BGP Route Reflection Revisited

Jong Han Park, AT&T Labs
Ricardo Oliveira, ThousandEyes
Shane Amante, Level-3 Communications Inc.
Danny McPherson, Verisign, Inc.
Lixia Zhang, University of California at Los Angeles

ABSTRACT

The original BGP design requires that all BGP speakers within an autonomous system be directly connected with each other to create a full mesh, and BGP update messages be propagated to directly connected neighbors only. This requirement leads to BGP session scalability problems in networks with large numbers of BGP routers. Route reflection was proposed as a quick fix to address this BGP session scalability problem and has been widely deployed in the operational Internet without a thorough analysis of its pros and cons. In this article, we first provide an overview of the route reflection design, summarize the discoveries from published literature, and discuss the trade-offs in using route reflection as compared to using a fully connected i-BGP mesh. Then we show that well engineered route reflector placement can overcome certain drawbacks, and that a few issues remain open for future study.

INTRODUCTION

In the original Border Gateway Protocol (BGP) design, fully meshed i-BGP sessions among all BGP routers in an autonomous system are used to disseminate BGP updates within one autonomous system. BGP update messages are forwarded only to directly connected neighbors to prevent the update messages looping. In a network with large numbers of BGP routers, this full-mesh requirement results in a large number of BGP sessions at each router. Furthermore, since BGP sessions are managed through manual configurations, this full-mesh requirement also leads to configuration changes at all routers whenever a router is added or removed.

Route reflection [1] was developed in 1996 as one of the two proposed solutions to address the above mentioned BGP scalability problem; the other one is AS confederations [2]. Between the two, route reflection has seen a larger deployed base. However, the design of route reflection did not go through thorough analysis studies before its deployment widely rolled out more than 10 years ago. Only recently have several studies appeared that analyze various impacts of route reflection on overall routing system performance. The results from these studies show that route reflection may potentially decrease the network's robustness against failures, introduce delayed routing convergence, reduce path diversity, adopt suboptimal routes, and even cause data forwarding loops.

In this article, we first provide a comprehensive overview of BGP route reflection's operations, and explain its pros and cons in detail. We then illustrate how one can use well engineered route reflector placement to overcome certain drawbacks in route reflection deployment and further scale the routing system, without any protocol or implementation changes. Finally, we identify remaining issues in achieving the goals of both efficient routing information dissemination and system scalability.

BGP ROUTE REFLECTION

In this section, we first present a brief review on BGP basics, followed by an overview of route reflection; interested readers are referred to [1, 3] for more detailed descriptions of BGP, including differences between i-BGP and e-BGP, and how route reflection operates. We then analyze the pros and cons of the basic route reflection design.

ROUTING IN THE INTERNET

The Internet is made of tens of thousands of different networks called autonomous systems (ASs), and BGP is used to communicate reachability information. Routers of different ASs set up BGP sessions in between to exchange BGP routing updates. Such BGP sessions are called *e-BGP* sessions, the operation of which is governed by routing policies. BGP sessions are also set up between routers within the same AS to exchange BGP routing updates, and these sessions are called *i-BGP* sessions.

In e-BGP, routers detect potential routing loops at the inter-AS level by inspecting the AS_PATH attribute carried in all BGP mes-

 0163-6804/12/$25.00 © 2012 IEEE

sages. A router will drop a BGP message if the AS_PATH in the message already contains its own AS number. To avoid routing loops in i-BGP, the original design requires that all BGP routers in the same AS be directly connected to each other via pair-wise i-BGP sessions and that the reachability information learned over one i-BGP session must not further distributed over other i-BGP sessions. This full-mesh i-BGP connectivity allows each BGP router to learn about reachability information directly from all other BGP routers in the same AS, eliminating the need to forward BGP updates learned from an i-BGP speaker to another i-BGP speaker, hence preventing potential routing loops. This standard full-mesh requirement works well for small-sized networks. However, this design does not scale as the number of BGP routers increases; the number of i-BGP sessions to be established and maintained within a given AS is

$$\frac{N \times (N-1)}{2},$$

where N is the number of BGP routers in the AS. Since operations such as creating or removing i-BGP sessions require operator intervention, this full-mesh i-BGP connectivity requirement also represents a high operational cost for large-sized networks.

