How to Take Back Your Address Space

Ioannis Avramopoulos
Princeton University

Submitted to HotNets-VI on August 3, 2007.
Superceded by Department of Computer Science, Princeton University Technical Report TR-808-07.

Abstract

Preventing adversaries from hijacking address space is important, but network operators are reluctant to deploy secure routing protocols. In this paper, we present Clout, a system that like secure routing protocols prevents prefix hijacking but, in contrast to these protocols, is easily deployable. Clout is deployable by unilateral action from a single party, or multilateral action from a moderate number of independent parties, without requiring changes to BGP or the data plane. In Clout, a collection of networks jointly defends a prefix by simultaneously announcing it in BGP, essentially *hijacking the hijacker*. Clout relies on the premise that the adversary can be outnumbered, a requirement attainable in practice. Deployment scenarios of Clout are also presented for emergency response and as a long-standing commercial service.

1 Introduction

There are frequent reports of *prefix hijacking*, i.e., the illegitimate control of address space, through the announcement of the victim prefixes in BGP. On one hand, prefix hijacking is so easy to perform that it often happens by accident. On the other hand, if a prefix is hijacked, the traffic destined to that prefix is delivered to the offender. The offender may discard the traffic to render the destination unavailable, modify the payload, impersonate the destination, inspect unencrypted traffic to read the payload, or inspect encrypted traffic to perform traffic analysis from the payload and the headers.

Preventing prefix hijacking is important. Among the proposed countermeasures to prevent prefix hijacking secure routing protocols, notably S-BGP [13] and soBGP [23], have received the bulk of the attention. However, ISPs have been reluctant to deploy them. Secure routing protocols are heavyweight requiring coordinated deployment by a large group of networks that must participate in a public key infrastructure, change BGP, and upgrade their routers. Furthermore, the possible deployment of a secure routing protocol would not entirely prevent prefix hijacking as that would still be possible through *collusion* attacks in which remote adversarial ASes would announce they are directly connected to attract the traffic.

In this paper, we present Clout, a system that like secure routing protocols prevents prefix hijacking but, in contrast to secure routing protocols, is easily adoptable. The basic idea in Clout is simple: A collection of ASes acting jointly can take back the address space the hijacker is trying to subvert by *hijacking the hijacker*. Clout is based on the same technique employed by the hijacker—given a prefix being defended, all member ASes simultaneously announce it in BGP. In this way, assuming that that the Clout group outnumbers the group of hijackers, traffic destined to the corresponding prefix is more likely to arrive at an AS belonging to the Clout group than an AS controlled by the adversary. Essentially Clout starts a war of attrition [19, 20] for the address space the adversary is trying to subvert.

Clout pulls traffic destined to the protected prefix away from the adversarial ASes and toward members of the collection of defending ASes. There are cases where the traffic is successfully protected as long as it reaches *any* member of the defending group. An example is when communication is of the *anycast* type [16]. For the rest of the cases, Clout must ensure that traffic reaches the origin. To that end, the members of the Clout group are connected with pairwise tunnels to form an overlay network. An overlay routing protocol running on top of this network directs traffic to the corresponding origin once it reaches any member. Vis-a-vis the straightforward approach of forwarding the traffic over the direct virtual link between the Clout member and the origin, overlay routing decreases the probability that the adversary is able to capture the traffic. This is demonstrated by the simulation results of Section 4. The intuition behind this result is that the probability that the Clout overlay remains connected so that the origin is reachable can significantly exceed the probability that individual links of the overlay are hijacked.

Clout does not require changes to BGP (as the announcement of extra prefixes requires a simple configuration change and a possible notification of the upstream provider) or the data plane (as routers increasingly offer IP tunneling at line rate). Furthermore, Clout is easily deployable by unilateral action from a single organization, or multilateral action from a moderate number of

independent organizations. For example, Clout is easily deployable by a content delivery network [3]. Alternatively, the owner of a prefix may deploy inexpensive BGP-speaking software routers [1] at diverse locations for self-defense, or a group of independent networks may decide to *join forces*. Therefore, Clout can be deployed in an independent fashion without relying on community consensus, for example, in response to a crisis. Furthermore, Clout retains its strength even if malicious ASes are colluding. Hence, Clout is an attractive alternative to secure routing protocols that also remains valuable even if a secure routing protocol is deployed.

