How Secure are Secure Interdomain Routing Protocols?

Anatol Dammer
Advisor: Dr. Nils Kammenhuber
Seminar Future Internet SS2010
Chair for Network Architectures and Services
Fakultät für Informatik, Technische Universität München
Email: anatol.dammer@mytum.de

ABSTRACT

Ever since the 1990s, the de facto standard for Internet inter-AS[1] routing has been BGP, the Border Gateway Protocol. Security issues caused or abetted by BGP, some of which have been known for considerable time, have become increasingly apparent. Long-running efforts of making BGP and inter-AS routing more secure have produced a number of proposals, none of which have managed to gain traction. This is at least partly due to the fact that even the most popular and well-regarded proposals fail to prevent strategic attacks. We provide an overview of several popular proposals and how they address, or fail to address, a range of attacks on inter-AS routing.

1. INTRODUCTION

In recent years, several high-profile attacks and outages caused by exploitation of BGP's flaws or simple misconfiguration have risen awareness of actually long-known deficiencies of inter-AS routing. In 1997, a misconfigured border router of one AS led to major Internet-wide disruptions lasting up to a few hours [3], in 2008 Youtube.com was unreachable for several hours for most of the Internet, due to misconfiguration at Pakistan Telecom [12, 14], and in 2010 IDC China briefly announced 40,000 prefixes owned by other entities [10], attracting traffic for those destinations. In 2002, 200–1200 routing prefixes per day were found to suffer from misconfiguration, with about 15 prefix hijacks occurring per day [11]. BGP has been the de facto standard for inter-AS routing ever since the 1990's, and the protocol has not changed fundamentally since then – this alone should raise a few flags, considering the explosive growth of the Internet and its increasingly complex dynamics. Also, it is clear that if simple misconfigurations can have such considerable impact on the Internet, the potential for deliberate, strategic attacks should be quite profound.

Introductions on BGP usually emphasize the fact that BGP relies on an optimistic approach to routing, basically trusting routing information received by peers blindly. As will become apparent, this is not the whole truth: while BGP by itself is certainly not a very secure protocol, attacks on inter-AS routing can also hugely benefit from other, partly non-technical, aspects like business relationships between network operators. The quantitative data by Goldberg et al. [5] shows how relatively simple attack strategies can easily diminish the benefits promised by proposals such as S-BGP, which at first might appear to provide very substantial gains in security. On the other hand, they also show how comparatively simple measures could actually prevent a large proportion of attacks.

After an introduction to inter-AS routing and BGP, this paper succinctly describes and then compares four approaches to improve several security aspects of inter-domain routing. Main source for this information is the paper by Goldberg et al. who ran simulations of various inter-AS-level attacks on an internetwork model based on Internet AS-graph data sets, and published quantitative information on how well those four major security proposals fared.

[1]Autonomous system, a collection of networks administered by one entity, e.g., a large corporation

2. INTER-AS ROUTING AND BGP

As its name suggests, the Internet is a network of networks. Due to the very large number of destinations reachable in the Internet, routing tables can not sensibly include all single destinations. This motivated a routing scheme where destinations are aggregated into *prefixes*. Also, since organizations often want to have sole authority over routing in their own networks, an organization's networks can be combined into one or more so-called *Autonomous Systems (AS)*, each carrying a unique number (*ASN*) assigned by IANA[2]. For example, large corporations and Internet service providers operate their own AS(es).

To establish connectivity to the Internet, an AS operator employs so-called *border* or *gateway* routers that exchange inter-AS routing information with other AS border routers, route traffic between the inner part of the AS and the Internet, and may also act as intermediaries for traffic between two other ASes. Border routers establish "peer" relationships with other border routers via BGP, and can then exchange prefix routing information, which may be called sending route or path "announcements", and make forwarding decisions based on this information. For example, a border router can *originate* prefixes, which means announcing a network prefix included in its own AS, or *propagate* routing information learned from other routers, offering the own AS as an intermediary willing to proxy traffic along such a path.

BGP is a path vector protocol; the routing information it disseminates includes the full path, specified by ASNs, to reach a destination. For this, a router *prepends* its own ASN to a path attribute in the BGP path announcement message[3].

[2]Internet Assigned Numbers Authority, `http://www.iana.org/`

[3]This is a slight simplification; the PATH attribute in BGP UPDATE messages can be more complex – for our purposes, this is irrelevant

doi: 10.2313/NET-2011-07-2_10

The case where an AS acts as an intermediary for traffic between two other ASes is a good starting point for introducing a very important aspect of inter-AS routing in the current Internet: business relationships. While intra-AS routing is mainly concerned with purely technical aspects such as finding and distributing shortest paths, inter-AS routing involves different, possibly competing, organizations and is thus heavily influenced by political and business decisions. A protocol for inter-AS routing has to offer support for enforcing policies based on such decisions. BGP offers support for *import* and *export* policies, which respectively control which routes from BGP peers are entered into a BGP router's local route database and which routes are announced to BGP peers.

