Computer Networks 70 (2014) 260–287

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

How secure are secure interdomain routing protocols?

Sharon Goldberg [a,*], Michael Schapira [b], Pete Hummon [c], Jennifer Rexford [d]

[a] Computer Science, Boston University, Boston, MA 02215, USA
[b] Hebrew University, Jerusalem, Israel
[c] AT&T, NJ, USA
[d] Princeton University, Princeton, NJ, USA

ARTICLE INFO

Article history:
Received 8 September 2013
Received in revised form 8 March 2014
Accepted 12 May 2014
Available online 13 June 2014

Keywords:
Security
Interdomain routing
BGP

ABSTRACT

In response to high-profile Internet outages, BGP security variants have been proposed to prevent the propagation of bogus routing information. The objective of this paper is to inform discussions of which variant should be deployed in the Internet. To do this, we *quantify* the ability of the key protocols (origin authentication, soBGP, S-BGP, and data-plane verification) to limit the impact of *traffic-attraction attacks*; *i.e.*, when an attacker deliberately draws traffic to its own network, in order to drop, tamper, or eavesdrop on packets. Our results and contributions are as follows:

(1) One might expect that an attacker could maximize the volume of traffic it attracts by using the following intuitive strategy: the attacker should announce, to as many of its neighbors as possible, the *shortest* path that is not flagged as bogus by the secure protocol. Through simulations on an empirically-determined AS-level topology, we show that this strategy is surprisingly effective, even when an advanced security solution like S-BGP or data-plane verification is fully deployed.
(2) Next, we show that these results *underestimate* the severity of attacks. In fact, counterintuitive strategies, like announcing longer paths, announcing to fewer neighbors, or triggering BGP loop-detection, can be used to attract even more traffic than the strategy above. We illustrate this using counterintuitive examples. We also demonstrate that these attacks are not merely hypothetical, by searching the empirical AS-level topology and identifying specific ASes that can launch these attacks.
(3) We prove that it is NP hard to find a traffic-attraction attack strategy that attracts the maximum volume of traffic.

Our results suggest that a clever export policy (*i.e.*, where the attacker announces a legitimate path to a carefully chosen set of neighbors) an often attract almost as much traffic as a bogus path announcement. Thus, our work implies that mechanisms that police export policies (*e.g.*, prefix filtering) are crucial, even if more advanced cryptographic solutions like S-BGP are fully deployed.

1. Introduction

The Internet is notoriously vulnerable to *traffic attraction* attacks, where Autonomous Systems (ASes) manipulate BGP to attract traffic to, or through, their networks [3,5,9,10,21,40,44–46]. Attracting extra traffic

* Corresponding author. Tel.: +1 617 353 8919.
E-mail addresses: goldbe@cs.bu.edu (S. Goldberg), schapiram@huji.ac.il (M. Schapira), jrex@cs.princeton.edu (J. Rexford).

http://dx.doi.org/10.1016/j.comnet.2014.05.007
1389-1286/

enables the AS to increase revenue from customers, or drop, tamper, or snoop on packets. While the proposed extensions to BGP prevent many attacks (see [6] for a survey), even these secure protocols are susceptible to a strategic manipulator who deliberately exploits their weaknesses to attract traffic to its network. Given the difficulty of upgrading the Internet to a new secure routing protocol, it is crucial to understand how well these protocols blunt the impact of traffic attraction attacks.

1.1. Quantifying the impact of attacks

We evaluate the four major security extensions that allow ASes to validate paths learned via BGP, ordered from weakest to strongest: origin authentication [39,41], soBGP [49], Secure BGP (S-BGP) [32], and data-plane verification [6,50]. We also evaluate an orthogonal security mechanism: prefix filtering [6]. While the stronger protocols prevent a strictly larger set of attacks than the weaker ones, these security gains often come with significant implementation and deployment costs. To inform discussions about which of these secure protocols should be deployed, we would like to quantitatively compare their ability to limit traffic attraction attacks. Thus, we simulate attacks on each protocol on an empirically-measured AS-level topology [1,8,12], and determine the percentage of ASes that forward traffic to the manipulator.

