

A BGP ATTACK AGAINST TRAFFIC ENGINEERING

Jintae Kim
Steven Y. Ko
David M. Nicol

Department of Computer Science and
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801, U.S.A.

Xenofontas A. Dimitropoulos
George F. Riley

School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332, U.S.A.

ABSTRACT

As the Internet grows, traffic engineering has become a widely-used technique to control the flow of packets. For the inter-domain routing, traffic engineering relies on configurations of the Border Gateway Protocol (BGP). While it is recognized that the misconfiguration of BGP can cause negative effects on the Internet, we consider attack methods that disable traffic engineering regardless of the correctness of configurations. We focus on the redirection of traffic as our attack objective, and present attack scenarios on some dominant sample network topologies to achieve this objective. We also evaluate and validate these attacks using two different discrete-event simulators, one that models BGP behavior on a network, and another that emulates it using direct-execution of working BGP code.

1 INTRODUCTION

The Internet is a very large-scale decentralized network. A packet sent from one computer to another may cross multiple networking administrative domains, so-called *Autonomous Systems* (AS). Internet routing infrastructure includes two different systems - an intra-domain routing system which operates within an Autonomous System and an inter-domain routing system, which provides connectivity between ASes. An AS routes traffic within its infrastructure any way it likes; all of the Internet's ASes coordinate the inter-domain routing by running the Border Gateway Protocol (BGP) protocol (van Beijnum 2002) on routers that connect ASes. BGP computes *routes* (i.e., sequences of ASes to visit) between every AS and every IP address.

Large ASes associated with telecommunications companies make a business of carrying inter-domain Internet traffic. Many ASes do not, typically those associated with a comparatively small organization (like a University, or large non-communication's company). Such ASes connect to the Internet through another AS known as a provider.

Some actually connect through multiple providers, a practice known as *multi-homing*. Through each such provider a multi-homed AS causes the creation of a different connection between it and every Internet destination. As a result, each multi-homed AS can elect to use whichever of these routes looks most attractive, based on its own definition of "best route" such as shortest hops, shortest latency, and etc. Activities that decide the best route of the AS's own definition and control the traffic to take certain routes are called *traffic engineering*. Traffic engineering can be done in many ways, but for the inter-domain routings, it solely relies on the configuration of BGP. Several policy-based inter-domain routing techniques have been developed to direct the traffic to a certain path that a network administrator specifies. However, slight mis-configurations or attacks that are designed to exploit vulnerabilities of BGP can cause different types of damage to the Internet. In fact, one of the biggest challenges faced by inter-domain routing technology is that it is vulnerable to mis-configurations and attacks that cause abnormal packet flows. It has been shown that mis-configurations often happen in practice, and that they can cause reachability or BGP convergence problems. For example, a well-known incident involving AS 7007 caused a crash of backbone networks throughout North America and Europe for several hours, by announcing that AS 7007 itself is the origin of the best path to almost the whole Internet (Bono 1997). Previous research also indicates that BGP is vulnerable to attacks, and thus can be exploited to compromise the inter-domain routing systems. The result of such attacks may be more severe than those of mis-configurations because such attacks will intentionally affect a larger number of hosts.

In this paper, we focus especially on vulnerabilities of traffic engineering under BGP, by exploring how attackers may be able to disable or disrupt inter-domain traffic engineering. The objective of this paper is to show that multiple colluding attackers can redirect the victim's traffic (both incoming and outgoing) as they like, regardless of

victim’s policy settings. By studying the actual Internet AS topology, we also show that local topologies vulnerable to these attacks are pervasive.

The rest of this paper organized as follows. We start by providing background information in Section 2. Section 3 describes our attack method. We analyze the BGP routing table and present the result of analysis in Section 4. Section 5 uses two types of simulators to demonstrate the effect of these attacks. One (using SSFNet) models BGP, and analyzes the impact on throughput that an attack may have. The other (using BGP++) embeds real BGP code into a simulated network environment. Playing the attacks in BGP++ with subsequent examination of the forwarding tables validates the possibility of these attacks. We conclude in Section 6.

2 Related Work

BGP attacks have recently started to gain attention from the network research community. Nordstrom and Dovrolis (2004) stated that BGP is vulnerable to a number of relatively easy attacks once one or more BGP speakers are compromised. It classified the attack objectives into the following four categories: black-holing, redirection, subversion and instability. It also proposed several attack methods in a concrete way, and extensively handled the major proposed countermeasures for BGP security.