To alleviate this i-BGP scalability problem, the vendor and operator communities quickly proposed two solutions in 1996: route reflection and AS confederations. Both solutions have been deployed in operational networks; in certain cases AS confederation deployment is combined with route reflection. Overall, route reflection has a wider deployed base and is the focus of this article.

Basic Operation of Route Reflection

The simplest model of route reflection deployment is to select one BGP router in an AS to be the route reflector (RR), and have all the other routers in the AS set up i-BGP sessions with the RR. The RR receives BGP update messages from each i-BGP speaker and forwards (or reflects) them to all the other i-BGP speakers. Because the RR forwards updates among i-BGP speakers, it eliminates the need for all i-BGP speakers to connect in a full mesh. To avoid a single point of failure, an AS generally sets up multiple RRs which are interconnected in a full mesh among themselves.

Figure 1 illustrates the difference between interconnecting i-BGP routers via full mesh and via RRs. Figure 1a shows an example of full-mesh i-BGP interconnections, where all i-BGP speakers are directly connected to each other. Figure 1b shows an example of route reflection deployment, where R_1 and R_3 serve as RRs and connect to i-BGP speakers R_2 and R_4, which are connected to both reflectors for redundancy. Since R_2 can learn R_4's BGP reachability information from the RRs and vice versa, R_2 and R_4 do not need to interconnect. R_2 and R_4 are *client* routers of R_1 and R_3. A client is an i-BGP speaker that connects directly to an RR to learn the reachability information collected by other routers in the AS. In the view of R_2 and R_4, R_1 and R_3 are *non-clients*. Note that R_2 and R_4

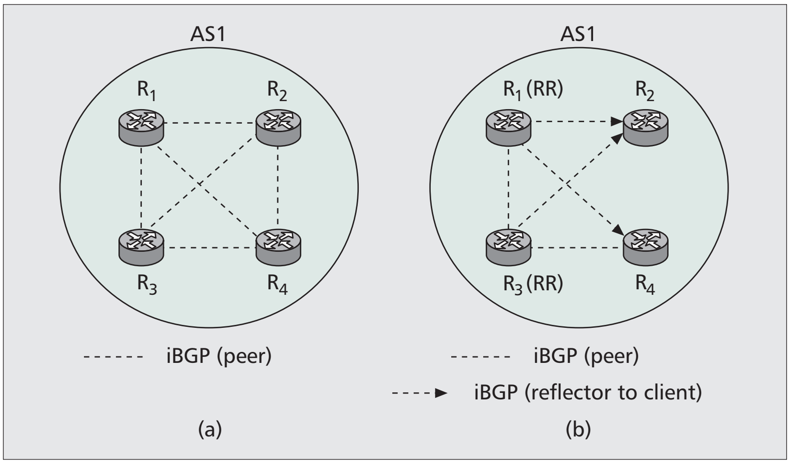

Figure 1. *Different i-BGP topologies: a) full-mesh i-BGP; b) i-BGP with route reflection.*

require no special configurations; they are not aware of R_1 and R_3 being RRs. Only R_1 and R_3 require configuration changes. The relation between R_1 and R_3 is non-clients, and they can pass the reachability information learned from one i-BGP speaker to others in the same AS.

