Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security

Phillipa Gill
Stony Brook University

Michael Schapira
Hebrew University at Jerusalem

Sharon Goldberg
Boston University

August 27, 2013

Abstract

With a cryptographic root-of-trust for Internet routing (RPKI [9]) on the horizon, we can finally start planning the deployment of one of the secure interdomain routing protocols proposed over a decade ago (Secure BGP [13], secure origin BGP [26]). However, if experience with IPv6 is any indicator, this will be no easy task. Security concerns alone seem unlikely to provide sufficient local incentive to drive the deployment process forward. Worse yet, the security benefits provided by the S*BGP protocols do not even kick in until a large number of ASes have deployed them.

Instead, we appeal to ISPs' interest in increasing revenue-generating traffic. We propose a strategy that governments and industry groups can use to harness ISPs' local business objectives and drive global S*BGP deployment. We evaluate our deployment strategy using theoretical analysis and large-scale simulations on empirical data. Our results give evidence that the market dynamics created by our proposal can transition the majority of the Internet to S*BGP.

Keywords. BGP security, interdomain routing, incentives, technology diffusion, simulations.

Bibliographic note. This work appeared at SIGCOMM'11.

1 Introduction

Th Border Gateway Protocol (BGP), which sets up routes from autonomous systems (ASes) to destinations on the Internet, is amazingly vulnerable to attack [5]. Every few years, a new failure makes the news; ranging from misconfigurations that cause an AS to become unreachable [20, 22], to possible attempts at traffic interception [8]. To remedy this, a number of widely-used stop-gap measures have been developed to *detect* attacks [11,17]. The next step is to harden the system to a point where attacks can be *prevented*. After many years of effort, we are finally seeing the initial deployment of the Resource Public Key Infrastructure (RPKI) [1,3,19,23], a cryptographic root-of-trust for Internet routing that authoritatively maps ASes to their IP prefixes and public keys. With RPKI on the horizon, we can now realistically consider deploying the S*BGP protocols, proposed a decade ago, to prevent routing failures by validating AS-level paths: Secure BGP (S-BGP) [13] and Secure Origin BGP (soBGP) [26].

1.1 Economic benefits for S*BGP adoption.

While governments and industry groups may have an interest in S*BGP deployment, ultimately, the Internet lacks a centralized authority that can mandate the deployment of a new secure routing protocol. Thus, a key hurdle for the transition to S*BGP stems from the fact that each AS will make deployment decisions according to its own local business objectives.

Lessons from IPv6? Indeed, we have seen this problem before. While IPv6 has been ready for deployment since around 1998, the lack of tangible local incentive for IPv6 deployment means that we are only now

starting to see the seeds of large-scale adoption. Conventional wisdom suggests that S*BGP will suffer from a similar lack of local incentives for deployment. The problem is exacerbated by the fact that an AS cannot validate the correctness of an AS-level path unless all the ASes on the path deployed S*BGP. Thus, the security benefits of S*BGP only apply after a large fraction of ASes have already deployed the protocol.

Economic incentives for adoption. We observe that, unlike IPv6, S*BGP can impact routing of Internet traffic, and that this may be used to drive S*BGP deployment. These crucial observations enable us to avoid the above issues and show that global S*BGP deployment is possible even if local ASes' deployment decisions are *not* motivated by security concerns. To this end, we present a prescriptive strategy for S*BGP deployment that relies solely on Internet Service Providers' (ISPs) local economic incentives to drive global deployment; namely, ISP's interest in attracting revenue-generating traffic to their networks.

Our strategy is prescriptive. We propose guidelines for how (a) ASes should deploy S*BGP in their networks, and (b) governments, industry groups, and other interested parties should invest their resources in order to drive S*BGP deployment forward.

1. Break ties in favor of secure paths. First, we require ASes that deploy S*BGP to actually use it to inform route selection. However, rather than requiring security be the first criterion ASes use to select routes, we only require secure ASes to *break ties* between equally-good routes in favor of secure routes. This way, we create incentives for ISPs to deploy S*BGP so they can transit more revenue-generating customer traffic than their insecure competitors.

