Securing BGP Incrementally

Martin Suchara
Princeton University
msuchara@cs.princeton.edu

Ioannis Avramopoulos
Princeton University
iavramop@cs.princeton.edu

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

ABSTRACT

Despite the pressing need to secure routing, none of the existing secure variants of BGP has been widely deployed. Due to the size and decentralized nature of the Internet, it became clear that any viable secure routing protocol must offer benefits also in its early stages of deployment. In order to determine when the protocols are *not* adoptable, we quantify the benefits offered by a partial deployment of an Idealized Secure BGP which is able to detect malicious routes with perfect accuracy. We also quantify the benefits of an imperfect version of the protocol. Subsequently, we conclude that even the best protocols which simply detect and avoid bogus routes do not offer good security performance except in limited scenarios. We offer alternative designs, and hope that our insights will result in a new secure routing protocol that will be more attractive to early adopters.

1. INTRODUCTION

Even though packet routing and data delivery depends critically on the Border Gateway Protocol (BGP), it was designed for use in a trusted environment. As a result, security and reliability of the Internet today is threatened. First, invalid announcements resulting from misconfigurations may cause propagation of invalid routes. Second, malicious autonomous systems (AS) may attempt to either blackhole or transparently intercept the traffic addressed to a particular prefix. In a typical attack, the adversary would announce a bogus route that looks attractive to a subset of ASes because it is more profitable or shorter than the real route.

The need for routing security resulted in an array of new variants of BGP. The high level mechanism of most of these proposals is simple: 1) identify which routes are malicious and 2) reject the implicated routes. A number of protocols attempt to identify invalid routes in the control plane through cryptographic means. These protocols include the well-publicized S-BGP and soBGP [6]. Other protocols detect offending routes through data plane probes [7], and one protocol combines the two methods [5]. Quite surprisingly, despite years of active research, none of the proposals has been widely adopted. We explain this by delineating the narrow conditions under which a partial deployment of an idealized version of these protocols performs well.

*CoNEXT'*07, December 10-13, 2007, New York, NY, U.S.A.

2. SECURITY BENEFIT EVALUATION

We use the term secure protocol to refer to a variant of BGP where participating ASes filter insecure routes, i.e., routes that contain a malicious AS number in the path. Idealized Secure BGP is an idealized variant that upper-bounds the security offered by these protocols.

Deployment and threat model: We assume that some ASes deploy the secure protocol while others run the legacy version. We use the terms participants and non-participants, respectively, to refer to these two groups. In our model, one malicious AS attempts to either intercept or blackhole packets destined to a particular address prefix owned by a victim AS. The adversary behaves in a way that minimizes the security benefit of the deployment of the secure protocol. The security benefit is determined by calculating the fraction of ASes which either accept a route containing the malicious AS or whose routing tables do not contain any route to the prefix of the victim.

Idealized Secure BGP: A routing oracle is an algorithm which is given a route and outputs true if and only if it contains the adversary. Idealized Secure BGP works just as legacy BGP except it uses the oracle to filter malicious routes. After a route is filtered, the next best route (if any) is used. We conclude that Idealized Secure BGP offers the greatest possible security benefit among all secure protocols. If secure protocol X differs from Idealized Secure BGP, it must either accept some malicious route, or reject some valid route, degrading its performance. Next, we define the optimal strategy of the adversary.

Strategy of the adversary: There are two strategies that the adversary can utilize. The first is a false origination attack. Non-participants in the neighborhood of the adver-

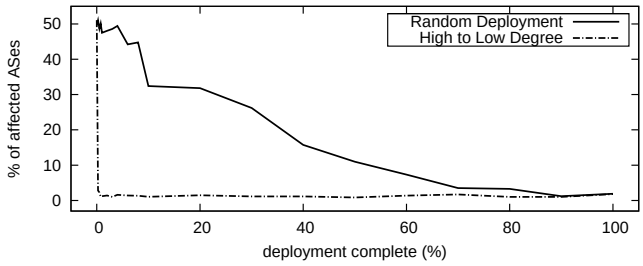

Figure 1: Evaluation of Ideal Secure BGP.

