Ensuring the Robustness of Internet Routing

Olaf Maennel
University of Adelaide, Australia
olaf@maennel.net

Steve Uhlig
University of Delft, Netherlands
suh@info.ucl.ac.be

Abstract— **Internet routing is probably the largest scale distributed calculation made on our planet. Its computation is based on routing protocols whose dynamics has been observed to be highly complex, and not well understood [1]. However, before trying to design new routing protocols we should understand and debug the existing ones.**

As a step towards improving the robustness of our routing system, we propose to first obtain an accurate picture of interdomain routing behavior. To achieve that we focus on developing an open source "BGP monitor and alarm system". In this position paper we argue that researchers and operators should work together on a *joint* set of tools that are capable of monitor the routing plane constantly.

How to monitor, interpret and predict the Internet routing system remains an unresolved problem in spite of considerable research effort during recent years (see [2] for an overview). Researchers and operators still lack appropriate tools to detect and locate problematic routing events or even anticipate planned routing changes.

The long list of problems associated with the Border Gateway Protocol (BGP) include slow BGP convergence [3], [4], interactions between different routing protocols [5], address hijacking [6], misconfiguration [7], unintended BGP states [8], and even diverging routing conditions [9]. However, there are no automated solutions in existence capable of handling such questions.

Obviously each problem possesses its own set of peculiarities, but the common underlying theme of these open questions is that they rely on the same data gathered from the network (such as routing information, and traceroutes). One of the more fundamental limitations is simply that it is hard to obtain the appropriate data. Often not enough data is available to enable accurate investigation into, and solution of BGP related issues (e.g., [10]). In such cases inference techniques may help to approximate the missing measurements. These techniques are in itself open research problems (e.g., [11]).

Ironically, a complementary problem stems from the huge amount of data that is already available. A number of efforts have already been undertaken by the community to collect necessary information (for BGP collection take for example RouteViews [12] and RIPE [13] projects). Here the issue arises because there are no tools available to appropriately aggregate the information so that the user can interpret the data.

Our aim is to work towards a methodology (and tool) that is capable of measuring and interpreting the routing system. With collaborators we are developing an open source "BGP alarm system", which is an automated tool that gathers and aggregates the data from various sources, in a manner that is beneficial for the user. Such a system will facilitate the analysis of routing behavior, which is a prerequisite for future changes to the routing system.

Besides long-term research goals it is essential for operators to monitor the routing system. Questions such as "how are my prefixes seen by others?", "Am I reachable?", "Send me a notification, if I misconfigured something?", "What paths do other ASes use to reach my customers?" and "how does it see the world?", "How much route diversity is available?", etc. must be answered constantly to detect routing problems.

To the best of our knowledge no system exists today to leverage all the routing data that is publicly available. Such a system might be used not only to better understand the global behavior of Internet routing, but also to diagnose its problems. We propose to build such a system that would create operational value to the available BGP data that lies dormant today.

Imagine a system that is scalable and equipped with an appropriate methodology to cluster information from various places in the network. So far, we envision a set of small tools, called "plug-ins", around some kind of "routing database". One of those "plug-ins" is responsible for downloading BGP data, another will clean the data (e.g., remove session resets, check the correctness of timestamps, detect monitor outages), yet another will cluster updates from various peers to locate routing instabilities. The plug-ins are building blocks and my depend on each other. For example a cluster plug-in may rely that the data was downloaded and is cleaned from session resets, while a plug-in that detects hijacking only needs new data as fast as possible. The central core handles those dependencies.

In summary we believe that discussing about the next generation routing protocol is pointless before we understand the problems of the current versions. Among the critical issues with routing that we have to understand is the issue with robustness. On the theoretical side a lot of progress has been made (e.g., [14], [15]), on the practical side a lot of work remains.

How can we diagnosing global routing problems (instabilities, inconsistencies), find out misbehaving ASes (prefix high-jacking or router misconfigurations), or understanding the global impact of specific events so as to identify improvements to the BGP protocol?

We urge researchers, operators and the regional route registries to stick together and help to contribute to an open source "BGP alarm system". It provides a well-debugged, scalable and validated source of data for all researchers and operators.

ACKNOLEDGEMENTS

This work is inspired by Randy Bush, Gert Dring, Anja Feldmann, Wolfgang Mühlbauer, Henk Uijterwaal, and the RIPE NG team. Many thanks also to Matthew Roughan. This work is partially supported by the Australian Research Council (ARC) grant DP0557066.

REFERENCES

[1] N. Feamster, T. Griffin, J. Rexford, and R. Bush. WIRED – Workshop on Internet Routing Evolution and Design, 2006.

[2] T. G. Griffin Interdomain routing links. http://www.cl.cam.ac.uk/users/tgg22/interdomain/.

[3] Z. M. Mao, R. Bush, T. G. Griffin, and M. Roughan, "BGP Beacons," in *Proc. ACM IMC*, 2003.

[4] Z. M. Mao, R. Govindan, G. Varghese, and R. Katz, "Route flap damping exacerbates Internet routing convergence," in *Proc. ACM SIGCOMM*, 2002.

[5] R. Teixeira, A. Shaikh, T. G. Griffin, and G. M. Voelker, "Network sensitivity to hot-potato disruptions," in *Proc. ACM SIGCOMM*, 2004.

[6] P. Boothe, J. Hiebert, and R. Bush, "How Prevalent is Prefix Hijacking on the Internet?," *NANOG 36*, February 2006.

[7] D. Wetherall, R. Mahajan, and T. Anderson, "Understanding BGP misconfigurations," in *Proc. ACM SIGCOMM*, 2002.

[8] T. G. Griffin and G. Huston, "BGP Wedgies," 2005. RFC 4264.

[9] T. G. Griffin and G. Wilfong, "Analysis of the MED Oscillation Problem in BGP," *In Proceedings of the International Conference on Network Protocols*, 2002.

[10] R. Teixeira and J. Rexford, "A measurement framework for pin-pointing routing changes," in *Proc. ACM SIGCOMM Network Troubleshooting Workshop*, 2004.

[11] A. Feldmann, O. Maennel, M. Mao, A. Berger, and B. Maggs, "Locating Internet Routing Instabilities," in *Proc. ACM SIGCOMM*, 2004.

[12] University of Oregon RouteViews project. http://www.routeviews.org/.

[13] RIPE's Routing Information Service. http://www.ripe.net/ris/.

[14] T. G. Griffin and J. L. Sobrinho, "Metarouting," in *Proc. ACM SIGCOMM*, 2005.

[15] N. Feamster, R. Johari, and H. Balakrishnan, "Implications of autonomy for the expressiveness of policy routing," in *Proc. ACM SIGCOMM*, 2005.