To provide an example: a network operator might like to only relay traffic between two parties if at least one of the parties pays for this service, usually by data volume. In addition to this *customer–provider* relationship, organizations such as major telecommunication companies also enter into so-called *peering* agreements: two organizations see themselves as peers in that they both benefit about equally from exchanging traffic, and are thus willing to mutually waive traffic fees. These relations allow for a classification of organizations into *Tiers*. "Tier 1"-providers have only customers and peers; because they do not have a "default route" to a provider, they constitute what is called the *Default-Free Zone* (DFZ) and are entirely reliant on peering agreements and customer contracts for connectivity. "Tier 2" providers, the most common providers in the Internet, have peering agreements but are also customers to Tier 1 providers. Tier 3 providers usually entirely rely on higher-tier providers, etc. Another concept that will be relevant later on are *stubs*, which are ASes that are only connected to one other AS and do not have any customers.

3. ROUTE SELECTION AND POLICIES

To understand the attacks that will be discussed later on, it is necessary to understand the criteria BGP uses to select routes and make forwarding decisions.

3.1 Route selection

Basically, a BGP router takes all routes it receives from its neighboring BGP routers, performs basic checks (the most relevant for us being routing loop detection), then runs all remaining routes through a decision process that decides if the routes are new or better than existing routes. Loop detection is based on the route path – if the own ASN is included in the path, the route information is discarded. Otherwise, a degree of preference for each route is calculated based on local preference, shortest AS path and tie-breaking rules, in that order. Local preference usually reflects policy decisions. Note that the path length comes second – a strong reminder of how important policy decisions are, and an aspect that will become important for attack strategies later on. After calculating the degree of preference, the best route for each destination is chosen and installed in a table that serves as input to the algorithms that make forwarding and route export decisions.

3.2 Policy scenarios

The aforementioned business relationships inherent to inter-AS routing have strong influence on which routes are exported by a router. ASes likely select and export routes

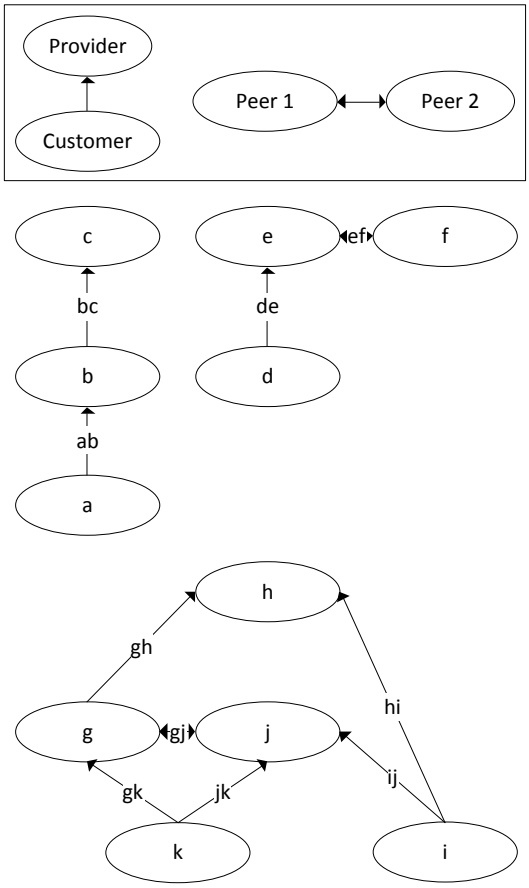

Figure 1: Routing policy examples

such that their own financial gain is maximized and financial loss is avoided unless absolutely necessary, e.g., to preserve connectivity. A few basic cases are illustrated in Figure 1. Here, AS b would export the route $a \to b$[4] to c to make its customer's AS available to the Internet (assuming c provides further connectivity), paying to its provider c but also getting paid by customer a. AS e would export the route $d \to e$ to its peering partner f – while e loses no money by relaying traffic to and from d over $e \to f$, it gains money from its customer d in the process. Likewise, f would not export a route to d, as doing so would mean using up capacities without gaining money from forwarding traffic over $e \to f$. In the last example, h will export the route $h \to i$ to g just like j will export $i \to j$ to g. AS g will then choose the peering link $g \to j$ to reach i, as this means avoiding costs for using the so-called transit or provider link $g \to h$. For some AS x, a *customer link* is a link to a customer of x, like $h \to i$ is a customer link to h.