Because RRs forward reachability information learned from an i-BGP speaker to another i-BGP speaker, routing messages travel more than a single i-BGP hop, and it becomes possible to create loops. To prevent such loops, two new attributes are added to BGP update messages: CLUSTER_LIST and ORIGINATOR_ID. An RR uses its router ID as the cluster ID. When forwarding a BGP update, if an RR finds its own cluster ID in the CLUSTER_LIST attribute of a received update, it discards the update; otherwise, it prepends its cluster ID in the CLUSTER_LIST attribute before forwarding the update. In addition, the first router that injects a routing update into the network will record its router ID in the ORIGINATOR_ID attribute. If a router receives an update with an ORIGINATOR_ID equal to its router ID, it discards the update. In Fig. 1b, R_2 will discard all updates reflected back to itself after checking that the ORIGINATOR_ID attribute contains its router ID.

Benefits of Route Reflection

Reduced Number of i-BGP Sessions — Route reflection can effectively reduce the number of i-BGP sessions in an AS. A non-RR router only needs to establish a small number (typically two for redundancy) of i-BGP sessions with the RRs. Although an RR router generally has a larger number of BGP sessions, one can control this number through well established engineering practices. Assuming a route-reflection-based AS with N i-BGP routers and K RRs, the number of i-BGP sessions for the network can be computed as

$$\frac{K \times (K-1)}{2} + \sum_{i=1}^{K} C_i,$$

where K is the number of RRs in the network and C_i the number of client i-BGP routers connected to the given route reflector RR_i. Typically

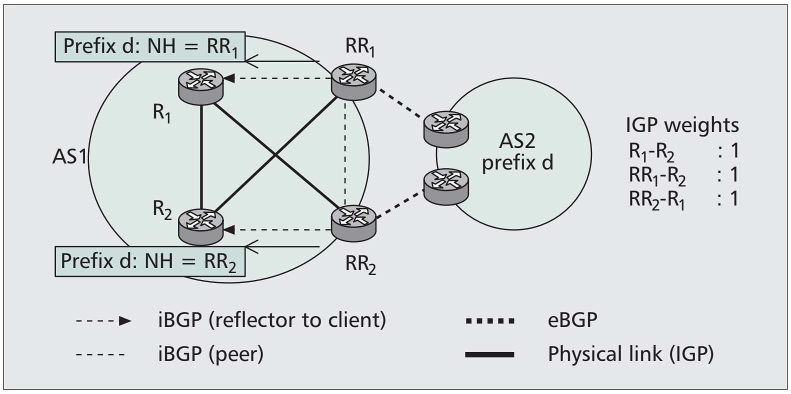

Figure 2. *Route reflection with data forwarding loop.*

K is a much smaller number than N in practice, making the total number of i-BGP sessions for a given RR much smaller than that of full mesh. For a given client, the number of i-BGP sessions is typically a constant (e.g., 2 for redundancy) regardless of network size.

Reduced Operational Cost — Creating, modifying, and removing BGP sessions require operator intervention. In the case of full-mesh i-BGP, any new router added to a network requires modifications to all the other routers' configurations. In the case of route reflection, adding or removing a client i-BGP router only requires configuration changes to the RRs to which the client connects, with no impact on the other routers.

Reduced RIB-in Size — A BGP router R maintains three different types of routing tables: RIB-in, Loc-RIB, and RIB-out. A RIB-in contains unprocessed (i.e., without applying import policy) routing information that has been advertised to R by each of R's BGP neighbors. After examining the reachability information and applying import policies across each RIB-in, the router decides a single best path for each destination D and stores this best path in Loc-RIB. R may or may not forward D's reachability information to its BGP neighbor routers depending on its export policy, but because the export policy to the i-BGP neighbors is mostly the same, R only needs a small number of RIB-outs (e.g., one per peer group that shares the same export policy) to store reachability information to be propagated to all its neighbors. On the other hand, the number of RIB-ins increases proportionally to R's number of BGP neighbors. If R has n neighbors each sending p prefixes, its total RIB-in size is on the order of $n \times p$. With full-mesh i-BGP sessions, n is the number of i-BGP neighbors in the full mesh. With route reflection, n for client i-BGP routers is the number of RRs to which the clients connect and is typically a small number.