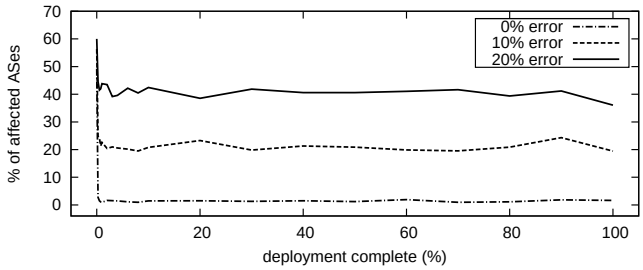

Figure 2: Evaluation of imperfect secure BGP.

sary will accept and propagate the malicious route because it is shorter than the true route. Depending on the variant of the secure protocol in use, the route may be filtered by the participants. The second strategy of the adversary is to spoof a short path in which the victim appears to originate its prefix. Because the route will be longer than in the first case, it may be less attractive to the non-participants. However, adding a few hops may allow the adversary make the route more attractive to the participants. Case in point: secure origin authentication accepting any routes with a genuine first hop. In general, the strength of the two attacks is not comparable. However, when Idealized Secure BGP is deployed, false origination is the optimal attack because participants always reject malicious routes.

Experimental setup: The security of Idealized Secure BGP is evaluated in a realistic setting. We use a dataset from RouteViews [2] to reconstruct the AS-level topology of the Internet. One randomly chosen AS is the victim and another is the adversary. Routing message propagation is simulated using a modified version of BSIM [1]. BSIM ensures that customer routes are preferred over peer routes, and peer routes over provider routes. If these rules do not result in a unique route selection, the shortest of the most preferred routes is chosen. Idealized Secure BGP participants modify this behavior and filter all routes that contain the AS number of the adversary. During our simulation, the victim announces its prefix first. Subsequently, the hijacker announces the same prefix utilizing the false origination attack. Experimental results are averaged over 500 runs.

Performance evaluation: The security benefit of Idealized Secure BGP is summarized in Fig. 1. First, the participants are selected uniformly at random with probability of participation ranging between 0% and 100%. Even if the protocol is widely adopted, a substantial fraction of the network accepts malicious routes or cannot reach the destination. We repeat the experiment several times changing the model of deployment. The performance of the protocol improves dramatically if ASes deploy in the order of their degree, e.g., when the participation level is 10%, one tenth of the ASes with the highest degree participates. The benefits of this deployment model were also observed in [4], which evaluates the security performance of PGBGP. We attribute the improved performance to the combination of two factors. First, high degree ASes learn many paths, so they often have at least one valid path. Second, if a high degree AS picks a good route, that route is propagated to many non-participating ASes. While deployment in the order of node degree is justified by the fact that larger well-connected ASes are more likely to deploy cutting edge technology, the assumption that a malicious route can be detected with perfect accuracy is not. To estimate the importance of accuracy of malicious route detection, we introduce false positives and false negatives into the decision process of Idealized Secure BGP. Fig. 2 shows that even if ASes deploy in the order of their degree, the performance degrades significantly as the error rate increases.

3. IMPLICATIONS AND FUTURE WORK

Our results suggest that while the best secure protocol offers substantial security benefits when deployed in the core of the network, the performance degrades significantly for more realistic protocols and/or deployment scenarios. While finding a protocol similar to Idealized Secure BGP may be possible, there are other promising approaches. First, overlay networks [3] can avoid malicious routes to a particular destination by re-routing traffic through an intermediate node. Second, it is possible to use false route announcements to override the routing tables of ASes which accepted malicious routes. Combination of these two approaches seems to be a particularly promising avenue for future research.

4. REFERENCES

[1] BSIM simulator: http://cs.unm.edu/~karlinjf.
[2] University of Oregon RouteViews project.
[3] D. Andersen, H. Balakrishnan, M. Kaashoek, and R. Morris. Resilient overlay networks. In *18th ACM SOSP*, October 2001.
[4] J. Karlin, J. Rexford, and S. Forrest. Pretty good BGP: Improving BGP by cautiously adopting routes. In *14th IEEE ICNP*, 2006.
[5] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. Katz. Listen and whisper: Security mechanisms for BGP. In *Proc. of USENIX/ACM NSDI*, 2004.
[6] R. White. Securing BGP through secure origin BGP. *Tech. report, Cisco Internet Protocol Journal*, 2003.
[7] D. Zhu, M. Gritter, and D. Cheriton. Feedback based routing. In *HotNets*, October 2002.