The rest of the paper is organized as follows. In Section 2, we clarify terminology and, in Section 3, we present Clout and its detailed workings. In Section 4, we use simulation to demonstrate that Clout achieves significant gains in security at moderate group sizes. Section 5 discusses deployment scenarios of Clout in response to an emergency or as a long-standing commercial service. Related work is discussed in Section 6 followed by the conclusion in Section 7.

2 Terminology

In this paper, we say that an adversary has *hijacked* a prefix, if the adversary is able to receive the traffic destined to that prefix. Oftentimes, the term *prefix hijacking* refers to a particular means the adversary can employ to receive the traffic to a prefix, namely, the announcement of the prefix in BGP such that the adversarial AS is the origin AS. Instead of originating a prefix, the adversary can also receive the corresponding traffic by a *path-spoofing* attack in which the AS that the adversary controls appears upstream of the origin. In this paper, we assume that the adversary always prefers to originate a prefix instead of spoofing a path to that prefix. If filters by the upstream providers do not prevent an adversary from originating a victim prefix, a path-spoofing attack is strictly weaker than a attack where the prefix is originated. Strictly weaker are also the aforementioned collusion (or *wormhole*) attacks. Finally, another means that the adversary can employ to receive traffic to a prefix is *sub-prefix hijacking*, a particular form of prefix hijacking, which we discuss in Section 3.4.

Note also that we use the terms *group* and *collection* interchangeably. Oftentimes, the term *group* refers to a collection of independent entities such as independent organizations. Clout may be deployed either by a single organization controlling multiple ASes or jointly by multiple organizations each controlling one or more ASes.

3 Clout

Clout prevents prefix hijacking through a coordinated announcement of the address space it protects from a collection of ASes. In this way, corresponding traffic sources are more likely to accept a route pointing to a member of the collection of defending ASes instead of a member of the collection of attacking ASes. Clout relies on the premise that the collection of defending ASes outnumbers the collection of attacking ASes. In the following, we present Clout in four steps starting with the basic idea and refining it at each step. We end the section with a variant of Clout that facilitates rapid deployment.

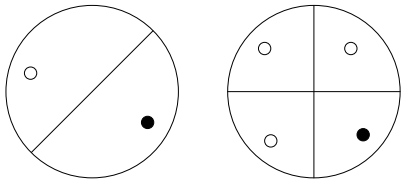

Figure 1: The left figure shows a partition of the network into two subsets when a prefix is announced from two origin ASes. The right figure shows a partition of the network into four subsets when a prefix is announced from four origins. The black dot corresponds to an adversarial AS. In the first case, approximately one half of the network points to the adversary whereas, in the second case, the percentage of the network pointing to the adversary is approximately one quarter.

3.1 Hijacking the Hijacker

The basic idea in Clout is simple: If address space is announced simultaneously from multiple origin ASes, each source of traffic destined to that address space will point to one of the origins, creating a partition of the sources according to the origin they point to (shown in Figure 1). This is the condition that an adversary exploits to hijack a prefix. Clout also exploits this condition to recover the traffic sources from the adversary. Clout achieves this by announcing the victim address space from multiple ASes acting jointly with the victim. As the size of the collection of defending ASes increases, the percentage of sources pointing to a Clout member, instead of an adversarial AS, also increases (see Figure 1). Clout relies on the premise that the Clout collection outnumbers the collection of ASes the adversary controls. Simulation results relating efficiency with the size of the defending group are given in Section 4.

The simultaneous announcement of address space from multiple origin ASes may cause the network to direct traffic from different sources attempting to communicate with the same IP address to different destinations. Although there are cases where a behavior like this is a limitation, oftentimes this behavior is desirable. An example is anycast [16], which has been used to improve the resilience of the root DNS servers [8]. This behavior is also compatible with existing content distribution

systems. For example, if content is replicated on multiple servers, the content provider would prefer that users independently reach their closest server. Therefore, assuming an appropriately engineered content distribution system, preventing an adversary from successfully disrupting or intercepting communication through a routing attack may be as simple as originating address space from multiple origins.

If replication of a service is infeasible, communication must take place between fixed endpoints. However, if a prefix is announced from multiple origins, traffic destined to an IP address of that prefix may reach a different AS than the one hosting the intended destination. In this case, traffic must be redirected to the intended destination. In the following, we present how Clout achieves this redirection.