Routing Protocol Security (rpsec) has also developed BGP attack trees in its Internet draft (Convery, Cook, and Franz 2004). They defined a few atomic goals to be achieved throughout the attacks, such as compromising MD5 authentication, establishing unauthorized BGP session with peer, originating unauthorized prefix into peer route table, etc. Furthermore, they constructed an attack tree for each atomic goal, so that each tree could describe comprehensive ways to achieve the goal. Although they tried to show that many different attacks exist, they did not discuss each attack in depth.

In addition to intentional attacks, Mahajan, Wetherall, and Anderson (2002) studied the effect of misconfigurations of BGP. They focused on two types of mis-configurations: origin mis-configurations (where an AS accidentally injects a prefix into the global BGP table) and export mis-configurations (where an AS exports a route in violation of its policies). By examining the streams of real BGP updates, they showed that these configuration errors are pervasive, and that they substantially increase the update load on routers.

3 BACKGROUND

3.1 Overview of BGP4 - The Border Gateway Protocol

BGP is the *de facto* standard protocol for the inter-domain routing. It specifies how each packet from one AS finds its way to another AS. A unit of “domain” in BGP is an IP prefix, each prefix is bound to some AS (and if the AS intends for other ASes to be able to address to it, the AS “announces” the prefix using BGP). Each AS has one or more external routers called BGP *speakers* that execute some implementation of the BGP specifications. When a packet exits an AS through a BGP speaker, that speaker looks at the packet’s destination address, and looks up in a *forwarding table* the port through which it needs to push the packet along. In principle, every BGP speaker must be able to forward a packet addressed to any IP prefix announced anywhere in the Internet. Toward that goal, execution of the BGP protocol builds a speaker’s forwarding tables. An entry in the forwarding table contains a complete route to a destination prefix. Given a packet’s IP address, a router searches for the matching prefix and forwards the packet accordingly. The forwarding table of a BGP speaker inside an AS is constructed by receiving UPDATE messages from BGP speakers of other ASes. Each router can announce a new route or withdraw an existing route through sending an UPDATE message. KEEPALIVE messages are exchanged between BGP speakers to make sure they are reachable. A certain timeout value is used to send these messages.

Traffic engineering in BGP protocol takes place when a BGP speaker receives two or more routes to a same destination IP prefix. A BGP speaker chooses one route based on its configuration. We will discuss traffic engineering further in Section 3.1.2 and Section 4.1

3.1.1 Peering Relationships

Two ASes are said to have a peering relationship when each of them has a BGP speaker, and the two BGP speakers communicate with each other. A peering relationship can be classified into one of the following three categories (Gao 2000).

- customer/provider : AS u is a provider of AS v iff u transits traffic for v (a customer) and v does not transit traffic for u .
- peer/peer : AS u and v have a peering relationship iff neither u nor v transits traffic for each other.
- sibling/sibling : AS u and v have a sibling relationship iff both u and v transit traffic for each other.

The customer/provider relationship is a commercial one based on inequality of size and capability. The sib-

ling/sibling relationship is one based on equality—one sibling agrees to carry traffic it has no direct interest in, received from another sibling. It is a sort of symmetric customer/provider relationship, but without one sibling paying another. The peer/peer relationship is one of convenience. One peer announces its own prefixes to another peer, in order to allow the second peer to directly send traffic it originates to the first.

3.1.2 Path Selection Mechanisms

When a BGP speaker receives two or more announcements describing how to reach a given IP prefix, it has to choose one among them. For this purpose, BGP has a path selection mechanism. That mechanism is an ordered sequence of priority comparisons. The first comparison in the sequence that indicates a preference for one announcement over another ends up defining the choice. The following sequence of comparison keys are from a CISCO router, but routers from other vendors take similar steps. This list is actually incomplete; the first four steps, shown below, are the only ones relevant to the attacks we consider.

1. The path with the largest “weight” (a Cisco-specific parameter, local to the router on which it’s configured.)
2. The path with the largest local preference. A local preference indicates the preferred path. Unlike the weight attribute which is only relevant to the local router, the local preference is shared among the external routers in a same AS. It is important to remember that a local preference is defined by the AS making the selection; typically the local preference is established at configuration, an announcement’s local preference is the local preference value assigned to the AS that announced it.
3. The path that was locally originated.
4. The path with the shortest number of hops (i.e. the shortest AS path).