[4]Note that route names were simply chosen alphabetically – in a BGP message, ASes prepend their ASN to the path, so *ba* would be a more "realistic" name for *ab*

doi: 10.2313/NET-2011-07-2_10

Figure 1 shows ASes in top-to-bottom hierarchical order with providers above their customers. This allows for easy illustration of the concept of *valley-free routing*, which directly follows from the business aspects of inter-AS routing. Simply put, paths are usually established such that packets never cross "valleys" in this hierarchical graph, such as the one created by the stub k. More precisely, packet flow conforms to the following scheme:

1. Travel upstream, i.e., towards a provider, across zero or more links
2. Traverse at most one peering link
3. Travel downstream, i.e., towards a customer, across zero or more links

The rationale for valley-free routing quickly becomes apparent if one considers each step and verifies that routes not conforming to the scheme would create financial loss for at least one AS.
In the following, we assume that every "honest", that is, non-malicious, AS follows these policies.

4. SECURITY PROPOSALS

Goldberg et al. mainly evaluate four different security protocols and plain BGP. While there are more specific proposals, the protocols they chose cover many proposals in terms of the security guarantees they provide[5]. Their order is strict from weakest to strongest security guarantees: any attack that is possible against a stronger protocol is also possible against all weaker protocols. An important factor to consider for all protocols is the substantial challenge of introducing a new protocol into the world of inter-AS routing, especially if computationally intensive cryptography would suddenly have to be performed by routers.

4.1 Origin authentication

Aiello, Ioannidis and McDaniel address the problem of address ownership [1]. In plain BGP, any AS can claim ownership of any prefix. This obviously provides ample opportunity for prefix hijacking attacks[6], and anomalies such as the one caused by AS 7007 in 1997 [3]. They state that origin authentication is a necessary but insufficient precondition for any inter-AS routing security infrastructure. Their fundamental work describes approaches to building a system that, from a database, can verify if a prefix announced by an AS has been assigned to that AS by an organization which in turn can provide a chain of address delegation up to IANA, the root authority for address assignment. In experiments, they found evidence that their approach should be deployable in terms of resource cost.

4.2 soBGP

On top of origin authentication, Secure Origin BGP (soBGP), described by Russ White et al. [15], proves validity of a path originated by an AS. Validity in this case means a path that physically exists in the Internet: The route path consists of real, interconnected ASes. Validation is provided by having routers disseminate signed local topology information, i.e., routers announce their peers to other routers, in effect establishing a global topology graph that every router knows. An attacker might still announce some path that is not actually available because it violates one of the standard policies of intra-AS routing described in section 3.2. While running attacks in an internetwork secured with soBGP requires knowledge of physically existing paths, such information can be obtained without too much effort – for example, from the very database that soBGP requires and maintains, as Goldberg et al. note. soBGP requires a PKI for origin authentication and path validation. Adjustments to BGP, such as a specific message type for exchange of security information, are suggested but, according to the authors, not necessary [16].

4.3 S-BGP

S-BGP, proposed by Kent et al. [9], provides path verification, meaning that an AS a can only announce a path $a \rightarrow b \rightarrow c$ if b actually announced $b \rightarrow c$ to a. S-BGP requires a PKI[7] that supports certificates for prefix ownership and granting authorization to ASes for announcing specific paths to specific prefixes. Simply put, path verification is achieved by a chain of signatures in route advertisements. This, combined with origin authentication provided by the PKI, seems to provide considerable security as a can only announce actually available paths that end with the rightful owner of a prefix. Besides a few other comparatively minor issues, an interesting aspect is that S-BGP does not ensure correct and honest application of policies by BGP peers. For example, nothing stops an attacker from announcing one path but actually forwarding incoming traffic that is meant for that path on an entirely different path.
BGP usually transmits messages in plaintext over TCP. S-BGP addresses this important security issue by using IPsec for all BGP messages. This ensures integrity, sender identity and even protection against message replay and DoS attacks which can be a significant problem with TCP.
The substantial amount of cryptography entailed by an Internet-wide deployment of S-BGP might seem challenging. One requirement for S-BGP was deployability and scalability; when the paper [9] was published in 2000, the authors concluded that deployment was feasible.

4.4 Data Plane Verification

A still relatively new research effort with groundwork by Wong et al. [17, 4] concerns itself with the actual path that data takes when it is forwarded by BGP routers. As mentioned, a router might advertise one path, but forward data on a different one. An AS might advertise an attractive path which would actually incur financial loss for the advertiser, and then use a cheaper path to forward the attracted traffic. S-BGP only protects the *control plane*, where routing information is exchanged. Goldberg et al. propose a verification scheme that works with shared secrets between routers along a route path. Basically, data packets are used as probes: a router can tag data packets with secrets shared with a router along the prospective route path. Only the expected

[5] A more comprehensive description can be found in [4]
[6] An attacker hijacks a prefix by directing traffic meant for that prefix to himself
[7] Public Key Infrastructure. For S-BGP, one PKI with two certification hierarchies is necesssary; the original paper thus describes two PKIs.

doi: 10.2313/NET-2011-07-2_10

recipient can return the correct "answer" to the tags and thus confirm that the packet reached the correct router. With an extension, entire paths can be verified.