3.2 Forming an Overlay

Consider a Clout group protecting prefix p of origin O. Consider also a member X of this group. Traffic arriving at X and destined to an address in p must be able to reach O. However, because X announces p as if X is its true origin, p may not be reachable from X. We reestablish the reachability of p from members of the Clout group using tunnels. For any two distinct members X and Y of the Clout group, including the origin O, we create two tunnels one from X to Y and one from Y to X; tunnel endpoints are addressed from different prefixes than the one being protected. If traffic destined to p arrives at any member of the Clout group except the origin, then this traffic is diverted either to the tunnel leading directly to O or possibly to another Clout member. Such forwarding decisions are decided by a routing protocol presented next.

3.3 Routing in the Overlay

A straightforward method to deliver traffic arriving at a Clout member X to the origin O is to divert it to the direct tunnel between X and O. However, if the Clout group is under attack, the probability that this delivery method is successful is small. The reason is that the adversary can successfully attack the tunnel endpoint at O with high probability. In particular, if X, O, and the adversary are chosen at random, the probability that X can reach the tunnel endpoint while the adversary is hijacking the endpoint is one half.

Although the probability that the direct virtual link between X and O is free from the adversary is one half, as we show in the next section, the probability that there is a path from X to O in the overlay graph that is free from the adversary can be significantly greater. We find such paths using an overlay routing protocol. The routing protocol selects paths from each member of the group to every origin (or prefix) protected by the group using information from probes between the members. Each AS periodically probes other member ASes to determine the availability of the corresponding virtual links. Probes are authenticated, using pairwise security associations, to increase the probability that they can detect whether a tunnel endpoint has been hijacked. Authentication prevents an adversary that has hijacked a tunnel endpoint from impersonating that endpoint. In this way, if the adversary wants to receive the traffic destined to a tunnel endpoint but wants to avoid being detected by the probes, then he must maintain a path to that tunnel endpoint at the same time he is hijacking the endpoint, which is only possible with a small probability [6]. If the adversary cannot maintain a path to the tunnel endpoint, the probes will indicate that the virtual link is unavailable and the adversary will be circumvented. The results of probes are disseminated to the rest of the members using a link state routing protocol similar to RON [4] and the members select routes that avoid the unavailable virtual links.

Note that there is a small probability that the adversary can hijack a tunnel endpoint and avoid detection by the probes. In such event, the adversary may be able to breach the availability of the corresponding communication by discarding the data traffic. Tieing the results of the probes with the fate of the data traffic will be able to recover the loss of availability [5]. The adversary may also be able to breach the confidentiality of the corresponding communication. However, to the best of our knowledge, no existing technique can entirely prevent undetectable *interception*. For example, secure routing protocols are vulnerable to *collusion* or *wormhole* attacks in which groups of remote adversarial ASes announce in BGP direct virtual links between them making the paths crossing the virtual links attractive. Through wormhole attacks the adversary can intercept remote traffic even if a secure routing protocol is deployed.

Note also that, because of Clout, data traffic takes detours in the network that may negatively affect performance. Such inefficiency can be corrected if the routing protocol selects good-performing paths using the results of the probes. Other techniques to avoid the inefficiency are also possible. For example, Clout can be deployed only in response to an alarm by a protocol that passively detects prefix hijacking.

Using the probes of a RON routing protocol is only one of the techniques that the Clout group can employ for path selection. We are exploring path-selection techniques that can be employed alternatively or in a complementary manner. For example, preferring virtual links corresponding to shorter BGP paths can further decrease the hijacking capabilities of the adversary.

3.4 Preventing Sub-prefix Hijacking

Thus far we have assumed that the ASes of the adversarial group and the ASes of the defending group announce the same prefix. However, unless care is taken, the adversary may defeat the defending group by employing a variant of prefix hijacking. In the following, we discuss this variant and the corresponding countermeasure employed by Clout to defend against it.

In order to gain control of a prefix, instead of originating that prefix in BGP, the adversary may break the prefix into multiple sub-prefixes and originate those instead. In this way, although the network will maintain routes to both the original prefix and the sub-prefixes, because of the *longest prefix matching* rule used in the data plane, traffic will be directed to the subprefixes. Therefore, using this technique, the adversary can gain full control of the traffic destined to the victim prefix.