*2. Make it easy for stubs to adopt S*BGP.* 85% of ASes in the Internet are *stubs* (*i.e.,* ASes with no customers) [7]. Because stubs earn no revenue from providing Internet service, we argue for driving down their deployment costs by having ISPs sign BGP announcements on their behalf or deploy a simplex (unidirectional) S*BGP [18] on their stub customers. In practice, such a simplex S*BGP must either be extremely lightweight or heavily subsidized.

3. Create market pressure via early adopters. We propose that governments and industry groups concentrate their regulatory efforts, or financial incentives, on convincing a small set of *early adopters* to deploy S*BGP. We show that this set of early adopters can create sufficient market pressure to convince a large fraction of ASes to follow suit.

1.2 Evaluation: Model and simulations.

To evaluate our proposal, we needed a model of the S*BGP deployment process.

Inspiration from social networks? At first glance, it seems that the literature on technology adoption in social networks would be applicable here [10, 12, 21, 24, 25, 27]. However, in social networks models, an entity's decision to adopt a technology depends only on its immediate *neighbors* in the graph; in our setting, this depends on the number of secure *paths*. This complication means that many elegant results from this literature have no analogues in our setting.

Our model. In contrast to earlier work that assumes that ASes deploy S*BGP because they are concerned about security [4, 6], our model assumes that ISPs' local deployment decisions are based solely on their interest in increasing customer traffic.

We carefully designed our model to capture a few crucial issues, including the fact that (a) traffic transited by an ISP can include flows from any pair of source and destination ASes, (b) a large fraction of Internet traffic originates in a few large content provider ASes [15,16], and (c) the cost of S*BGP deployment can depend on the size of the ISP's network. The vast array of parameters and empirical data relevant to such a model mean that our analysis is *not* meant to *predict* exactly how the S*BGP deployment process will proceed in practice; instead, our goal was to evaluate the efficacy of our S*BGP deployment strategy.

Theorems, simulations and examples. We explore S*BGP deployment in our model using a combination of theoretical analysis and simulations on empirical AS-level graphs [2,7]. Every example we present comes directly from these simulations. Instead of artificially reducing algorithmic complexity by subsampling [14], we ran our simulations over the full AS graph. Thus, our simulations ran in time $O(N^3)$ with $N = 36K$, and we devoted significant effort to developing parallel algorithms that we ran on a 200-node DryadLINQ cluster [28].

1.3 Key insights and recommendations.

Our evaluation indicates that our strategy for S*BGP deployment can drive a transition to S*BGP. While we cannot predict exactly how S*BGP deployment will progress, a number of important themes emerge:

1. Market pressure can drive deployment. We found that when S*BGP deployment costs are low, the vast majority of ISPs have incentives to deploy S*BGP in order to differentiate themselves from, or keep up with, their competitors. Moreover, our results show this holds even if 96% of routing decisions (across all source-destination AS pairs) are *not* influenced by security concerns.

*2. Simplex S*BGP is crucial.* When deployment costs are high, deployment is primarily driven by simplex S*BGP.

3. Choose a few well-connected early adopters. The set of early adopters cannot be random; it should include well-connected ASes like the Tier 1's and content providers. While we prove that it is NP-hard to even *approximate* the *optimal* set of early adopters, our results show that even 5-10 early adopters suffice when deployment costs are low.

*4. Prepare for incentives to disable S*BGP.* We show that ISPs can have incentives to *disable* S*BGP. Moreover, we prove that there could be deployment oscillations (where ASes endlessly turn S*BGP on and off), and that it is computationally hard to even *determine* whether such oscillations exist.

5. Minimize attacks during partial deployment. Even when S*BGP deployment progressed, there were always some ASes that did not deploy S*BGP. As such, we expect that S*BGP and BGP will coexist in the long term, suggesting that careful engineering is required to ensure that this does not introduce new vulnerabilities into the interdomain routing system.

