Determining the Cause and Frequency of Routing Instability with Anycast

James Morrow Hiebert
Dept. of Computer and Info Science
University of Oregon,
Eugene, OR 97403-1202
jamesmh@cs.uoregon.edu

Peter Boothe
Dept. of Computer and Info Science
University of Oregon,
Eugene, OR 97403-1202
peter@cs.uoregon.edu

Randy Bush
Internet Initiative Japan, Inc.
Kapa'au, HI 96755
randy@psg.com

Lucy Lynch
Computing Center
University of Oregon
Eugene, OR 97403-1212
llynch@uoregon.edu

1 Introduction

We present a methodology by which an autonomous system (AS) can estimate the stability of their BGP routes without requiring access to restricted BGP data. We demonstrate a novel measurement approach using DNS anycast as an indicator of routing instability. Using this method, even end-users may monitor their ISP's routing stability, without the continual use of expensive ICMP traceroutes or access to generally restricted routing information. We then perform a case study from within a large ISP in order to quantify and determine the cause of the routing instability. We correlate external and internal BGP events with variations in the final anycast destination. We conclude that anycast is extremely sensitive to anomalous BGP events and that by monitoring anycast it may be possible for large networks to receive early warning of BGP instability.

Our contributions are as follows:

- We present a novel methodology for detecting routing instability. This technique is highly sensitive and can be an early indicator of bad routing behavior,
- we provide a case study which demonstrates good correlation between routing events that indicate likely problems and anycast route switching events, and finally
- we show that switching events are relatively infrequent.

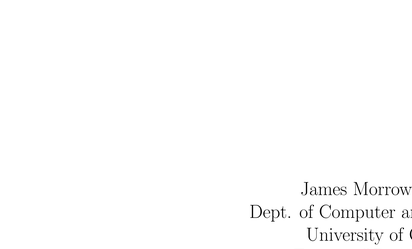

FIGURE 1.1:

This poster is an abbreviation; for further details and clarifications on our study, please refer to the complete article [1].

2 Methodology

The contributions of this work required data from two primary sources. We performed a case study, demonstrating our methodology. This experiment consisted of twenty-six days of anycast route monitoring. Secondly, to validate our results, we compared our data to a collection of routing announcements from the same AS as our anycast monitoring.

2.1 Anycast Route Monitoring

Our own experiments using anycast as a tool provided the primary data set used to determine and demonstrate the presence or absence of routing stability. Because anycast routing is more sensitive to routing changes it provides an excellent "canary in the coal mine" for routing instability.

Our experiment consisted of obtaining one particular datum, distributed across time and space (location in the Internet). We placed a collector at various hosts across the Internet. Each collector would make a single query to each of the anycasted DNS root servers (c, f, i, j, k and m) at intervals of 2 and 20 seconds for UDP and TCP queries respectively. The query, which appeared as `$ dig @X.root-servers.net. hostname.bind chaos txt`, was crafted to determine which X-root cluster received the query. By knowing which cluster received the query we were able to detect the frequency of route switches to each cluster at a frequency of half that of our query interval (4 and 40 seconds for UDP and TCP respectively).

Output of the script would consist of a series of 4-tuples: timestamp, DNS root (e.g. c, f, i, j, k, or m), transport, and DNS root-server. One such set of output was received from each of our collectors at regular time intervals.

2.2 eBGP and iBGP Data for Validation

When a routing switch is detected using our anycast data collectors, for our analysis we require router announcement data to determine whether the switch occurred due to external or internal factors.

To analyze the source of the switches, we utilized existing collections of routing announcements from the Route Views project[4] and from a large ISP with whom we were cooperating to perform this study. As shown in figure 1.1, routing announcements were collected by setting up a router as a peer which would only receive BGP announcements (and record them) without sending announcements back. We obtained both external BGP (eBGP) and internal BGP (iBGP) data through two separate router sessions.

Finally we correlated the switching events discovered through our anycast experiment with routing events which were evident in the cache of routing data.

2.3 Correlation of Routing Events

We consider a time-clustered group of up to three types of measurable events to constitute a single logical event: one or more eBGP events at the ISP's peering routers, one or more iBGP announcements by the ISP's edge routers, or one or more switches in the responding anycasted cluster seen by our anycast probe (henceforth referred to as a *switching event*). The set of routing announcements considered is limited to announcements of the relevant anycast prefixes.

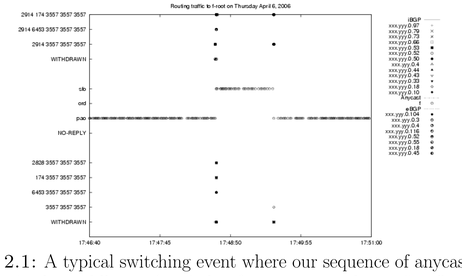

FIGURE 2.1: A typical switching event where our sequence of anycast probes (the points in the middle rows) switches from one root cluster "pao" to another "sfo" for a brief period of time. The graph shows a 5 minute slice of all three data sets: eBGP traffic on the top, our anycast probes to f-root in the middle, and iBGP traffic on the bottom.At the same time as the initial switch and the switch back there is visible eBGP traffic (above) and iBGP traffic (below) where routers are announcing best path changes. The x-axis is time, the y-axis is separated according to AS-path announced, and each point style represents a different router.

We correlate the three types of occurrences within a similar time frame as one routing event. We define a routing event to be any of the three aforementioned occurrences separated by no more than a time window of t seconds. In other words, each *event* is separated by a period of quiescence greater than or equal to the time window. This windowing method was used by Wu et al. [5] and 70 seconds was determined to be an appropriate window. Thus for all of our experiments we used a t value of 70 seconds.