Counteracting sub-prefix hijacking is possible if the address space is announced at the finest possible granularity in a process known as *deaggregation*. Deaggregation at the finest granularity prevents the adversary from announcing a more specific prefix than the defender. The finest granularity that address space can be deaggregated is determined by ISP filtering rules according to which any BGP announcement of a prefix more specific than a $/24$ is discarded. Note that deaggregation often happens today possibly as a defensive countermeasure against sub-prefix hijacking [18]. Clout also relies on deaggregation to protect against sub-prefix hijacking.

3.5 Enabling Rapid Deployment

Whenever a new prefix is added to the list of prefixes protected by a Clout group, the members' upstream providers (if they have any) need to be contacted to adjust their route filters so that the new prefix is permitted. Sometimes this process may be slow. We present here a technique that can circumvent this step. Making use of this technique may sacrifice control over the address space to the adversary's advantage. From a security standpoint, it is, therefore, intended to be used only during the transitive period where the filters are adjusted.

According to this technique, instead of announcing the new prefix p in BGP as if the corresponding Clout member is the origin of the prefix, the Clout member announces a one-hop path leading to p. This is sufficient to circumvent the route filter at the upstream provider. The establishment of pairwise tunnels with the rest of the Clout members and the overlay routing protocol are retained. Note that since the adversary's hijacking strategy remains unchanged, this protocol sacrifices security. Therefore, once an upstream provider adjusts the filter, the corresponding Clout member should start announcing the prefix as if the member is the prefix's origin.

4 Evaluation

In this section, we use simulation based on a realistic Internet topology to demonstrate that Clout can effectively prevent prefix hijacking. We only evaluate the case where traffic to the victim prefix must be delivered to the corresponding origin. If delivery of the traffic to any member of the Clout group suffices, then it is easy to calculate Clout's effectiveness. In particular, if N is the size of the Clout group and M is the size of the adversary's group, then, on average, the percentage of ASes that accept routes to the Clout group is $N/(N + M)$. In the evaluation, we assume that the probes can alway detect whether a virtual link has been hijacked.

4.1 Methodology

Our simulator is based on BSIM [10], which simulates BGP policy-based routing on an AS topology annotated with business relationships. We used an annotated AS topology from June 2007 available from CAIDA [2].

In the simulation, we simultaneously announce a common prefix p from multiple ASes. The set of ASes announcing p is divided into two groups; the adversarial and the defending group. For each member X of the defending group, we count the number of ASes having accepted a route leading to X. Then, we construct the overlay graph of the defending group. Each member X of the defending group announces a prefix p_X simultaneously with the adversarial group. If X is reachable from Y, where Y is another member of the defending group, we add a link from Y to X. Once the overlay graph is constructed, we count the number of ASes that are able to reach the origin by some path in the overlay graph.

We vary the number of ASes in the defending group covering a range between 1 and 35. The number of ASes controlled by the adversary is 1 or 2. Given the size of the defending and the adversarial groups, we repeat the simulation 50 times and show the average.

A simplifying assumption we have made is that the ASes in the defending group and the ASes controlled by the adversary are chosen at random from the set of all ASes. In practice both the adversary's group and the defending group may try to optimize the relative location of their members subject to the availability of vulnerable ASes for the adversary and of synergetic ASes for the defending group. In this paper, we are interested in the average outcome.

4.2 Results

Figure 2 shows the percentage of ASes that are able to reach the origin of the defending group either directly

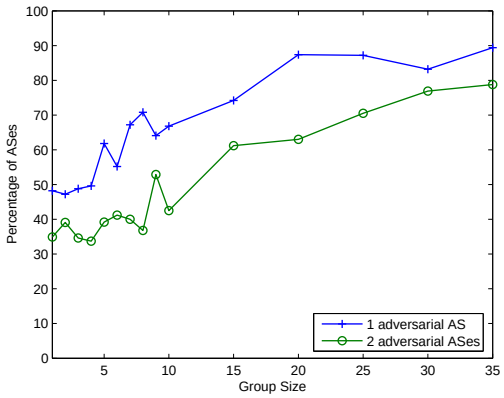

Figure 2: Percentage of ASes able to reach the origin of the defending group when the origin and the defending group are under attack as the size of the defending group varies.

or via the Clout overlay while the origin and the Clout overlay are under attack. The number of ASes participating in the group varies in the range from 1 to 35. The group of the adversary is assumed to consist of one or two members.