Bibliographical notes. An extended abstract of this work appeared at SIGCOMM'11, and the full version of this work is available as online at `http://www.cs.bu.edu/~goldbe/papers/SBGPtrans_full.pdf`. Video of a presentation of this work is available at `http://www.nanog.org/meetings/nanog49/presentations/Tuesday/sec_bgp.wmv`.

References

[1] ARIN. ARIN resource certification. `https://www.arin.net/resources/rpki.html`.

[2] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: Mapped? In *IMC*, 2009.

[3] R. Austein, G. Huston, S. Kent, and M. Lepinski. Secure inter-domain routing: Manifests for the resource public key infrastructure. draft-ietf-sidr-rpki-manifests-09.txt, 2010.

[4] I. Avramopoulos, M. Suchara, and J. Rexford. How small groups can secure interdomain routing. Technical report, Princeton University Comp. Sci., 2007.

[5] K. Butler, T. Farley, P. McDaniel, and J. Rexford. A survey of BGP security issues and solutions. *Proceedings of the IEEE*, 2010.

[6] H. Chang, D. Dash, A. Perrig, and H. Zhang. Modeling adoptability of secure BGP protocol. In *SIGCOMM*, 2006.

[7] Y.-J. Chi, R. Oliveira, and L. Zhang. Cyclops: The Internet AS-level observatory. *ACM SIGCOMM CCR*, 2008.

[8] J. Cowie. Rensys blog: China's 18-minute mystery. `http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml`.

[9] IETF. Secure interdomain routing (SIDR) working group. http://datatracker.ietf.org/wg/sidr/charter/.

[10] M. Jackson and L. Yariv. Diffusion on social networks. In *Confrences des journes Louis-Andr Grard-Varet Public Economy Theory Meeting*, 2005.

[11] J. Karlin, S. Forrest, and J. Rexford. Autonomous security for autonomous systems. *Computer Networks*, oct 2008.

[12] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *ACM SIGKDD*, 2003.

[13] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *JSAC*, 2000.

[14] V. Krishnamurthy, M. Faloutsos, M. Chrobak, L. Lao, J.-H. Cui, and A. G. Percus. Sampling large internet topologies for simulation purposes. *Computer Networks (Elsevier)*, 51(15):4284–4302, 2007.

[15] C. Labovitz. Arbor blog: Battle of the hyper giants. `http://asert.arbornetworks.com/2010/04/the-battle-of-the-hyper-giants-part-i-2/`.

[16] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. In *SIGCOMM*, 2010.

[17] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Shang. Phas: Prefix hijack alert system. In *Usenix Security*, 2006.

[18] M. Lepinski and S. Turner. Bgpsec protocol specification, 2011. `http://tools.ietf.org/html/draft-lepinski-bgpsec-overview-00`.

[19] C. D. Marsan. U.S. plots major upgrade to Internet router security. *Network World*, 2009.

[20] S. Misel. "Wow, AS7007!". Merit NANOG Archive, apr 1997. `http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html`.

[21] S. Morris. Contagion. *Review of Economics Studies*, 2003.

[22] Rensys Blog. Pakistan hijacks YouTube. `http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml`.

[23] RIPE. RIPE NCC Resource Certification. `http://www.ripe.net/certification/`.

[24] E. Rogers. *Diffusion of Innovations, 5th Edition*. Free Press, 2003.

[25] T. Valente. *Network Models of the Diffusion of Innovations (Quantitative Methods in Communication Subseries)*. Hampton Press, 1995.

[26] R. White. Deployment considerations for secure origin BGP (soBGP). draft-white-sobgp-bgp-deployment-01.txt, June 2003, expired.

[27] H. P. Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 2001.

[28] Y. Yu, M. Isard, D. Fetterly, M. Budiu, U. Erlingsson, P. K. Gunda, and J. Currey. Dryadlinq: a system for general-purpose distributed data-parallel computing using a high-level language. In *Usenix OSDI*, 2008.