe unrelated routing changes and when there are no such routing changes at that AS, no routing updates associated with that prefix can be observed [2].

The set of updates is cleaned by deleting the signals which were not caused by the withdrawal. Another cause for these signals is routers which are not configured to support route flap damping. Route flap damping is explained in more detail in section VIII.

Another attribute which we analyze is the amount of time it takes between the first message we can see after the beacon has been withdrawn and the last message that we see in response to the message being withdrawn. For any beacon event there will be some neighbor that sends the associated update first and some neighbor that takes the longest, this is called the relative convergence period. For instance if there is some neighbor which reports the first event at 1076450513 (2004- 02- 10 17:01:53) and the last neighbor reports that event at 1076450539 (2004- 02- 10 17:02:19), then the relative convergence period for this event is 26 seconds. Our analysis for the relative convergence period is explained in section VII.

V. PATH EXPLORATION

As a path vector protocol, BGP undergoes path exploration after router changes. During the convergence period, a router may send multiple updates before eventually settling down on a new stable path, which increases the number of updates which propagate throughout the network [3]. Without path exploration, we would expect to see 12 updates per day for each beacon, 6 for the announcements and 6 for the withdrawals. However, we found nearly four times that number on average due to path exploration. This was true for for RIPE and PSG beacon prefixes. This is due to routers not being set up with route flap damping.

VI. ROUTE FLAP DAMPING

The BGP uses two main algorithms to determine the shortest path between nodes. These are the distance-vector routing protocol and the link-state protocol. One major problem with the distance vector protocol is the count to infinity problem.

The count to infinity problem happens when a prefix is withdrawn and other prefixes around it had that prefix included in their paths. This stems from loops occurring in paths at routers. If A tells B that it has a path somewhere, there is no way for B to know if that path actually has B in it. Suppose both A and B have a path to C, if C is withdrawn, A will have in its routing table a path to C through B and will tell B about this path. B will then see that it's path to C is down but that it has a path to C through A which is still up (which actually includes B), B will switch to that path and inform A that its path to C is down. This will then reverse the process which will continue until it reaches infinity (in which case the algorithm corrects itself due to the "Relax Property).

This event or any short instances of it will cause a number of updates to be propagated to RouteViews. There are a few solutions to this problem, one of which is route flap damping. This mechanism is aimed at reaching routing stability in the BGP. Flap damping punishes unstable routes or routes that change frequently by suppressing them. This mechanism has in accordance with it a minimum route advertisement counter which is aimed to deter unstable routes from being used on a short time scale. It specifies a minimum amount of time a router needs to wait before sending consecutive updates referring to the same prefix to the same neighbor. [2]

For the long term scale route flap damping keeps a penalty value associated with each route received from its neighbor. Whenever the route changes, a penalty value is added; this value decay exponentially over time. If the penalty value ever exceeds the cutoff threshold, the route is suppressed or is no longer available to forward traffic. Subsequently, if any updates for this router are received, they will not be propagated. Once the penalty decays to below the reuse threshold, the route is considered usable again [2].

For our analysis of route flap damping we examined announcements which were sent in very short intervals containing the same messages. We took the total number of distinct messages within one withdrawal interval and the total number of

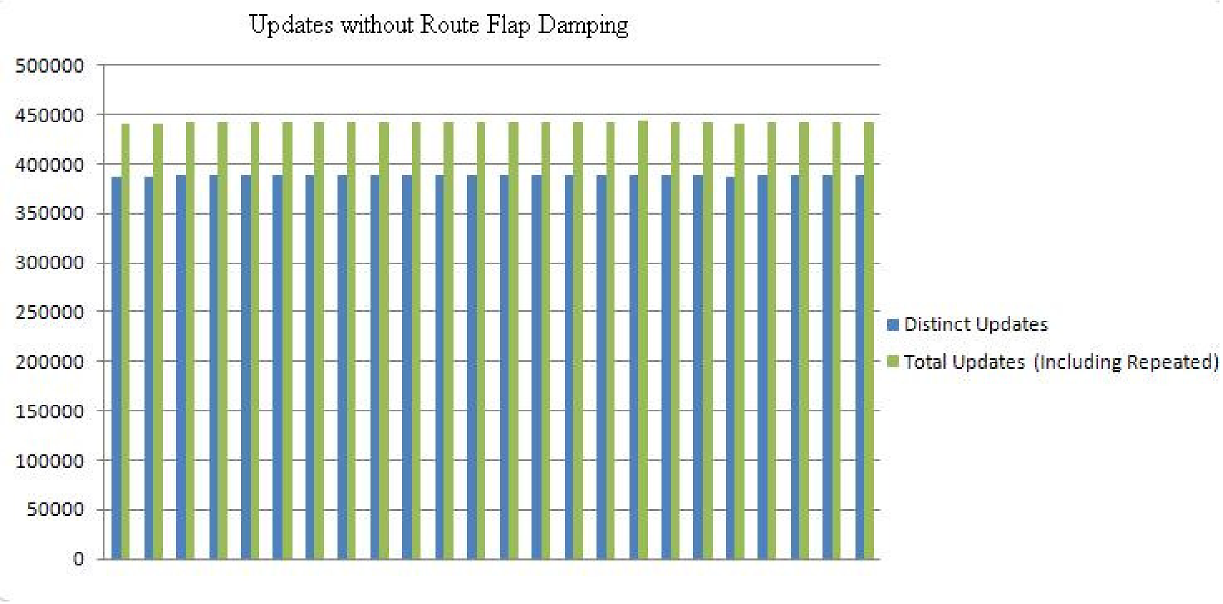

Figure 3: Distinct and Global Updates vs Total Number of Updates

announcements associated with the anchor prefix for that same interval and used them to predict the total number of announcements within that interval. The difference is the amount of messages which were needlessly sent. This would contain all of the repeated and unnecessary announcements due to route flap damping not being enabled at some routers. We found that these numbers were correlated to predict the number of repeated messages with a coefficient of regression of 98%.