We notice that for a small number of participants, approximately up to 5 in the case of a single adversarial AS and approximately up to 10 in the case of two adversarial ASes, there are no gains in security as shown by the plateaus in the graphs. Gains are achieved after the corresponding threshold values of 5 and 10 are reached.

As the group size increases beyond these values, the security gains increase as well. The percentage of ASes able to reach the origin in the case of a single adversarial AS is approximately 90% when the size of the defending group is 35. The percentage of ASes able to reach the origin in the case of two adversarial ASes is approximately 80% when the the defending group also consists of 35 members.

Note that the use of the overlay routing protocol significantly increases the percentage of ASes able to reach the origin in comparison to a scheme where a group member blindly forwards the traffic to the origin. For example, it is straightforward to show that, in the blind scheme, the percentage of ASes able to reach the origin cannot exceed 50% on average.

Also note that the reason the defending group must outnumber the adversarial group by a significant number is that the measure of success is different for the two groups. In the case of the adversarial group, we have assumed that an attack is successful as long as the traffic reaches any member of the adversarial group. In the case of the defending group, traffic must be delivered to the origin through the overlay network that is also under attack. However, as the size of the defending group grows, the adversary is less able to partition the overlay graph and capture the traffic.

5 Deployment Scenarios

Clout is deployable by unilateral action from a single organization, such as an ISP or an overlay provider. This feature is particularly attractive since coordination among ISPs has proved difficult in practice. Furthermore, Clout is implementable by simple configuration changes in the control plane, to announce additional prefixes, and in the data plane, to configure tunnels. Therefore, Clout is amenable to deployment in an independent fashion without relying on community consensus or prior steps by the vendors, which is advantageous at times of crisis such as a natural disaster or a physical attack. It is during such periods when critical functions must rely on the Internet[1] that a cyber attack would be felt the most. Independent deployment would also enable timely response against an attack that targets the Internet itself.

Clout can also be deployed as a long standing commercial service to ensure, for example, the confidentiality and integrity of communications as an alternative, or possibly in a complementary manner, to cryptographic techniques such as IPsec [12] and SSL [7]. Commercial deployments can be initiated either by ISPs or overlay providers such as content delivery networks (CDNs) [3].

The deployment of Clout by a CDN, for example, is appealing for several reasons, which we explain next. First, a CDN is typically comprised of numerous networks. Therefore, a CDN can form large Clout groups that can significantly outnumber and, therefore, outrival the adversarial groups. Second, a CDN can proactively store replicated content in the group members, obviating the need to deliver the traffic to the origin. Serving requests locally through proactive content replication significantly increases the resilience of a Clout group. Third, to avoid harming its reputation, a CDN would be eager to confirm that the true owner of a prefix is the one requesting the protection service. Finally, a CDN can offer this service without an additional investment in infrastructure.

6 Related Work

Previous countermeasures against prefix hijacking can be classified according to whether they try to prevent or detect prefix hijacking.

Network operators are known to *deaggregate* their address space (i.e., break a prefix into smaller prefixes and announce the smaller prefixes in BGP) as a defensive countermeasure against prefix hijacking [18]. This coun-

[1]For example, in the aftermath of 9/11, a New York City hospital relied on an external link for the retrieval of medical records that was temporarily broken by the collapse of the Twin Towers [15].

termeasure prevents prefix hijacking from subduing the entire Internet. However, the adversary retains significant control over the hijacked address space.

Secure routing protocols such as [13, 21, 23] employ heavyweight proactive techniques to prevent prefix hijacking based, in part, on cryptography. However, ISPs have been reluctant to deploy these protocols.

Protocols such as [9, 14, 24] employ anomaly-detection techniques to detect prefix hijacking. However, these protocols provide only part of the solution without specifying the recovery procedure following the detection of prefix hijacking.

PGBGP [11] employs a combination of detection and prevention techniques against prefix hijacking. Anomaly detection is used to flag routes as suspicious that are subsequently depreferenced for a configurable time interval.

That the routing system should not be concerned with preventing prefix hijacking but rather with ensuring availability using availability-centric routing (ACR) was argued in [22]. Clout remains valuable even if the protection of availability is the only objective of the routing system. Furthermore, ACR is a mostly sender-driven technique whereas Clout is receiver-driven.

