

Incrementally-Deployable Security for Interdomain Routing

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

Joan Feigenbaum
Yale University
joan.feigenbaum@yale.edu

Abstract

The Internet’s interdomain-routing system is extremely vulnerable to accidental failure, configuration errors, and malicious attack. Any successful approach to improving interdomain-routing security must satisfy two requirements for incremental deployability: backwards compatibility with the existing routing protocol and installed base of routers and incentive compatibility with the desire of each domain to improve its part of the routing system even if other domains have not taken similar steps. We propose an incrementally deployable approach based on a Routing Control Platform (RCP) that makes routing decisions on behalf of the routers in a domain, without requiring changes to the routers or protocols. The RCP runs anomaly-detection algorithms that identify, and avoid, suspicious routes, allowing a domain (or a small group of cooperating domains) to significantly improve interdomain routing security.

1 Introduction

The Internet consists of tens of thousands of separately-administered networks, called Autonomous Systems (ASes), that exchange routing information using the Border Gateway Protocol (BGP). In essence, BGP is the glue that holds the disparate parts of the Internet together, making the correct and stable operation of the protocol of paramount importance. Yet, an AS can easily introduce false information into BGP about how to reach destination addresses, leading other networks to send traffic in the wrong direction. Such BGP “hijacking” is an effective way for an attacker to snoop on traffic en route to a legitimate destination, impersonate a Web site (e.g., to perform identity theft on unsuspecting Web clients), block access to certain sites, or anonymously send spam. Several high-profile “BGP hijacking” incidents have taken place over the years [3, 4, 11, 12], including the recent hijacking of the popular YouTube site by Pakistan Telecom [12].

Despite the many serious problems with BGP, no viable alternative has achieved any significant deployment. Pro-

posals for a secure interdomain-routing protocol have been stymied, at least in part, by the inability to have a “flag day” on which routers throughout the Internet upgrade to the new protocol. We believe that any successful approach for improving the security of the interdomain-routing system must be both *backwards compatible* with the BGP protocol and the installed base of routers and *incentive compatible* with the desire of each AS to improve its part of the routing system even if other ASes have not yet taken similar steps. Failure to satisfy these two requirements can severely hamper the deployment of a new security technique, however effective it would be if widespread adoption were achieved.

For example, Secure BGP (S-BGP) [10] uses digitally signed statements to verify the authenticity of address allocations and route announcements. However, S-BGP has not been widely deployed, because the protocol (i) relies on the deployment of an Internet-wide public-key infrastructure, (ii) requires fundamental changes to BGP and the routers themselves, (iii) does not provide security advantages until all domains in a path are using the new protocol, and (iv) does not address the need for each AS to specify policies and verify that they are followed consistently. In our work, we advocate an incrementally deployable approach to interdomain-routing security, with the end goal of ubiquitous deployment of a secure routing protocol. Our work has the following two main ingredients:

Routing Control Platform (RCP): The RCP provides a way for a domain to move the control for making BGP-routing decisions out of the individual routers and into a small set of servers. In addition to simplifying network management, the RCP enables a domain to apply enhanced security policies or even upgrade to a secure routing protocol, while continuing to use the traditional BGP protocol to interact with legacy routers. Participating ASes have an incentive to exchange routing information directly via their RCPs in order to use the improved protocols.

Algorithms for identifying and avoiding suspicious routes: The RCP provides a logical place to run data-analysis algorithms that identify suspicious routes by analyzing streams of BGP messages. The RCP can also intentionally avoid selecting these suspicious routes, and in-




Figure 1. RCP deployment in a single AS

stead assign “normal-looking” routes to the routers. When multiple ASes deploy RCPs, they can cooperate in a distributed fashion to identify suspicious routes advertised by other parts of the Internet, or actively “hijack the hijacker” to draw data traffic toward the legitimate destinations.

In this paper, we briefly describe the design, implementation, and evaluation of our RCP system, called Morpheus [16], and present algorithms that participating ASes can run to substantially improve BGP security [9, 19, 18, 2, 1, 17].

2 Routing Control Platform (RCP)

Rather than modifying the installed base of routers, an AS can apply new secure-routing policies at a small collection of servers that select BGP routes on behalf of the routers. In this section, we first describe how a Routing Control Platform (RCP) [6] operates without requiring any changes to the routers or any support from neighboring ASes. Then, we present the design of our prototype system [16] that enables network administrators to compose AS-level policy objectives into a single, coherent policy.