Reduced Number of BGP Updates — With a significant reduction in the number of i-BGP neighbors, a client router naturally receives a significantly reduced number of updates. A route reflector R_r receives routing updates from all its neighbors, but since BGP only propagates the best path to each destination, R_r further propagates only those updates that change its best path selections. In sharp contrast to a full-mesh i-BGP setting where all BGP updates are propagated to all routers, RRs effectively shelter their client routers from a large percentage of incoming updates.

Incremental Deployability — Last but not least, route reflection allows coexistence of RRs with conventional BGP routers that do not understand route reflection. A conventional BGP router B can be connected to RRs as a client or non-client (in which case B must also be connected to all other RRs). This allows a network to perform a gradual migration from the full-mesh i-BGP model to the route reflection model.

Caveats of Route Reflection

Compared with the full-mesh i-BGP interconnections, although route reflection provides an effective alternative to address the i-BGP scalability problem, it also brings several negative impacts on overall routing system performance as listed below.

Robustness — With full-mesh i-BGP, a single router failure has limited impact on the rest of the network. That is, only the failed router cannot send or receive updates from the full mesh; the rest of the routers in the network are not affected. In the case of route reflection, if a route reflector R_r fails, not only does R_r itself lose reachability learned from its neighbors; the client routers that used R_r to communicate with other routers would no longer be able to send or receive routing updates. To avoid such single points of failures, RRs are normally deployed in pairs, and each client router is usually connected to two or more RRs.

Prolonged Routing Convergence — An AS with route reflection can experience longer routing convergence compared to full-mesh i-BGP interconnections. In the full-mesh i-BGP case, a BGP update travels only one i-BGP hop to reach all other i-BGP routers. However, with route reflection, an update message may traverse more than one RR before reaching the final i-BGP router. Since each RR runs the best path selection process, there are both processing delay and transmission delay to cross a route reflector. These additional delays in update propagation time can lead to a longer overall convergence delay.

Besides the increased delay in routing message propagations, redundant route reflectors also introduce multiple parallel paths to a given destination. For example, in Fig. 1b, R_2 can see up to three paths during the convergence process after a destination announced by R_4 becomes unreachable:

- R_2-R_1-R_4,
- R_2-R_3-R_4
- R_2-R_1-R_3-R_4

Had all the routers been connected in a full mesh, R_2 would have only one path to reach it, and the convergence could be faster.

Data Forwarding Loop — In a simple route reflection configuration where a single RR connects to all client routers, there should be no data plane loops. However in real deployment, because all client routers must connect to multiple RRs to avoid a single point of failure, this redundant connectivity to RRs can potentially introduce subtle data plane loops that defeat intuitive inspection, as we show by the following example borrowed from [4].

When a client router receives a data packet, it looks up the destination address and forwards the packet to the egress next-hop router. Depending on the IGP connectivity, there can be multiple router hops between this client router and the egress next-hop router, as is the case in Fig. 2. In Fig. 2, RR_1 and RR_2 can reach prefix d in AS2, and both announce this reachability to their clients R_1 and R_2. As far as BGP routing is concerned, there is no routing loop. However when R_1 receives a data packet whose destination address is d, it will try to send the packet to the egress next-hop RR_1 via R_2, expecting R_2 to further forward this packet to RR_1. On the other hand, R_2 believes that the egress router for destination d is RR_2 and sends the packet back to R_1, expecting that R_1 will forward the packet to RR_2. As a result of the inconsistencies between the control plane topology and physical connectivity, i.e., R_1 is connected to RR_1 on the control plane but connected to R_2 physically, and vice versa, packets heading to destination d would end up bouncing back and forth between R_1 and R_2.