4.5 Defensive Filtering

Defensive filtering is not actually a novel security protocol but more of a best practice that can also be used on top of other security proposals. It describes filtering of route announcements that, according to predefined rules or heuristics, are estimated to be invalid or malicious. Defensive Filtering is particularly interesting in the case of stubs. As mentioned before, stubs are ASes without any customers. This means that they can only legitimately announce prefixes they themselves own – according to the assumed BGP policies from section 3.2, they can not sensibly serve as transit networks for other prefixes. Thus, providers of stub ASes should keep a list of prefixes owned by their connected stubs and discard any announcements for other prefixes, thereby greatly diminishing or even eliminating the potential damage attacks or misconfiguration by a stub could cause to other networks.

5. METHODOLOGY

Before we turn to the quantitative analysis of the effects various attacks have on the aforementioned security proposals, a short introduction of assumptions and methodology is necessary.

5.1 Threat model, data set, quantification

Goldberg et al. chose traffic attraction and traffic interception attacks for their analysis. While other attacks surely are relevant in today's Internet, it will become apparent that resilience to those two attacks is a critical aspect of inter-AS routing security proposals and serves well as a test case. Traffic attraction denotes the scenario where an AS tries to attract traffic destined for a prefix it does not actually own, usually trying to maximize the number of ASes that route through the attacker. This can be motivated by a number of reasons: performing a DoS attack on the prefix by dropping the attracted traffic (routing blackhole), modifying or examining traffic (interception) and, again, non-technical goals such as increasing revenue or causing financial damage by "forcing" traffic through paths the affected parties would rather avoid. Interception requires, on top of attraction, that intercepted data eventually reaches its correct destination.
Goldberg et al. ran their attack simulations on internetwork models based on data from CAIDA[8], who offer an inter-AS graph from inferred AS business relationships and available BGP peering data. All attacks they ran could have been performed just as well on the corresponding ASes in the real Internet, provided the CAIDA model was accurate enough in those cases. Success of attacks was measured by running attacks on multiple, random pairs of attacking ASes and victim ASs, measuring the fraction of ASes whose traffic the attacker managed to attract and computing the distributions of these fractions.
The authors tried to assume the worst case, attacking each protocol with the worst possible attack, i.e., the optimal strategy for the attacker.

[8]Cooperative Association for Internet Data Analysis, `http://www.caida.org/home/`

5.2 Underlying assumptions, caveats

Goldberg et al. made several choices that understate the effect of their attacks while at the same time making reasonable assumptions on aspects that might benefit attacks, such as assuming that ASes announce all paths except those "forbidden" by the policies stated in section 3.2. They also assume a static AS graph, which is certainly not true for the real Internet, but probably justified by their argument that AS graph changes occur on a much longer timescale than BGP execution.
A significant caveat is their assumption that no monitoring services are used for defense against attacks. Such services, e.g. offered by Renesys and RIPE (RIS), monitor inter-AS routing with a large number of probes placed at various points in the Internet and make BGP peering data available publicly or to their customers[9]. Users of such services can spot suspicious local changes in their routing information or use the data to search for larger anomalies in inter-AS routing. Also, for some attacks, Goldberg et al. grant some knowledge of global routing configuration to the attacker, justifying this with the assumption that the attacker acts strategically and with preparation. Important is also the fact that only single attacking ASes were considered – colluding ASes have interesting attack options as well, such as tunneling route announcements between each other that then offer shorter, bogus, paths [7]. S-BGP can not prevent this attack if the routers sign each other's paths.

6. ATTRACTION ATTACKS

The strategy for the first set of attacks, traffic attraction attacks, is as follows: announce the shortest possible paths that are allowed by the respective security protocol to all BGP peers to attract traffic, disregarding the routing policies we are assuming for honest ASes. That means, for plain BGP the attacker would announce the victim prefix as his own, *originating* it. In case of origin authentication, the attacker will announce a *direct link to the owner* of the prefix and soBGP requires at least a *physically existing path*. For S-BGP, the attacker has to choose the *shortest path to the victim that is actually available* to him. As Goldberg et al. point out for the case of S-BGP, if the attacker decides to actually forward traffic on the path he could already announce without S-BGP raising an alarm, the attack is not detected by data plane verification either.
Figure 2 shows the probability an attacker can attract at least 10% of ASes in the internetwork with his announcements. See Figure 3 for a more detailed plot, showing the cumulative probability for some fraction of ASes routing through the attacker. Note the high probabilities of success for this relatively unsophisticated attack strategy, especially considering that these are lower bounds – Goldberg et al. even prove that finding the optimal attack strategy is NP-hard.