2.1 Incremental Deployment of the RCP

Despite the complexity of BGP’s policy-configuration and path-selection logic, the protocol itself is relatively simple. Two routers communicate via BGP by establishing a session over a Transmission Control Protocol (TCP) connection. The two routers use the BGP session to exchange update messages that either advertise a new route for an address block (or *destination prefix*) or withdraw an old one. We propose retaining the protocol’s simple state machine and message format, while radically changing the decision logic and policy configuration. We follow the same basic approach that has allowed Ethernet technology to change substantially in the past fifteen years while retaining (or perhaps *because of* retaining) the framing format.

In particular, the RCP communicates with the individual routers via BGP sessions, as shown in Figure 1. The RCP forms internal BGP (iBGP) sessions with the routers to learn routing information about external destinations and to send each router a *customized* routing decision for each destination prefix. In contrast to a routing registry that stores information about prefix ownership or routing policies, the RCP stores the BGP-update messages themselves and selects the routes in real time on behalf of the routers in the AS. To the rest of the network, the RCP looks just like a router that sends and receives BGP messages in the standard format, although internally the RCP may incorporate new kinds of decision logic for selecting paths. The RCP also monitors the Interior Gateway Protocol (IGP), such as Open Shortest Path First (OSPF) or Intermediate System-Intermediate System (IS-IS), running inside the network to react in real time to changes in the topology.

To be viable in practice, the RCP must be fast and reliable enough to make BGP-routing decisions for the many routers in a large IP network. Fortunately, these issues are tractable in practice, for the following reasons. First, a modern commodity computer has ample memory and processing resources to store and compute routes for a large AS. Maintaining tens of thousands of BGP sessions and storing millions of BGP routes is quite doable with today’s PCs. Second, the RCP can amortize storage and computation overhead in selecting routes for the many routers in the AS. Most of the BGP decision process is the same for all routers in the AS, and the RCP need only store a single copy of each learned route. Third, simple replication of the RCP functionality on multiple computers and placement of these replicas in diverse locations are sufficient to achieve high reliability. Fourth, the RCP does not need to carry data packets or communicate with end hosts in any way, making it possible to erect a strong security perimeter around the server replicas. These systems issues have been explored in depth in earlier RCP prototyping efforts [5, 15].

Although building and deploying an RCP is *feasible*, an AS must have a sufficient incentive to deploy the RCP replicas and configure BGP sessions to the legacy routers. We believe that the initial incentive lies in the prospect of avoiding the substantial complexity of configuring the routers and operating the network. Managing the configuration of the many routers in a large network currently requires constant vigilance on the part of operators. Although network-management systems can often automate the most frequent tasks, these systems are essentially scripts that “robot-ize” the process of typing at the command-line interface of the routers. Having to work around and within the constraints of the existing routing protocols makes network management much more complex and error-prone than necessary. In addition, the legacy routers are quite difficult to change, requiring long interactions with vendors and standards bod-

ies. In contrast, with an RCP, an individual AS can easily deploy a wide range of new, customer-facing (and revenue-generating) services [14].

2.2 Design and Evaluation of Morpheus

The RCP offers an AS substantial flexibility in selecting BGP routes for the underlying routers. However, to make the RCP approach viable in practice, network administrators must have a working system that is easy to configure and extend. Early RCP prototyping efforts either mimicked the existing BGP decision process [5] or supported each new routing policy largely from scratch [14]. In contrast, our goal is to create a system architecture that allows administrators to specify simple policy objectives and compose them in a flexible way. For example, administrators often have policy objectives for routing security, business relationships with neighboring ASes, traffic engineering, stability, and scalability. Today, these objectives are intertwined in complex ways with the BGP policy configuration on individual routers. We want the administrator to be able to specify each objective separately, and also specify the relative importance of the objective in making route-selection decisions. In addition, we want to enable a marketplace for third-party software that implements new modules that realize individual policy objectives.

