Better Abstractions for Better Routing?

J.E. Van der Merwe
AT&T Labs - Research

A significant part of the complexity of Internet routing can be attributed to the complexity and obscurity of existing router configuration "languages" [1], [2], [3]. This complexity is exacerbated by the fact that routing configuration is distributed across hundreds or thousands of network elements, with subtle interdepencies [4] that make it difficult to reason about routing in the network [5].

Recently proposed routing architectures [6], [7], [8], [9] espouse a logically centralized approach to network management and routing configuration management. It has been argued that a centralized approach will lead to a simplification of routing protocol configuration, not only because it is centralized but also because it allows the network operator/designer to think about the network in a holistic fashion rather than as a large collection of routers.

While in full agreement with these sentiments, herewith some observations:

"Logically centralized" normally implies "physically distributed". Indeed for redundency, performance and scalability reasons, locically centralized approaches will by neccesity have to be physically distributed to some degree. This means that care must be taken even in this new paradiagm to ensure that the desireable properties of centralized reasoning is maintained and that the problem is not simply moved from one set of distributed network elements to another.
I would further argue that, together with centralized/holistic reasoning about routing configuration, there should be an emphasis on raising the level of abstraction allowing operators to think about "what" policies and functionality they are trying to realize, rather than about "how" to realize them. Coming up with the right level of abstraction that enhances, rather than obscures, understanding seems to be the biggest challenge in this regard.

In ongoing work we are starting to address some of these challenges:

Distributed RCP/IRSCP: An obvious benefit of the fully distributed nature of routing configuration in the current Internet is the fact that individual routers are fully autonomous entities that can function independently from any other entities in the network. For example, a set of routers that get isolated from a part of the network can continue to function among themselves, albeit in a greatly reduced universe. With a simplistic centralized approach, where routers rely on some other entity to provide them with the information required to forward traffic, it is not difficult to come up with failure scenarios where a set of routers isolated from their intelligence is prevented from this type of graceful failure modes. In recognition of this fact we have built a distributed IRSCP which allows the "distributedness" of a deployment to match the appropriate reduncy requirements, while at the same time ensuring the consistent decision making amongst the different replicas. With the basic substrate and a rudimentary mechanism in place to allow holistic/centralized reasoning and configuration, we are now exploring different ways to deal with configuration management that would be more amenable to dynamic connectivity management tasks [10].

Intent-based Network/Routing Configuration: We take both an bottom-up and a top-down approach to raising the level of abstraction for routing configuration management. First, in the more constrained context of our IRSCP work we have developed a high level parameterized interface that allows operators to perform reasonably complex functions without having to deal with the resulting (underlying) IRSCP configuration [10]. The challenge is, however, to generalize this to "any" routing related function that an operator might want to perform in the network [11]. Towards this end we have started work on analyzing the actual network (and routing) configuration changes that operators perform over time to understand whether any patterns can be identified that might be candidates for intent-based routing configuration.

REFERENCES

[1] A. Feldmann and J. Rexford, "Ip network configuration for intradomain traffic engineering." IEEE Network Magazine, September/October 2001.
[2] D. Maltz, G. Xie, J. Zhan, H. Zhang, A. Greenberg, and G. Hjalmtysson, "Routing design in operational networks:a look from the inside." Proceedings of ACM SIGCOMM, August 2004.
[3] N. Feamster and H. Balakrishnan, "Detecting bgp configuration faults with static analysis." 2nd Symposium on Networked Systems Design and Implementation (NSDI), May 2005.
[4] T. G. Griffin and G. Wilfong, "On the correctness of ibgp configuration." Proceedings of ACM SIGCOMM, August 2002.
[5] N. Feamster and J. Rexford, "A model of bgp routing for network engineering." Proc. ACM SIGMETRICS, June 2004.
[6] O. Bonaventure, S. Uhlig, and B. Quoitin, "The Case for More Versatile BGP Route Reflectors." Internet Draft draft-bonaventure-bgp-route-reflectors-00.txt, July 2004.
[7] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and J. van der Merwe, "The Case for Separating Routing from Routers." FDNA Workshop, Aug 2004.
[8] M. Caesar, D. Caldwell, N. Feamster, J. Rexford, A. Shaikh, and J. van der Merwe, "Design and implementation of a Routing Control Platform," in *Proc. NSDI*, 2005.
[9] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang, "A clean slate 4D approach to network control and management," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 5, 2005.
[10] J. V. der Merwe et.al., "Dynamic Connectivity Management with an Intelligent Route Service Control Point." ACM SIGCOMM Workshop on Internet Network Management (INM), October 2006.
[11] M. Caesar and J. Rexford, "Bgp policies in isp networks." IEEE Network Magazine, November 2005.