6.1 Findings

Goldberg et al. draw several conclusions from the results above. This paper concentrates on two significant and simple findings; for a full list with several intriguing findings see the full version of the source paper [6].

[9]Customers being regular business customers in this case, not traffic customers

doi: 10.2313/NET-2011-07-2_10

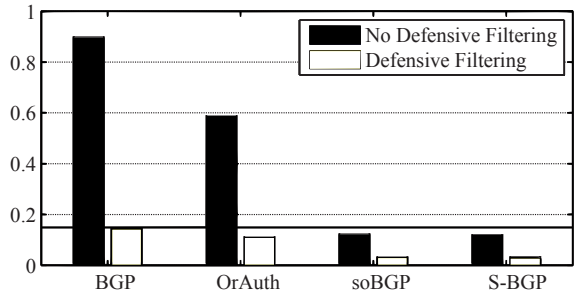

Figure 2: Lower bounds on the probability of attracting at least 10% of ASes in the internetwork [5]

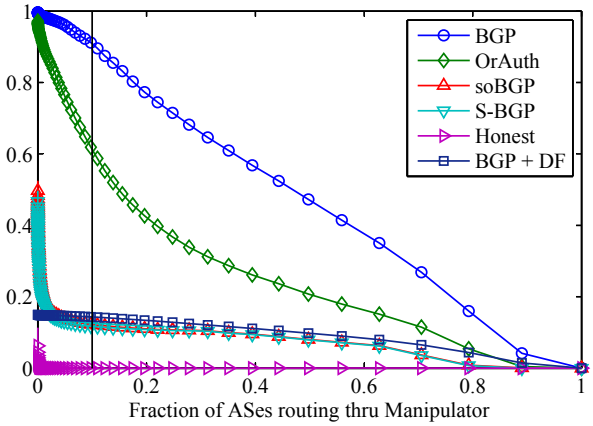

Figure 3: CCDF for the "Shortest-Path Export-All" attack strategy [5]

6.1.1 *Defensive filtering*

The first result that is quite striking and one of the most significant findings of the paper is apparent in Figure 2. The plot shows the large influence defensive filtering of stub announcements has in preventing attacks. Defensive filtering combined with plain BGP works almost as well as S-BGP alone – without requiring any changes to routing protocols, PKIs or other computationally intensive cryptography. This result will reappear when we discuss other attack strategies.

6.1.2 *Export policies*

The only minor difference between soBGP and S-BGP serves as a hint to another important finding. While S-BGP does restrict possible paths the attacker can announce, and thus forces the attacker to announce longer paths compared to, e.g., soBGP, this does not make the attack much less efficient. Goldberg et al. show that this is just a side effect of a very important point – path lengths are often less relevant for a route's attractiveness than export policies. This is easily understood by considering the case where an attacker ignores his policy of not incurring financial loss and announces provider paths to his provider. A provider will likely, according to the BGP route selection process and policies, prefer a customer route before even considering path lengths!
Because route announcements are not binding, with the exception of data plane verification, an attacker can use the announcement of a path that is attractive to other ASs but costly for the attacker, but then forward attracted traffic on a cheap or free path, if at all.

6.1.3 *Tier 2 attackers*

A somewhat surprising result is that the most efficient attackers for traffic attraction are ASes located in Tier 2. While Tier 1 is often still viewed as the "backbone" or "core" of the Internet[10], with short path lengths to most destinations, path length is trumped by policy considerations once again. Tier 1 networks are always providers or peers, never customers. This makes them less attractive for all lower tiers, as those would usually have to pay for forwarding traffic to a Tier 1 or occupy peering capacities. Tier 2 networks provide an ideal combination of good connectivity and attractive customer links. For the same reason, Tier 1 ASes are more vulnerable to traffic attraction attacks than Tier 2's – ASes that want to reach a Tier 1 can only be customers or peers of their destination and as such are more likely to accept alternative paths introduced by an attacker which are cheaper or even earn them money, in case of customer paths.

7. INTERCEPTION ATTACKS

Like attraction attacks, interception attacks aim at attracting as much traffic as possible, but also at preserving a path to the victim on which the intercepted traffic is ultimately delivered. The attacker typically wants to snoop traffic or modify it, ideally without the victim noticing anything out of the ordinary. This means that the attacker must not cause routing blackholes, which happen when the attacker attracts traffic meant for his victim but has no available route to the victim – typically, because he attracts the traffic from his providers to his victim as well. Interestingly, Goldberg et al. provide proof that in many scenarios, blackholes are impossible: see Table 1.
An attacker who wants to preserve a customer path to a victim can announce any path to any neighbor type, while there are counterexamples that show that for example peer paths can not always be preserved if an attacker indiscriminately announces paths to providers. This makes attackers in Tier 1 ideal interceptors – they do not have provider paths, and thus do not have to worry as much about introducing routing blackholes as lower-Tier-ASes.