To support flexible composition of policy objectives, we consider the following route-selection problem an AS faces: Given a set of available routes $R = \{r_1, r_2, \dots, r_n\}$ for a destination prefix p , choose a best route r^* for each router according to a set of criteria $C = \{c_1, c_2, \dots, c_k\}$. The set of criteria (i.e., policy objectives) includes route characteristics such as security, performance, and stability. Each policy objective is implemented as a classifier module that, given a BGP-learned route, computes a tag that is a sort of “score” for the route according to this particular objective. A classifier may base its decision on the contents of the route announcement, as well as external information (such as the business relationship with each neighboring AS or a registry of prefix ownership) and internal state (such as the past history of ASes that originated BGP routes for each prefix). Each classifier operates independently, and can be implemented independently.

Unfortunately, the many policy objectives may be conflicting in the sense that no BGP route is best with respect to all criteria. Therefore, the route-selection process must ensure that the network administrator has the flexibility to make arbitrary trade-offs among the criteria. The decision process computes a cumulative score as a weighted sum of the tags for each route and picks the route with the highest total score as the best route. For enhanced flexibility, multiple decision processes with different settings of the weights can be run in parallel. This allows an AS to offer differ-

ent services to different classes of customers. For example, some customers may put a premium on security, even at the expense of performance, whereas others may prioritize performance and stability over other objectives.

Our solution to the route-selection problem—combining the route classifiers and the weighted decision processes—is a system we call Morpheus [16] as it gives ISPs the power to “shape” their routing policies. Our prototype is implemented as an extension to the XORP [7] software router. We extended XORP by implementing (i) the weighted-sum decision process(es) as a replacement for the conventional BGP decision process, (ii) support for computing different route assignments for different routers, and (iii) several example classifiers based on security and business objectives. Our experiments on a 3.2GHz Intel Pentium-4 PC, running Linux 2.6.11, demonstrate that (i) the classifiers run in 20–100 microseconds, depending on the complexity of the policy objective and (ii) the decision process requires about 50 microseconds to select a best route. These performance results are sufficient to support a large AS within hundreds of routers and thousands of BGP sessions.

Classifiers and weights offer a much more intuitive way for network administrators to specify and compose their policy objectives. Still, humans are not good at setting a large number of weights directly to reflect their preferences. Instead, studies show that humans do a much better job in expressing their preferences through pairwise comparisons between alternatives, even though the results of these comparisons are often inconsistent [13]. Based on this observation, Morpheus leverages Analytic Hierarchy Process (AHP) [13], a technique in decision theory, to provide a simple, intuitive configuration interface. Network operators specify their policy preferences through pair-wise comparisons of example routes, and AHP automatically derives the appropriate weights to use when making online decisions. AHP also provides an appealing way for an AS to give its customers more influence over route selection, as a value-added service. In particular, the AS may allow its customers to specify preferences for certain policy objectives (e.g., related to the customers’ own performance and security goals) while retaining control over other objectives (e.g., related to the AS’s business relationships with neighboring domains).

3 Algorithms that Improve BGP Security

Morpheus is a natural platform for running new algorithms that identify and avoid suspicious routes. Over the past few years, we have designed and evaluated several anomaly-detection algorithms that individual ASes can run to improve their own security [9, 19] as well as new techniques for small groups of ASes to collaborate to achieve additional security gains [17, 2, 1]. In this section, we briefly describe two of these techniques [9, 2].

3.1 Avoiding Suspicious Routes

A single AS, acting alone, cannot rely on secure extensions to BGP (such as S-BGP) to detect or prevent bogus route announcements. However, an AS can run anomaly-detection algorithms that, over time, learn how to distinguish “normal” BGP routes from “suspicious” ones. Such an algorithm can run as a classifier in Morpheus to ensure that normal routes are (strongly) favored over suspicious ones. Our solution, called Pretty Good BGP (PG-BGP) [9], maintains a history of (destination prefix, originating AS) pairs and an inferred AS-level topology, to detect prefix hijacking and invalid paths, respectively. Upon receiving a new BGP route announcement, PG-BGP assigns a low “score” if the route disagrees with past history. Fortunately, a large AS typically learns multiple routes for the same destination prefix (e.g., from different neighboring ASes), allowing the AS to select a normal route over a suspicious one. In the meantime, the arrival of a suspicious route can trigger an alarm for the network administrator to investigate. This separation of timescale—real-time avoidance of suspicious routes coupled with human-timescale investigation of the anomalies—enables the AS to have an automated response that prevents attacks from propagating.