Clout is similar in spirit to Crowds [17], a system that anonymizes web browsing. They both leverage the collective resources of a group to ensure the security of the individual members.

7 Conclusion

Prefix hijacking is easy to perform. Unfortunately it may also have serious consequences as the traffic destined to the victim prefix is delivered to the offender. Therefore, preventing prefix hijacking is important. In this paper, we presented Clout, a system that prevents prefix hijacking by having a group of ASes simultaneously announce the prefix being protected in BGP, essentially *hijacking the hijacker*. Clout relies on the premise that the group protecting the prefix significantly outnumbers the group of the adversary. In practice, large Clout groups can be based on existing content delivery networks, they can be built independently by the prefix owners, or prefix owners may decide to form coalitions. An attractive feature of Clout is that it is amenable to independent deployment based on unilateral action by individual networks or multilateral action by groups of networks of moderate size. Clout can be deployed in a timely manner in response to an ongoing crisis without waiting for the operations community to reach consensus.

References

[1] `http://www.vyatta.com/`.

[2] `http://as-rank.caida.org/data/`.

[3] `http://www.akamai.com/`.

[4] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilent overlay networks. In *Proc. ACM Symposium on Operating System Principles*, Oct. 2001.

[5] I. Avramopoulos and J. Rexford. Stealth probing: Efficient data-plane security for IP routing. In *Proc. USENIX Annual Technical Conference*, May/Jun. 2006.

[6] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the Internet. In *Proc. ACM SIGCOMM*, Aug. 2007.

[7] T. Dierks and C. Allen. The TLS protocol version 1.0. RFC 2246, IETF, Jan. 1999.

[8] T. Hardie. Distributing authoritative name servers via shared unicast addresses. RFC 3258, IETF, Apr. 2002.

[9] X. Hu and Z. M. Mao. Accurate real-time identification of IP prefix hijacking. In *Proc. IEEE Symposium on Security and Privacy*, May 2007.

[10] J. Karlin, S. Forrest, and J. Rexford. *PGBGP simulator*. `http://www.cs.unm.edu/~karlinjf/pgbgp/`.

[11] J. Karlin, S. Forrest, and J. Rexford. Pretty Good BGP: Improving BGP by cautiously adopting routes. In *Proc. IEEE ICNP*, Nov. 2006.

[12] S. Kent and R. Atkinson. Security architecture for the Internet protocol. RFC 2401, IETF, Nov. 1998.

[13] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (Secure-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, Apr. 2000.

[14] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang. Phas: a prefix hijack alert system. In *Proc. USENIX Security Symposium*, Aug. 2006.

[15] C. Partridge, P. Barford, D. Clark, S. Donelan, V. Paxson, J. Rexford, and M. Vernon. *The Internet under Crisis Conditions: Learning from September 11*. The National Academies Press, 2003.

[16] C. Partridge, T. Mendez, and W. Milliken. Host anycasting service. RFC 1546, IETF, Nov. 1993.

[17] M. Reiter and A. Rubin. Anonymity loves company: Anonymous Web transactions with Crowds. *Communications of the ACM*, 42(2), 1999.

[18] P. Smith, R. Evans, and M. Hughes. RIPE routing working group recommendations on route aggregation. Document ripe-399, RIPE, Dec. 2006.

[19] M. Walfish, H. Balakrishnan, D. Karger, and S. Shenker. DoS: Fighting fire with fire. In *Proc. ACM SIGCOMM HotNets Workshop*, Nov. 2005.

[20] M. Walfish, M. Vutukuru, H. Balakrishnan, D. Karger, and S. Shenker. DDoS defense by offense. In *Proc. ACM SIGCOMM*, Sept. 2006.

[21] T. Wan, E. Kranakis, and P.C. van Oorschot. Pretty secure BGP (psBGP). In *Proc. Network and Distributed System Security Symposium*, Feb. 2005.

[22] D. Wendlandt, I. Avramopoulos, D. Andersen, and J. Rexford. Don't secure routing protocols, secure data delivery. In *Proc. ACM SIGCOMM HotNets Workshop*, Nov. 2006.

[23] R. White. Securing BGP through secure origin BGP. *The Internet Protocol Journal*, 6(3), 2003.

[24] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis. A lightweight distributed scheme for detecting IP prefix hijacks in real-time. In *Proc. ACM SIGCOMM*, Aug. 2007.