	May announce to		
Preserve path of type	Customers	Peers	Providers
Customer	✓	✓	✓
Peer	✓	✓	×
Provider	✓	×	×

Table 1: Blackhole prevention [5]

7.1 Three different strategies

The first strategy for interception is, like in section 6, shortest path export all – for each security protocol, announce the shortest possible paths to all neighboring BGP routers. Attacks with this strategy on less secure systems such as BGP are more likely to cause blackholes compared to, e.g., S-BGP

[10] A notion that has been outdated for some time now, actually, since before the introduction of BGP

doi: 10.2313/NET-2011-07-2_10

because S-BGP forces the attacker to announce an available path – which can not be a blackhole. This implies an easy way to circumvent the problem of blackholes: instead of announcing shortest paths, announce shortest available paths not only in case of S-BGP. While this prevents blackholes, this strategy appears to be less-than-ideal in internetworks without S-BGP. A hybrid strategy of using shortest path export all per default, checking if a path to the victim is still available, and switching to shortest available path export all if necessary seems like a sound strategy.

7.2 Results

Results for interception attacks on plain BGP are plotted in Figure 4. Goldberg et al. provide no results for these interception attacks on any of the security proposals. For plain BGP, the attacks are obviously quite successful. Results are likely to be similar or identical for the security proposals, as shortest available path export all will circumvent all proposals up to and including S-BGP.

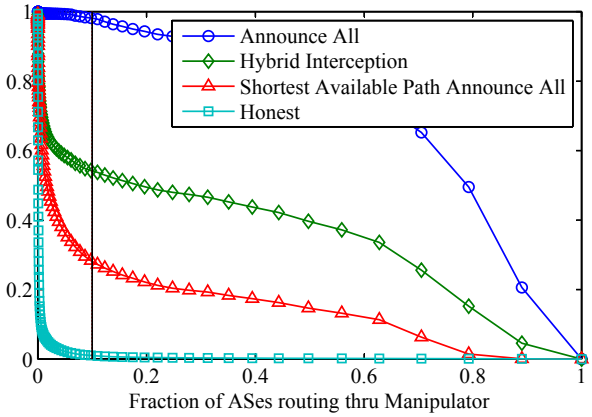

Figure 4: Interception attacks on (plain) BGP [5]

8. COUNTERINTUITIVE ATTACKS

Attacks on inter-AS routing are not always obvious, and understanding attacks is made more complicated by the heavy influence of non-technical considerations. Goldberg et al. found three interesting AS subgraphs in their data set for which they demonstrated very counterintuitive attacks that were astonishingly successful in their simulations; demonstrating that shortest path export all is not optimal for attackers. Figures used in this section show the amount of providers etc. for some ASes; these are in plain text next to the AS in the graph. Colored numbers in triangles state the number of customer ASes which route through the attacker via the AS the triangle's arrow points to.

8.1 Announcing longer paths

For this example, we assume that soBGP, S-BGP or data plane verification is implemented in the internetwork. Figure 5 shows the AS subgraph this attack will be run on. On top, the green arrows indicate a scenario where the attacker m intercepts traffic to v from $a2$ and $a3$ by using the shortest path export all strategy by announcing the path $m \rightarrow a1 \rightarrow v \rightarrow prefix$. Including $a3$'s customers, this attack manages to attract 2546 ASes. The attacker can do even better, though. If m announces $m \rightarrow a2 \rightarrow a3 \rightarrow v \rightarrow prefix$, this longer path will actually be preferred by m's provider $a1$ over its own *direct peering link* to v! Because in this specific case $a1$ has considerably more customers than $a2$, the attacker increases attracted traffic – *threefold*, as shown in the lower part of Figure 5! Note that because $p1$ and $p2$ are now using customer links to reach v instead of their peering links, they are in principle willing to announce this path *to anyone*. To avoid this attack scenario, one would probably have to implement checks that ASes follow standard path export policy – m is not announcing false paths, claiming ownership of prefixes it does not own or announcing one path but forwarding on another, thereby circumventing all security proposals up to and including data plane verification. The sole exception are stub attackers when defensive filtering is in place.