Our experiments evaluating PG-BGP with real BGP measurements show very positive results. First, analysis of traces of BGP update messages confirms our intuition that most suspicious routes disappear relatively quickly, within an hour or at most a day. As such, PG-BGP’s automated response is often sufficient to protect against these kinds of attacks without further action by the network administrator. Second, simulations of PG-BGP on the Internet’s AS-level topology show that modest deployments are surprisingly effective, not only at protecting the participating ASes but *even at protecting the rest of the Internet*. Our experiments show that, if the largest 60 ASes (mostly tier-1 and tier-2 providers) deployed PG-BGP, an attacker can only convince 2.5% of the ASes in the Internet to direct traffic along a bogus route. These gains are possible because (i) the large ASes have substantial path diversity, allowing them to select a valid route rather than the suspicious one and (ii) most of the other ASes in the Internet select a route announced to them (either directly or indirectly) through one of these large providers. These large providers are the most technically sophisticated, typically running the most recent router software and security patches, making it more likely that they would deploy security enhancements like PG-BGP.

Though PG-BGP is capable of mitigating short-term attacks and misconfigurations autonomously, an additional mechanism is needed for longer attacks. The Internet Alert Registry (IAR) [8] is a distributed notification system that can inform an AS when its prefixes are hijacked, or when it is (perhaps inadvertently) hijacking other prefixes. Once

informed, the networks can confirm the validity of a suspicious route and take action to fix it if necessary. Our IAR system is an opt-in service that runs the PG-BGP algorithm on public feeds of BGP updates from RIPE and RouteViews, and distributes e-mail alerts to the affected ASes. Numerous network administrators have subscribed to the IAR, and others often browse our IAR Web site to read the reports posted there.

3.2 Collaboration in Small Groups

PG-BGP is very effective at enabling a large participating AS to select valid routes to destinations throughout the Internet. However, if no large providers run PG-BGP, a small AS (e.g., a stub network such as a corporate or university campus) with limited route diversity does not necessarily learn *any* valid route to a destination during an attack. In addition, an AS may want to ensure that senders throughout the Internet can successfully reach its destination prefixes, without requiring (say) 60 large providers to deploy a solution like PG-BGP. Ideally, we would like a way for an even smaller group of ASes—say, 5-10 ASes—to be effective in protecting their own destination prefixes from attack. This problem is challenging because of the large number (tens of thousands) of *non-participants* that unknowingly select and propagate bogus routes originating by attackers.

To be effective the participating ASes must (i) expose additional path diversity, to ensure that they have valid routes to the destination and (ii) be proactive in coaxing non-participants into selecting valid routes. Our solution consists of two main mechanisms. First, the participating ASes form a secure overlay network we call an SBone. In contrast to conventional overlays, an SBone connects *networks* rather than end hosts, collects path-quality measurements that are robust to adversaries, and avoids mapping virtual links on to compromised paths through the Internet. Second, all participating ASes originate BGP announcements for the prefixes the group wants to protect, and then forward the traffic over the secure overlay to the legitimate destination. “Shouting” the group’s prefixes—essentially “hijacking the hijacker”—substantially improves availability, in exchange for a small increase in routing-table size and path lengths. To limit the overhead, the group members can shout *reactively* after detecting an attack using an anomaly-detection scheme like PG-BGP and the IAR.

Our experiments, evaluating both SBone and Shout on a snapshot of the Internet’s AS-level topology, suggest that the two mechanisms are very effective, especially when a few large ISPs participate in the group. For example, if a group of 5-10 ASes includes at least three large ISPs, members of the group can communicate over valid paths through the SBone under 95% of attacks by a single adversary. A group with five large ISPs can even defend against attacks

by multiple adversaries. Shouting allows the group to attract traffic sent by non-participating ASes even during an attack. If the group has ten members (including three or more large ISPs), 95% of *all ASes in the Internet* are able to reach the group's destination prefixes during an attack. That is, the vast majority of *non-participating ASes* are successfully coaxed into picking a valid route to the destination. Understandably, sometimes the path is longer (in terms of AS-level hops), but for the ten-member group (including three large ISPs), paths are just 15% longer on average—a small price to pay to avoid a compromised path.