Performing a quantitative comparison requires some care. It does *not* suffice to say that one protocol, say S-BGP, is four times as effective as another protocol, say origin authentication, at preventing a specific type of attack strategy; there may be other attack strategies for which the quantitative gap between the two protocols is significantly smaller. Since these more clever attack strategies can just as easily occur in the wild, our comparison must be in terms of the worst possible attack that the manipulator could launch on each protocol. To do this, we put ourselves in the mind of the manipulator, and look for the optimal strategy he can use to attract traffic from as many ASes as possible.

However, before we can even begin thinking about optimal strategies for traffic attraction, we first need a model for the way traffic flows in the Internet. In practice, this depends on local routing policies used by each AS, which are not publicly known. However, the BGP decision process breaks ties by selecting shorter routes over longer ones, and it is widely believed [18,27] that policies depend heavily on economic considerations. Thus, conventional wisdom and prior work [15,17,27–29] suggests basing routing policies on business relationships and AS-path lengths. While this model (used in many other studies, *e.g.*, [3,19,30]) does *not* capture all the intricacies of interdomain routing, it is still very useful for gaining insight into traffic attraction attacks. All of our results are obtained within this model.

1.2. Thinking like a manipulator

If routing policies are based on AS path lengths, then intuition suggests that it is optimal for the manipulator to use the following "smart" attack strategy: announce the shortest path that the protocol does not reject as bogus, to as many neighbors as possible. Depending on the security protocol, this means announcing: (a) a direct connection to the victim IP prefix (*i.e.*, a "prefix hijack" as in [9,40]), or (b) a bogus edge to the legitimate destination AS, or (c) a short path that exists but was never advertised, or (d) a short path that the manipulator learned but is not using, or (f) a legitimate path that deviates from normal export policy (*i.e.*, a "route leak" as in [44]). Indeed, we use simulations on a measured AS-level topology to show that this "smart" attack strategy is quite effective, even against advanced secure routing protocols like S-BGP and data-plane verification.

Worse yet, we use counterexamples show that our simulations *underestimate* the amount of damage manipulator could cause, because the "smart" attack is not optimal. In fact, the following bizarre strategies can sometimes attract even more traffic than the "smart" attack: announcing a longer path, exporting a route to fewer neighbors, or using "path poisoning" to trigger BGP's loop-detection mechanism (*cf.*, [31]). In fact, we present counterexamples that show that prefix hijacking (*i.e.*, originating a prefix you do not own) is *not* always the most effective attack against BGP! These counterexamples are not merely hypothetical—we identify specific ASes in the measured AS-level topology that could launch them. Moreover, we prove that it is NP-hard to find the manipulator's optimal attack, suggesting that a comprehensive comparison across protocols must remain elusive.

1.3. Our findings and recommendations

While we necessarily underestimate the amount of damage a manipulator could cause, we can make a number of concrete statements. Our main finding is that secure routing protocols only deal with one half of the problem: while they do restrict the paths the manipulator can announce, they fail to restrict his export policies. Thus, our simulations show that, when compared to BGP and origin authentication, soBGP and S-BGP significantly limit the manipulator's ability to attract traffic by announcing bogus short paths to all its neighbors. However, even in a network with S-BGP or data-plane verification, we found that a manipulator can still attract traffic by cleverly manipulating his export policies. Indeed, we found that announcing a short path can be less important than exporting that path to the right set of neighbors (an attack strategy that has also been called a "route leak" [11,44]). Thus:

- Advanced security protocols like S-BGP and data-plane verification do *not* significantly outperform soBGP for the "smart" attacks we evaluated.
- Prefix filtering of paths exported by stub ASes (*i.e.*, ASes with no customers) provides a level of protection that is at least comparable to that provided soBGP, S-BGP and data-plane verification.
- Tier 2 ASes are in the position to attract the largest volumes of traffic, even in the presence of data-plane verification and prefix filtering (of stubs).
- *Interception attacks* [3,9,45]—where the manipulator silently intercepts traffic and delivers it to the destination—are easy for many ASes, especially large ones.