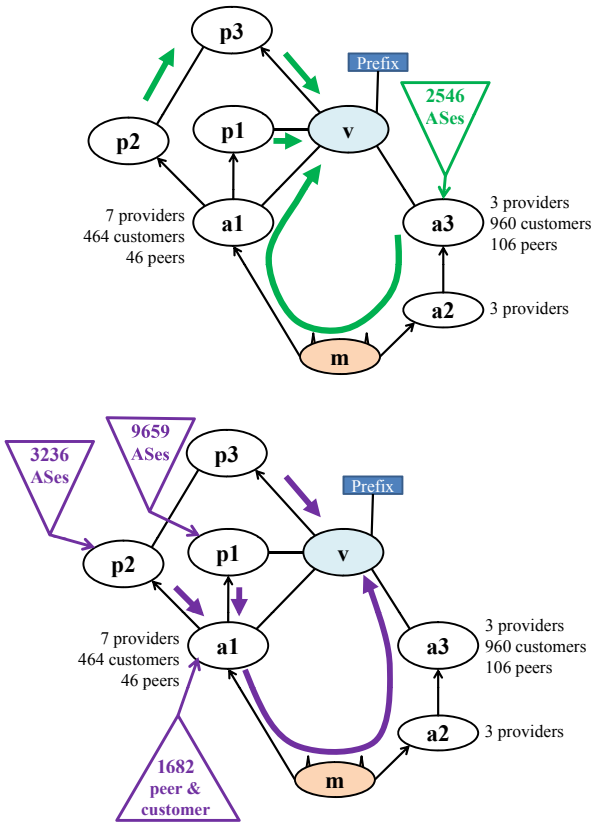

Figure 5: Announcing a longer path [5]

8.2 Exporting less

Figure 6 again shows shortest path export all in green: m announces $m \rightarrow v \rightarrow prefix$ to Tier 2 provider $T2$ and both $T1a$ and $T1b$ choose their customer link to $T2$ for reaching v: $T2 \rightarrow m \rightarrow v \rightarrow prefix$. If m stops this announcement, $T2$ has to use the peering link $T1c$ and, following policy guidelines, stops propagating his route to v to his providers $T1a$ and $T1b$. $T1a$ and $T1b$ now have to use their peering links with m to reach v. So far, nothing seems to have been accomplished by v; actually, traffic form $T2$ is now no longer attracted. What makes this attack superior in this case is the fact that the Tier 1 networks now announce shorter paths to v to their customers, attracting more traffic. For this specific

doi: 10.2313/NET-2011-07-2_10

case, traffic attraction could be increased *fourfold*. So, by *forcing* Tier 1 ASs, which have a large number of customers, to use shorter paths, the attacker massively increases the attracted traffic. This attack, just like the previous one, requires no overtly malicious activity – only strategic route export policies. It works in presence of all security protocols discussed here.

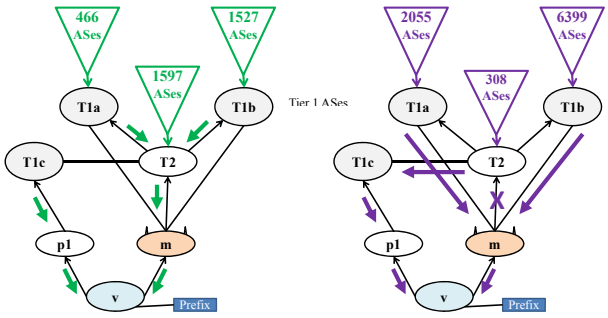

Figure 6: Exporting less [5]

8.3 False loops

The last attack described here aims at creating a black-hole. On the left in Figure 7, the attacker chooses the very aggressive attack strategy of originating the prefix that rightfully belongs to v. $T1a$ will choose the route $a3 \rightarrow a2 \rightarrow a1 \rightarrow m \rightarrow prefix$ because it is a customer path. In this dataset, Goldberg et al. showed that 32010 ASes could be attracted this way, which is the majority of ASes in that dataset. Now the attacker aims at something similar to the strategy in section 8.2: shortening the path of which $T1a$ thinks that it leads to v through m. In this AS subgraph, m can achieve this by announcing $m \rightarrow a2 \rightarrow prefix$ to $a1$, which will forward its customer's route to $T1a$ and $a2$. At $a2$, BGP loop detection will reject this path as invalid. $T1a$ thus loses its path over $a2$ and starts using the manipulated peering path $a1 \rightarrow m \rightarrow a2 \rightarrow prefix$, drawing more traffic into the trap set up by m; 32370 ASes in this case. This slight increase is due to the increased attractiveness of the path, which is now shorter[11]. S-BGP catches this attack because it recognizes the illegal paths announced by m.

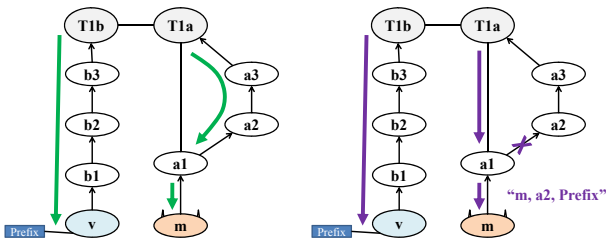

Figure 7: False loops [5]

[11]Actually, the situation is slightly more complicated, see [6] for a detailed description. The reason for the increased effectiveness of the attack is the same.