With three boolean attributes—visible **iBGP** and **eBGP** traffic of the relevant prefixes, and an **anycast switch**–there are seven combinations of the presence or absence of each attribute (the eighth being where nothing happens). Our primary concern is with the four cases where an anycast switch is present (iBGP/eBGP, iBGP only, eBPG only, and neither). We also include the remaining three cases in our tables.

3 Results

After having performed our experiments and collected all data, we analyzed one dataset which consisted of the twenty-six day period from March 17, 2006 to April 12, 2006.

3.1 Quantification of Switching Behavior

We searched the twenty-six day period and classified each event according to the corresponding traffic found at that time (as defined in section 2.3). The resulting counts are shown in table 3.1.

Early results (not shown) detected a substantial volume of routing traffic which did not correspond to an anycast switching event. We detected 135 eBGP events, 44 iBGP events, and 3 eBGP/iBGP events. This seemed to indicate that there was a low correlation between routing traffic and switching events. However, in examining the data directly we noticed that a significant percentage of the routing messages (98.5% eBGP, 77.3% iBGP, and 100.0% eBGP/iBGP) were reannouncements of existing paths. Presumably these re-announcements were due to MED values changing or some other relatively insignificant BGP chatter.

iBGP	Anycast	eBGP	sum of events	c	f	i	j	k	m
		X	2	0	1	0	1	0	0
X			10	8	1	1	0	0	0
X		X	0	0	0	0	0	0	0
	X		4	2	2	0	0	0	0
	X	X	2	0	0	0	2	0	0
X	X		8	8	0	0	0	0	0
X	X	X	5	0	1	0	4	0	0
Total Switching Events			19	10	3	0	6	0	0

TABLE 3.1

The results in table 3.1 show only eBGP or iBGP events for when a *new or different* path is announced. When the insignificant BGP chatter is eliminated, our data shows a total of 31 important routing events. Of these events 19 of them (61.3%) were detected by a switching event in the anycast cluster. This demonstrates that our technique is very effective at detecting routing changes.

It is also important to note that during many of these switching events there was no loss of service to the anycast DNS servers, and thus even the most savvy network operators would not detect a change.

3.2 Ability to Detect Obscure Routing Events

In addition to detecting most routing events which would have been found by monitoring BGP messages, our technique also detects some routing events which would be otherwise impossible to discover. Tables 3.1 shows several switching events for which there was no corresponding eBGP or iBGP traffic. Presumably these switches were due to routing changes further down the AS path. Our method detects some such events, while directly monitoring one's own BGP feeds would fail to do so.

Another reason that most of these events are undetectable is that, even though the routes to the anycast DNS servers switch, the user continues to receive service. In only 7 of our switching events did the "dig" application register a NO-REPLY from the anycast DNS server. Throughout the entire experiment a total 68 NO-REPLY events were registered. There is no evidence of strong correlation between NO-REPLYs and anycast switches. This is consistent with research by Li et al. [2] who conclude that BGP activity has little correlation with packet loss.

3.3 Ability to Foresee Behavior

In order to test whether anycast switches can be correlated with more than just BGP announcements for the anycasted DNS prefixes, we examined the volume of BGP traffic and its temporal proximity to the detected switching events.

Figure 3.1 shows the average rate of external and internal BGP messages (which announce a new or different path) relative to the time of all anycast switches. The overall average rate of BGP messages is also plotted (the horizontal lines at the bottom of the graphs). The big spike in the center of the graph indicates that, right around the time of the anycast switches, we also observed a ten-fold increase in BGP traffic. *This indicates that anycast switches are a good forecaster for BGP traffic, and that DNS anycast switches are not isolated events.*

This lack of isolation lends credence to the possible implications of the conclusions by Rexford et al. [3]. Their research demonstrates that routing to popular destinations is generally stable and that most Internet routing instability stems from little-used routes. Our work in monitoring these popular DNS routes is consistent with these conclusions. We see that when routing to one of these popular destinations changes–an unusual event–that the event is part of a more major routing occurrence.

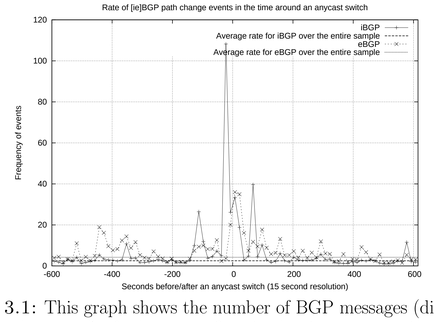

FIGURE 3.1: This graph shows the number of BGP messages (divided into internal and external) of relevant prefixes in the time surrounding switching events. This graph is constrained to those messages announcing new or different paths.

References

[1] J. Hiebert, P. Boothe, R. Bush, and L. Lynch. Determining the cause and frequency of routing instability with anycast. In *Proc. of Asian Internet Engineering Conference (AINTEC) 2006*, November 2006.

[2] J. Li, R. Bush, Z. Mao, T. Griffin, M. Roughan, D. Stutzbach, and E. Purpus. Watching data streams toward a multi-homed sink under routing changes introduced by a BGP beacon. In *Proceedings of Passive and Active Measurement Conference 2006*, March 2006.

[3] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang. BGP routing stability of popular destinations. In *ACM SIGCOMM Internet Measurement Workshop*, 2002.

[4] University of Oregon Route Views Project. http://routeviews.org, 1997-present.

[5] J. Wu, Z. M. Mao, J. Rexford, and J. Wang. Finding a needle in a haystack: Pinpointing significant BGP routing changes in an IP network. *NSDI*, 2005.