Reduced Path Diversity — For a given BGP router, path diversity is a measure to quantify the number of different routes available to reach a given destination. High path diversity for each destination prefix can increase the resiliency against failures and offer opportunities for traffic engineering. Since an RR only propagates its best route for a given destination, all the client routers of the given RR use the same single best route to the destination as chosen by the RR. Figure 3 shows such an example: although both R_1 and R_2 are directly connected to AS2 to reach destination prefix d, if the reflector RR chooses R_1 as the best path to d, R_3 has to use that path as well. Furthermore, when the link between R_1 and R_4 fails, R_3 will have to wait for some time until RR learns about the failure and switches to an alternative path to d, and then propagates the new path to all its clients. In contrast, full-mesh i-BGP interconnections not only allow R_1 and R_2 to use their direct connection to AS2 to reach prefix d, but also allow R_3 to learn both paths and choose between them, and be able to switch to the other path as soon as it learns about the failure from R_1 directly. Recently, a number of measurement studies addressing the amount of path diversity in a given AS [5] and the impact of architectural impact of more scalable i-BGP architectures such as route reflection [6] have appeared in the literature, and interested readers are further referred to these studies.

There have been several recent efforts to increase the path diversity in i-BGP to reduce the convergence time. Reference [7] by Raszuk *et al.* suggests increasing path diversity within an AS by modifying the best path selection in RRs so that different RRs will advertise different paths to client routers. Another proposal is adding a *best external* option [8] in BGP. By using the best external option, a border BGP router can propagate more than one best external path to i-BGP neighbors inside an AS. This can increase the number of paths observed by i-BGP routers and decrease the number of hidden paths. Yet another proposal by Walton *et al.* [9] suggests allowing any BGP router to propagate more than a single best path to increase the overall path diversity.

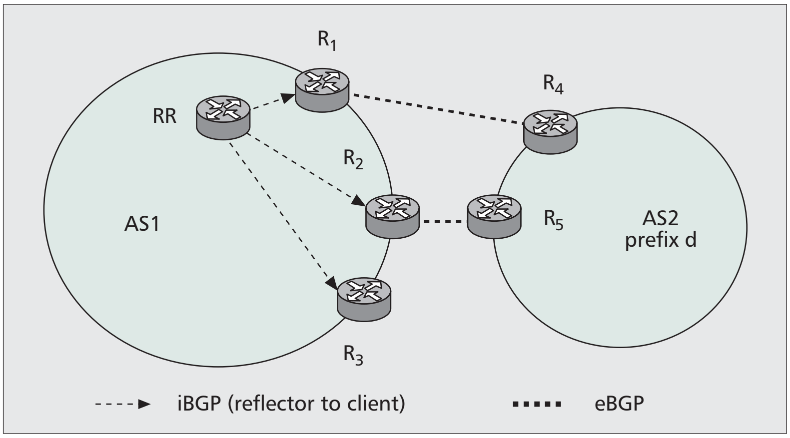

Figure 3. *RR chooses* its *best route.*

Sub-Optimal Routes — An RR selects its best paths to reach the destination prefixes using its local routing information, and propagates these selected paths to its clients. It is most likely that not all the best paths chosen by the reflector would be the best paths for *each* of all its clients. Therefore, some client routers end up using suboptimal paths to some destinations. For example, in Fig. 3, AS1 has two paths to reach prefix d in AS2, R_1-R_4 and R_2-R_5. Assuming that the link lengths in Fig. 3 reflect the IGP distances of the routers, the route reflector RR would pass to R_1, R_2, and R_3 its own best path to prefix d in AS2, which is through R_1-R_4 (because RR itself is closer to R_1 than R_2). R_2 will still use its own best path through R_2-R_5 because of the BGP best path selection rule that prefers the path learned from e-BGP over that learned from i-BGP. However, R_3 will use the path R_1-R_4, the only path learned from the RR. R_3's shortest path to prefix d should have been through R_2-R_5, had AS1 used full-mesh i-BGP interconnections.

