

BGP Behavior Analysis During the August 2003 Blackout

Z. Wu, E. Purpus, J. Li

Computer and Information Science Department, University of Oregon

120 Deschutes Hall

1202 University of Oregon

Eugene, OR 97403-1202

{zww|epurpus|lijun}@cs.uoregon.edu

1. Introduction

By studying BGP under the stress of large-scale events such as large-scale power outages, worm attacks, and other events that affect Internet routing, we can begin to understand how BGP reacts to adversity. While many questions regarding BGP's global behavior will still remain unanswered, studying BGP under specific events will expose the resiliency of BGP under the often hostile environments of the Internet.

The East coast blackout in 2003 is an event which gives us an opportunity to study BGP behavior under significant stress. Shortly after 4 p.m. EDT on August 14, 2003, a major power outage simultaneously hit dozens of cities in the eastern United States and Canada. The blackout affected the connectivity of 3,175 networks according to the Renesys report by Cowie *et al* [1], clearly a large-scale event. Since so many networks became unavailable during this event, the routes to these networks were also affected, thus making the event significant within the context of studying BGP behavior.

Our goal in this research is to characterize and analyze BGP behavior surrounding the blackout event and to deduce the resiliency and responsiveness of BGP during such a large-scale event. We hope our BGP behavior analysis techniques in this blackout event can be generalized and applied towards analyzing BGP's resiliency in other security, mis-configuration, or outage events. The BGP behavior during this blackout event can be similar in many aspects to those events that also affect BGP routing through disrupting the service of many Internet routers, such as the distributed denial of service (DDoS) attacks.

In this paper, we analyze BGP behavior from two angles: global-level statistical BGP behavior analysis and prefix-level analysis. Global level analysis may reveal the existence of anomalies or differences in BGP behavior and prefix level analysis may reveal details about the nature of an anomaly or difference. Furthermore, unusual behavior at the prefix level may be hidden at global level by the aggregate nature of the global statistics.

Our paper makes the following contributions. We propose several global metrics to characterize macroscopic BGP behavior and discover that although no significant difference is observed in the number of total BGP updates or other metrics, the number of explicit withdrawal messages makes the blackout stand out. We also introduce per-prefix AS-path graphs to study BGP behavior from a microscopic point of view, and demonstrate that this is a promising approach to observe prefix-level routing changes or diagnose routing problems. To the best of our knowledge, we are the first to introduce this technique.

2. Methodology

Using BGP update data collected before, during, and after the blackout event, we gather statistics that describe the state of the Internet at both the global and prefix level.

Our data set comes from the University of Oregon RouteViews archives [4]. By peering with ASes from around the world, RouteViews monitors continuously collect the BGP update messages received from these peers, resulting in a huge archive of BGP updates regarding reaching prefixes all over the Internet.

BGP peers exchange update messages between each other to update the path information for particular prefixes. These messages come in the form of either announcements or withdrawals, where paths are added or removed respectively. In announcement messages, each prefix updates include an AS-path to describe the path taken to reach that prefix. This provides the path information necessary to construct prefix-level routing graphs.

Per-prefix AS-path graphs allow us to look at BGP behavior from a microscopic viewpoint. Starting with an empty graph, we filter all the update messages for a particular prefix during a given window of time. For every announcement message, we add the path to the graph if it does not already exist, based on the AS-path field for this prefix. The router-id of the announcement is also associated with that path to differentiate between multiple routers from the same AS. Note that every node on the path is an AS. For every withdrawal message, we match the router-id of the withdrawal with the path in the graph associated with the same router-id, and remove this path from the graph. At this point, any unconnected AS nodes are removed from the graph.

This graph construction process allows us to dump the current state of the graph at any point in time. In our analysis, we create per-prefix AS-path graphs every five minutes.

3. Evaluation

3.1 Global Statistics

Number of BGP Updates

The number of BGP updates has been a primary metric in observing the dynamics of BGP in previous major events such as worm attacks [2], [3]. During the blackout event, we observe no significant changes in the volume of updates. Because the power outage is an ON/OFF event, networks experiencing the outage do not cause *consistent* routing oscillations. In contrast, networks under worm attack experience congestion or peculiar traffic patterns for an extended period of time, resulting in continuous route flapping and a much higher number of observed update messages.

Number of Explicit Withdrawals

Although power backup systems are common, we expect that the power outage will have caused some networks to become unreachable due to the sustained nature of the outage, extending beyond backup capabilities. BGP sessions that involve networks affected by the blackout can thus be broken and result in many withdrawal updates. The number of networks affected during the blackout directly affects the number of withdrawal updates. In the case of the blackout, a spike of almost 200% in the number of withdrawals occurred immediately after the power outage.