VII RELATIVE CONVERGENCE PERIOD

We next analyzed the relative convergence period which was associated with each beacon withdrawal period. This is the time between the first and last message that we see at RouteViews in response to a beacon being withdrawn. For any event there will be some neighbor that sends the associated update first and some neighbor that takes the longest. This reveals information about the BGP and how long it takes messages to propagate through it.

Ideally it would be helpful to know the end-to-end convergence times of a message, however this would require clock synchronization on a lot of different servers [2]. This task is difficult to fulfill, so another measure is to use the relative convergence time. We have run analysis on these times across all of our beacon events and have formulated a way to predict the amount of time this period will take based on the number of total announcements which come in. This prediction is accurate to 91%. Of our experiments conducted on 4 days with a mean relative convergence of all events of 25.91, we predict to an accuracy of 91% that the mean relative convergence on the fifth day will be 27.24. See figure 4 for more details.

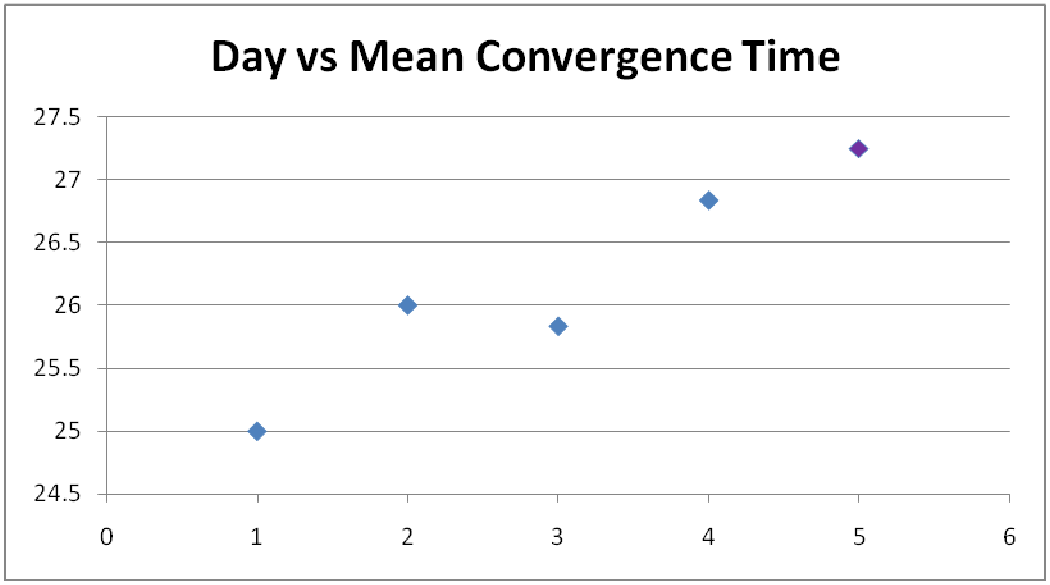

Figure 4: Day vs Mean Convergence Time

VIII. CONCLUSION

This paper describes BGP beacons from both RIB and PSG that have been set up for public use to analyze BGP behaviors. We have described the functions of the announce and withdrawal of these beacons and how they can help to analyze how routers update their path and how they send out updates to other routers informing them of these updates. We have presented how to gather data on updates in the BGP using RouteViews [4] and then how to process the raw data from that point on. We have limited our analysis to the PSG beacons.

We then explain the process of route exploration that takes place at each router when a node goes down and how some of those updates within the BGP could have been eliminated if route flap damping had been used. We then look at the total convergence times between in events from the first message we receive to the final message in an event and the correlation between the number of messages that come in and the total amount of time that takes.

References

[1] Ricardo Oliveira, Beichuan Zhang, Dan Pei, Lixia Zhang, *"Quantifying Path Exploration in the Internet"*, to appear in IEEE/ACM Transactions on Networking, June 2009

[2] Z. M. Mao, R. Bush, T. Griffin, and M. Roughan, “BGP beacons,” in ACM SIGCOMM Internet Measurement Conference (IMC), 2003.

[3] Ricardo Oliveira, Rafit Izhak-Ratzin, Beichuan Zhang, Lixia Zhang,*"Measurement of Highly Active Prefixes in BGP"*, in IEEE GLOBECOM, St. Louis, USA, November 2005

[4] “The RouteViews project,” http://www.routeviews.org/.

[5] “PSG Beacon List.” [Online]. Available: http://www.psg.com/~zmao/BGPBeacon.html

[6] “The RIPE Routing Information Services,” http://www.ripe.net/ris/docs/beaconlist.html