In the next section, we explain how one can address some of the negative side effects by following the guidelines in [1].

Circumventing the Drawbacks through RR Placement

In a network with route reflection, a client router can connect to any RR in the same network. However, as discussed earlier, improperly configured client-reflector relations may lead to subop-

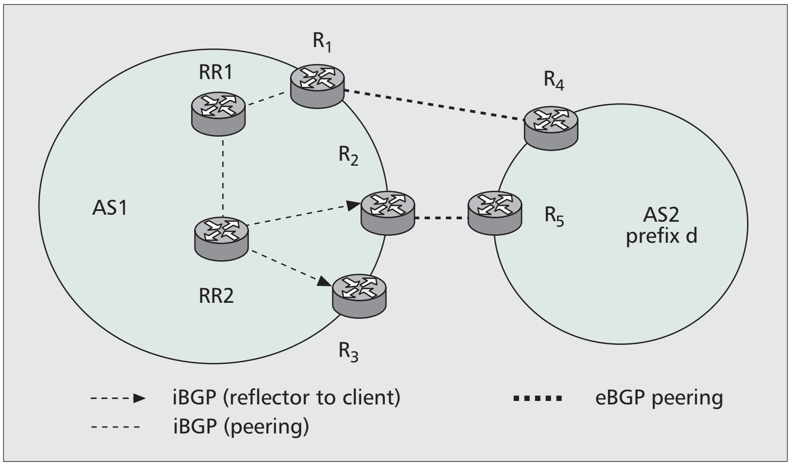

Figure 4. *POP based route reflection.*

timal routing paths. In practice, a pair of RRs is placed in each of its major points of presence (POPs, i.e., offices that the given ISP peers with its peers or customers), so client routers connect to the RRs residing in the same POP, making the logical i-BGP topology following the underlying geographic locations to mitigate the negative impacts of route reflection.

Given that an RR is located in the same POP with its clients, its best path selections should be the same as those made by its clients, at least at the granularity of the POP level. Thus, some of the negative impacts from deploying route reflection mentioned in earlier, such as reduced path diversity and suboptimal routing, should no longer exist at the POP level. For example, the sub-optimal route problem illustrated in Fig. 3 can be avoided by placing an RR in each POP. As shown in Fig. 4, if RR1 is placed in the same POP with R_1, and RR_2 in the same POP with R_2 and R_3, both R_2 and R_3 can use the path R_2-R_5 to reach prefix d.

We make an observation that there is a tradeoff between the number of deployed route reflectors and routing optimality. Although technically a route reflector can maintain thousands of sessions (and therefore clients), connecting so many client routers across multiple POPs may lead to less optimal routing decisions. On the other hand, placing RRs at every POP introduced its own scalability concerns.

Large Internet service providers (ISPs) have routers at a large number of POPs, which may be located in different continents. Route reflection requires that all RRs be connected in a full mesh, putting a pair of RRs in every POP brings back the initial problem of managing full-mesh i-BGP sessions among a large number of RRs in a global scale. An ISP can circumvent the above issue by building a *hierarchy* of RRs.

Hierarchical Route Reflection

A hierarchical route reflection structure can be built by recursive application of route reflection. Since route reflection is an effective means to move i-BGP sessions away from full mesh, one can apply the same idea again at the RR level; that is, for a set of M POP level RRs that requires

$$\frac{M \times (M-1)}{2}$$

full-mesh i-BGP connections, one can simply set up an RR S to connect up the M RRs as its clients. As we already learned, for the overall routing system performance, this RR S should be placed as geographically close to all its clients as possible. However, since the RRs are located at different POPs, no single location can satisfy this requirement. This problem can be alleviated to a large degree through the deployment of multiple levels of route reflections. For example, although there is no location that is close to the POP level RRs on both east and west coasts of the United States, one could have two higher-level RRs, one on the east coast and one on the west coast, that are closer to the POP-level RRs. To ensure the propagation of global BGP routing reachability to all i-BGP routers, one only needs to create full-mesh i-BGP connections among all the top-level RRs. Although hierarchical RR further reduces the total number of sessions, this benefit does not come for free, as we explain next.