Interval of BGP updates

The interval between BGP updates (also called inter-arrival time) can also indicate the level of BGP activity and the general stability of routing paths. When no changes occur to the routing paths, no updates will be exchanged. Conversely, frequent path changes can lead to frequent routing information exchange through BGP updates, thus leading to smaller update intervals. In our measurement, we define the update interval as the time between two updates for the *same* prefix from the *same* RouteViews peer.

We analyze the interval distribution during the blackout period and compare it to two randomly chosen periods of the same length. Our analysis found three spikes for all three periods, around interval values of 30, 60, and 90 seconds, and most intervals (about 70-80%) are less than 100 seconds. Both the blackout period distribution and the reference distributions have these spikes, and do not differ significantly from one another. We believe the spikes are the effect of the configuration parameter – Minimal Route Advertisement Interval (MARI) timer – that was introduced to control BGP traffic overhead on routers by defining the minimal amount of time between advertisements from a single router for a particular prefix.

Top Active Prefixes

Due to the ON/OFF nature of the power outage event, we suspect that during the blackout period those affected network prefixes were *not* significantly more active than those unaffected prefixes. To evaluate our hypothesis, we study the hourly rate of BGP updates for the top 1,000 active prefixes and find that there are no significant variations during the blackout period.

3.2 Prefix-Level Statistics

Global-level BGP analysis can show certain changes during a large-scale failure such as the East coast blackout event, but such an analysis is aggregated over all networks. In order to pinpoint problematic prefixes or perform any prefix-level operations, it is important to zoom in and watch BGP behavior at individual prefix level. In order to show prefix-level changes during the blackout, we construct prefix graphs every five minutes before, during and after the blackout, and record their characteristics at these points. We discuss the results for one particular prefix which displays common characteristics which we observed in many of the prefixes we studied.

Number of Nodes and Edges

Studying the change in the number of nodes during the blackout, we see a sudden drop in the number of nodes between the initial outage and the start of the power restoration. During the recovery phase, the paths are quickly re-announced and the number of nodes remains fairly steady thereafter.

In and Out Degree

Node degrees in a per-prefix AS-path graph indicates the connectivity of the graph. Variations in node degrees thus will reflect changes regarding how nodes are connected at different time. In a graph where nodes have high in-degrees, many ASes choose paths through a common AS. In a graph where nodes have out-degrees greater than one, multiple routers for a single AS choose different paths to reach a given prefix (particularly common for large ASes that span large regions).

During the blackout period, we see a slightly different pattern to the number of nodes: a sudden *increase* in the in-degree followed by a sudden decrease and then a return to the original value. These two sudden changes reflect the changes in connectivity caused by the blackout. Suddenly, 15 ASes all switch to paths that have a common AS as the next hop, but then as the total number of possible paths drops, the maximum in-degree also drops. The first spike in in-degree indicates that many BGP routers had that particular common AS as a secondary route. When the preferred routes became unavailable to these routers, many of them switched to the same route through this common AS. The recovery phase returns the graph to its original shape and maximum in-degree also returns to its original value.

4. Conclusions

While the Internet continues to thrive, the resiliency of the fundamental Internet routing infrastructure is not well understood. BGP as the *de facto* inter-domain routing protocol over the Internet must therefore be studied, especially when it is facing serious failures or attacks.

In this paper, we analyze BGP's behavior at the August 2003 East coast blackout, during which 3,175 networks lost their connectivity, and identify a number of metrics for analyzing BGP behavior. Global metrics studied include number of updates, explicit withdrawals, inter-arrival time, and top active prefixes over a time window of two weeks surrounding the blackout period. We also introduced per-prefix AS-path graphs as another analysis method, and studied changes in graph characteristics such as the number of nodes, in-degrees, and out-degrees. Our results show that during the blackout event, there was an apparent increase in the number of explicit withdrawals, and at the prefix level, the per-prefix AS-path graphs of some prefixes also show a sharp decrease in the number of edges and nodes, as well as changes in node degrees. Nonetheless, our results also show that BGP can recover from the blackout event in a timely manner. Overall, the negative impact of the power outage was limited only to the affected areas, and BGP has performed well during this large-scale power outage.

5. Acknowledgments

We would like to thank Lixia Zhang and Randy Bush for their helpful feedback about this research.

References

- [1] J. Cowie, A. Ogielski, B. Premore, E. Smith, and T. Underwood. Impact of the 2003 blackouts on Internet communications. Technical report, Renesys, November 2003.
- [2] J. Cowie, A. Ogielski, B. Premore, and Y. Yuan. Global routing instabilities during Code Red II and Nimda worm propagation. Technical report, Renesys, 2001.
- [3] M. Lad, X. Zhao, B. Zhang, D. Massey, and L. Zhang. An analysis of BGP update surge during Slammer attack. In *Proceedings of the 5th International Workshop on Distributed Computing (IWDC)*, December 2003.
- [4] University of Oregon Route Views Project. <http://antc.uoregon.edu/route-views/>.