4 Conclusions

The Internet's interdomain routing system is critical infrastructure underlying much of the world's communication. Protecting the routing system from malicious attacks and accidental misconfigurations is of paramount importance. Yet, wholesale deployment of a secure version of BGP is immensely difficult in practice. Instead, we advocate an incrementally-deployable approach to improving interdomain routing security. Our solutions are backwards compatible with legacy routers and allow each AS to decide whether to participate. The Routing Control Platform (RCP) allows an AS to apply flexible routing policies, including new secure protocols and anomaly-detection schemes, to select routes on behalf of the routers. Our PG-BGP anomaly detector, implemented as a classifier module in our Morpheus RCP prototype, allows an AS to detect and avoid suspicious routes. Our SBone and Shout solution allows a small group of 5-10 ASes to ensure that non-participating ASes can successfully deliver traffic to the members of the group. Together, these solutions offer substantial incentives for early adopters, making them an important part of an incrementally-deployable approach to improving the security of the interdomain routing system.

Acknowledgments

The authors gratefully acknowledge the support of the U.S Department of Homeland Security Science and Technology Directorate under Grant W911NF-05-1-0417.

References

- [1] I. Avramopoulos and J. Rexford. A pluralist approach to interdomain communication security. In *Proc. NetEcon Workshop*, June 2007.
- [2] I. Avramopoulos, M. Suchara, and J. Rexford. How small groups can secure interdomain routing. Technical Report TR-808-07, Princeton University Computer Science Department, December 2007.
- [3] R. Barrett, S. Haar, and R. Whitestone. Routing snafu causes Internet outage. *Interactive Week*, April 25 1997.
- [4] P. Boothe, J. Hiebert, and R. Bush. How prevalent is prefix hijacking on the Internet?, February 2006. NANOG 36, <http://www.nanog.org/mtg-0602/boothe.html>.
- [5] 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. Networked Systems Design and Implementation*, May 2005.
- [6] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and J. van der Merwe. The case for separating routing from routers. In *Proc. ACM SIGCOMM Workshop on Future Directions in Network Architecture*, August 2004.
- [7] M. Handley, E. Kohler, A. Ghosh, O. Hodson, and P. Radoslavov. Designing extensible IP router software. In *Proc. Networked Systems Design and Implementation*, May 2005.
- [8] Internet Alert Registry. <http://iar.cs.unm.edu/>.
- [9] J. Karlin, S. Forrest, and J. Rexford. Autonomous security for autonomous systems. *Computer Networks, Special issue on Complex Computer and Communications Networks*, October 2008.
- [10] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (S-BGP). *IEEE J. Selected Areas in Communications*, 18(4), April 2000.
- [11] Rensys Blog. Con-Ed steals the 'Net. http://www.renesity.com/blog/2006/01/coned_steals_the_net.shtml.
- [12] Rensys Blog. Pakistan hijacks YouTube. http://www.renesity.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml.
- [13] T. L. Saaty. *Models, Methods, Concepts and Applications of the Analytic Hierarchy Process*. Kluwer Academic Publishers, 2001.
- [14] J. van der Merwe et al. Dynamic connectivity management with an intelligent route service control point. In *Proc. ACM SIGCOMM Workshop on Internet Network Management*, September 2006.
- [15] P. Verkaik, D. Pei, T. Scholl, A. Shaikh, A. Snoeren, and J. van der Merwe. Wrestling control from BGP: Scalable fine-grained route control. In *USENIX Annual Technical Conference*, June 2007.
- [16] Y. Wang, I. Avramopoulos, and J. Rexford. Design for configurability: Rethinking interdomain routing policies from the ground up. *IEEE J. Selected Areas in Communications*. To appear.
- [17] H. Yu, J. Rexford, and E. Felten. A distributed reputation approach to cooperative Internet routing protection. In *Proc. Workshop on Secure Network Protocols*, November 2005.
- [18] J. Zhang and J. Feigenbaum. Finding highly correlated pairs efficiently with powerful pruning. In *Proc. Conference on Knowledge and Information Management*, pages 152–161, 2006.
- [19] J. Zhang, J. Rexford, and J. Feigenbaum. Learning-based anomaly detection in BGP updates. In *Proc. ACM SIGCOMM MineNet Workshop*, August 2008.