Impacts of Hierarchical Route Reflection

Increased Hop Distance and Paths — Under full-mesh i-BGP, any i-BGP speaker can reach any other i-BGP speaker with one i-BGP hop. Under a hierarchical route reflection, the distance for an update to travel from one i-BGP speaker to another is at least two hops (client-reflector-client), and in many cases longer. In addition to increased numbers of i-BGP hops, this hierarchical route reflection also leads to increased number of alternative paths that updates may travel through.

Additional Path Diversity Reduction — Multi-level hierarchical route reflection topology can also further reduce path diversity, because the total number of routes to a destination d is limited by the total number of the RRs at the highest level that d's reachability is propagated. As one approaches the top of the hierarchy, the number of RRs reduces.

Summary and Future Work

In this article, we described the route reflection solution along with its advantages and disadvantages that have been identified over time. In the past, the number of BGP sessions that a router can handle was relatively small. Thanks to software and hardware technology advances, today's routers on the market are capable of handling thousands of i-BGP sessions [7], removing one of the reasons for route reflection deployment. However the operational cost from configuring and maintaining full-mesh i-BGP sessions remains a strong motivation for deploying route reflections in a large network. Our study suggests that several open issues remain, and several potentials also exist, to make route reflection an effective solution toward future routing scalability. We identify the following items for future work.

REMAINING ISSUES WITH ROUTE REFLECTION

We sort the route reflection induced side effects identified in Section II-D into two categories. The first one concerns routing convergence. Route reflection deployment in a global-scale ISP desires a hierarchical structure, which can prolong routing propagation and worsen routing convergence. The second category concerns mitigating path diversity reduction. Route reflection may reduce the overall path diversity. However, this route reflection induced reduction can be mitigated by a well engineered RR placement as mentioned in [6]. Another interesting approach is to utilize redundant RRs that can address robustness, path diversity, and suboptimal paths all at once as discussed in [7, 9]. By design, an RR plays a more important role than a client router, thus it requires redundancy against a single point of failure. Redundant RRs can then be utilized to increase path diversity and reduce suboptimal routing.

ROUTE REFLECTORS AND ROUTE SERVERS

In conducting this research we also observe that a number of similarities exist between a route reflector in i-BGP and a route server which is used to distribute reachability information in large exchange points in e-BGP. We believe that a clear articulation of similarities and differences between the two can further improve our understanding of route reflection and how to use it most effectively to distribute routing information in i-BGP context.

SEPARATING CONTROL PLANE FROM DATA PLANE

Lastly, as the Internet continues to grow in size, ISPs also grow rapidly over time and its overall topology becomes more complex. A recent trend in scaling and simplifying network management is to decouple a network's control plane from its data plane. We observe from the operational practice that route reflection can be used as a simple, incrementally deployable means to steer a network toward separating its control plane from the data plane.

REFERENCES

[1] T. Bates, E. Chen, and R. Chandra, "RFC 4456: BGP Route Reflection: An Alternative to Full Mesh Internal BGP," Apr. 2006.

[2] P. Traina, D. McPherson, and J. Scudder, "RFC 5065: Autonomous System Confederations for BGP," Aug. 2007.

[3] S. Halabi and D. McPherson, "Internet Routing Architectures, 2nd Ed.," Aug. 2000.

[4] T. G. Griffin and G. Wilfong, "On the Correctness of iBGP Configuration," *SIGCOMM Comp. Commun. Rev.*, vol. 32, no. 4, 2002, pp. 17–29.

[5] J. Choi *et al.*, "Understanding BGP Next-hop Diversity," *Proc. 14th IEEE Global Internet Symp.*, Shanghai, China, Apr. 2011, pp. 846–51.