9. RELATED WORK

The security proposal SPV [7] was not considered by Goldberg et al.; except for origin authentication and use of IPSec, it provides similar guarantees as S-BGP. However, Butler et al. find its reliance on probabilistic arguments in some cases too problematic and refer to Raghavan et al. [13], who found that a majority of ASes can forge routes in SPV with high probability.

Another surprisingly multifaceted, but not very high-profile, attack on inter-AS routing that was not discussed by Goldberg et al. is link cutting [2].

Some of the proposals described here are already under way, an example being a PKI for origin authentication [8].

10. CONCLUSION

This paper described quantitative comparisons by Goldberg et al. of four inter-AS routing security proposals, which show that even quite sophisticated and seemingly secure proposals can still be circumvented by surprisingly easy attacks. Especially two findings are important: first, traffic attraction attacks *can* be mitigated. For example, defensive filtering alone would probably significantly reduce the number of possible attraction attacks, see Figure 2. Second, strategic configuration of export policies by an attacker can easily circumvent even the most sophisticated proposals – which only makes the Internet-wide implementation of defensive filtering more important for improving inter-AS routing security.

Goldberg et al. used mostly convincing methods for their analysis. While they omitted some interesting attack and defense strategies, only focused on traffic attraction and interception and had to concede that the specific subgraphs used for their counterintuitive, but very effective, attacks were hard to find, their general findings seem sound. On the non-technical side, issues such as single points of trusts in PKIs needed for example for S-BGP were not addressed.

In conclusion, inter-AS routing remains remarkably insecure. While work is under way to improve the situation, currently, effective tools like defensive filtering are not universally used due to the fact that providers do not directly benefit from its implementation *on their own network*. Sophisticated security schemes in development might require major overhaul of Internet routing architecture and significantly increase resource use while still failing to address relatively simple attacks. Unfortunately, it seems that apart from using route monitoring services and implementing best practices such as defensive filtering, there is not much an AS operator can do to improve BGP security today – except to wait for the rest of the Internet to follow suit with implementing best practices.

doi: 10.2313/NET-2011-07-2_10

11. REFERENCES

[1] W. Aiello, J. Ioannidis, and P. McDaniel. Origin Authentication in Interdomain Routing. In *Proceedings of the 10th ACM conference on Computer and Communications Security*, CCS '03, pages 165–178, New York, NY, USA, 2003. ACM.

[2] S. M. Bellovin and E. R. Gansner. Using Link Cuts to Attack Internet Routing. Technical report, ATT Research, 2003.

[3] V. J. Bono. 7007 Explanation and Apology. `http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html`.

[4] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, 98(1):100–122, 2010.

[5] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How Secure are Secure Interdomain Routing Protocols? *SIGCOMM Comput. Commun. Rev.*, 40:87–98, August 2010.

[6] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How Secure are Secure Interdomain Routing Protocols? Technical Report MSR-TR-2010-18, Microsoft Research, 2010.

[7] Y.-C. Hu, A. Perrig, and M. Sirbu. SPV: Secure Path Vector Routing for Securing BGP. *SIGCOMM Comput. Commun. Rev.*, 34:179–192, August 2004.

[8] IETF. Secure Inter-Domain Routing Working Group, 2011. `http://datatracker.ietf.org/wg/sidr/charter`.

[9] S. Kent, C. Lynn, J. Mikkelson, and K. Seo. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18:103–116, 2000.

[10] C. Labovitz. China Hijacks 15% of Internet Traffic? `http://asert.arbornetworks.com/2010/11/china-hijacks-15-of-internet-traffic/`.

[11] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP misconfiguration. *SIGCOMM Comput. Commun. Rev.*, 32:3–16, August 2002.

[12] D. McPherson. Internet Routing Insecurity::Pakistan Nukes YouTube? `http://asert.arbornetworks.com/2008/02/internet-routing-insecuritypakistan-nukes-youtube/`.

[13] B. Raghavan, S. Panjwani, and A. Mityagin. Analysis of the SPV Secure Routing Protocol: Weaknesses and Lessons. *SIGCOMM Comput. Commun. Rev.*, 37:29–38, March 2007.

[14] RIPE. YouTube Hijacking: A RIPE NCC RIS case study. `http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study`.

[15] R. White. Architecture and Deployment Considerations for Secure Origin BGP (soBGP). `ftp://ftp-eng.cisco.com/sobgp/drafts/draft-white-sobgp-architecture-01a.txt`.

[16] R. White. Securing BGP Through Secure Origin BGP. `http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_6-3/securing_bgp_sobgp.html`.

[17] E. L. Wong, P. Balasubramanian, L. Alvisi, M. G. Gouda, and V. Shmatikov. Truth In Advertising: Lightweight Verification of Route Integrity. In *Proceedings of the twenty-sixth annual ACM symposium on Principles of Distributed Computing*, PODC '07, pages 147–156, New York, NY, USA, 2007. ACM.

doi: 10.2313/NET-2011-07-2_10