[6] J. H. Park *et al.*, "A Comparative Study of Architectural Impact on BGP Next-hop Diversity," *Proc. 15th IEEE Global Internet Symp.*, Orlando, FL, Mar. 2012.

[7] R. Raszuk *et al.*, "Distribution of Diverse BGP Paths," http://tools.ietf.org/html/draft-ietfgrow-diverse-bgp-path-dist-07, Nov. 2011.

[8] P. Marques *et al.*, "Advertisement of the Best External Route in BGP." http://tools.ietf.org/html/draft-ietf-idr-best-external-05, Jan. 2012.

[9] A. R. D. Walton *et al.*, "Advertisement of Multiple Paths in BGP," http://tools.ietf.org/html/draft-ietf-idr-add-paths-06, Sept. 2011.

BIOGRAPHIES

JONG HAN PARK (jpark@research.att.com) received his B.S. and M.S. in computer science and networked systems from the University of California, San Diego (2000) and the University of California, Irvine (2005). Recently he finished his Ph.D. in computer science from the University of California at Los Angeles (UCLA, 2011) and joined AT&T Labs, Middletown, New Jersey. Currently at AT&T Labs, he works on various topics including BGP and network traffic data monitoring, performance analysis, and CDN architecture design.

RICARDO OLIVEIRA (ricardo@thousandeyes.com) received his B.S. in electrical engineering from the Engineering Faculty of Porto University, Portugal (2001), and his M.S. and Ph.D. in computer science from UCLA (2005, 2009) where he was the recipient of the Cisco Outstanding Graduate Student Research Award and the Edward K. Rice Outstanding Doctoral Student Award. During his Ph.D. he did several internships, including AT&T Labs, Thomson, and Juniper Networks in Silicon Valley, and published dozens of papers in international conferences and journals, including ACM SIGCOMM, ACM SIGMETRICS, and *IEEE Transactions on Networking*.

SHANE AMANTE (Shane.Amante@level3.com) is currently senior IP architect at Level 3 Communications, Inc. where he is responsible for backbone network architecture. Previously, he worked on edge network architecture and was responsible for the design and deployment of Level 3's HSIP, IPVPN, EVPL, VPLS, and Switched MetroEthernet services. He is an active participant in the IETF, is co-author of several drafts and RFCs, and was formerly the co-chair of the L2VPN WG. Prior to Level 3, he worked as a network architect at Ricochet, OneSecure, Qwest Communications, and Icon CMT performing a variety of roles, including network design, deployment, and automation of IP and managed security services.

DANNY MCPHERSON (dmcpherson@verisign.com) is chief security officer for Verisign where he is responsible for strategic direction, research and innovation in infrastructure, and information security. Prior to joining Verisign, he was vice president and chief security officer at Arbor Networks. He is a member of the Internet Architecture Board (IAB) and Internet Research Steering Group (IRSG), and co-chairs the IETF's L3VPN WG. He also serves on the ICANN Security and Stability Advisory Council (SSAC) and the FCC's Network Reliability and Interoperability Council (NRIC).

LIXIA ZHANG [F] (lixia@cs.ucla.edu) is a professor in the Computer Science Department of UCLA. She received her Ph.D. in computer science from Massachusetts Institute of Technology and was a member of the research staff at Xerox PARC before joining UCLA. In the past she served as vice chair of ACM SIGCOMM, a member of the editorial board for *IEEE/ACM Transactions on Networking*, a member of the Internet Architecture Board, and co-chair of the Routing Research Group under IRTF. She is a fellow of ACM, the recipient of the 2009 IEEE Internet Award, and the holder of the UCLA Postel Chair in Computer Science.

The operational cost from configuring and maintaining full-mesh i-BGP sessions remains a strong motivation for deploying route reflections in a large network. Our study suggests that several open issues remain, and several potentials also exist, to make route reflection an effective solution toward future routing scalability.