4 ATTACK MODEL

We start this section by explaining traffic engineering. Afterwards, we describe the goal and mechanisms of our attacks. Note that we assume that an attacker has compromised and taken control of one or more BGP routers. A BGP router is just another device on the Internet, accessible like another device on the Internet, albeit with an operating system different than most. Nevertheless attackers can gain access to the BGP machines through mechanisms like password spoofing, brute force password scanning or default password detection, because a large number of routers simply use the default passwords. Many routers also use widely-

used interfaces like telnet and SSH, so they share all known vulnerabilities of the interfaces (Subramanian et al. 2004).

4.1 Traffic Engineering

This subsection gives an overview of traffic engineering in practice, mostly described in (van Beijnum 2002, Quoitin et al. 2003).

4.1.1 Incoming Traffic

A network administrator can see how much traffic his routes attract over the different connections to the global Internet. While the traffic may be distributed equally over existing connections (depending on which prefixes are announced through them), a problem can occur if one connection attracts more traffic than it can handle or the connection capacities are significantly different. In this case, an administrator can encourage upstream ASes to select a different connection, by re-announcing prefixes over the first connection but with the AS identity prepended to the path more than once. For upstream ASes where the two paths to the prefix are equivalent in preference using comparisons 1-3 above, comparison 4 (shortest path) may now settle the issue in favor of the targeted connection. This technique is known as AS path padding.

4.1.2 Outgoing Traffic

The most effective way to influence the BGP path selection process is to adjust the *local preference*. This can be very effective when a certain route is always better than the other routes. For example, if routes over the main connection are preferred over routes that use a slower backup connection, higher local preference can be assigned to routes received from a BGP neighbor linked with the main connection.

Since local preference is defined by the AS that uses it, attacks on local preference must either be on that AS itself, or to cause elimination of connections to ASes having more advantageous local preference than the connection the attack wishes to have used.

4.2 Attack Description

The objective of our attacks is to redirect the victim AS’s traffic load to the path specified by the attackers, regardless of the victim’s BGP configuration. Since an AS normally prefers the high-bandwidth connection, redirecting the traffic to the low-bandwidth connection can decrease the network performance (decreasing the capability of the victim AS). It may also be that an AS multi-homes to provide itself an expensive but infrequently used backup connection, so an attack that forces an AS to use the expensive connection and suffer economically. Another motivation for redirecting

traffic is to allow one to eavesdrop or modify the messages sent to the victim; the packets are still forward to their original destination, making the attack difficult to detect.

4.2.1 Incoming Traffic

Assume that an attacker has compromised a neighboring AS of the victim. If the attacker wishes to intercept traffic destined for the victim, but finds that traffic to the victim is routed through a different AS, it can immediately announce that it is the origin of the best path to the victim (one way of which is described below). Then the ASes connected to the attacker will send their traffic through the attacker, and the usage of the link between the attacker and the victim will increase.

Figure 1 shows the attack with a simple example topology. AS 1 is the gateway AS that provides outside connections to the attacker, AS 2, and the victim. Assume that the victim wants to force all incoming traffic to pass through the normal link from AS 2; as we describe in Section 4.1.1, this is usually accomplished by the victim padding multiple copies of its own id onto the path it announces. However, if the attacker wants to force all incoming traffic to go through its link (backup link in Figure 1), then it can engineer an announcement that should be more attractive to AS 1 than what AS 1 receives from AS 2, e.g., drop the padding on the victim's own announcement, or state that the attacker is the victim itself. Assuming that comparison 1 (weight) is equivalent, we see that comparison 2 (local preference) might still cause AS 1 to choose to use AS 2. While this is certainly possible, the customary use of local preference is by a customer to select between providers, which is not the case here. Therefore if comparison 2 does not resolve the route selection, then the false announcements we named will use rule 3 to resolve in favor of the false announcement. Moreover, deceived ASes further propagate the false announcements to other ASes.

This attack relies upon use of the shortest-path rule to allow false announcements by the attacker to attract inbound traffic for the victim. By this principle more complex topologies might also be attacked. For example, if the victim in Figure 1 homes with even more ASes that peer only with provider AS 1, then as the attacker impersonates the victim (in terms of AS identity and prefixes announced), it will again intercept all of the victim's traffic.

Another type of attack can intercept some (if not all) of a victim's inbound traffic, but in much more general topologies. This is illustrated in Figure 2. The stub AS 8 that has a single upstream provider, AS 7, is the victim. The compromised AS 4, which can reach the victim AS 8 along the path 5 7 8, sends an update to AS 1 announcing that it originates one or more prefixes of the victim. AS 1 that receives the false update, runs the BGP decision process and selects the compromised AS 4 to route traffic

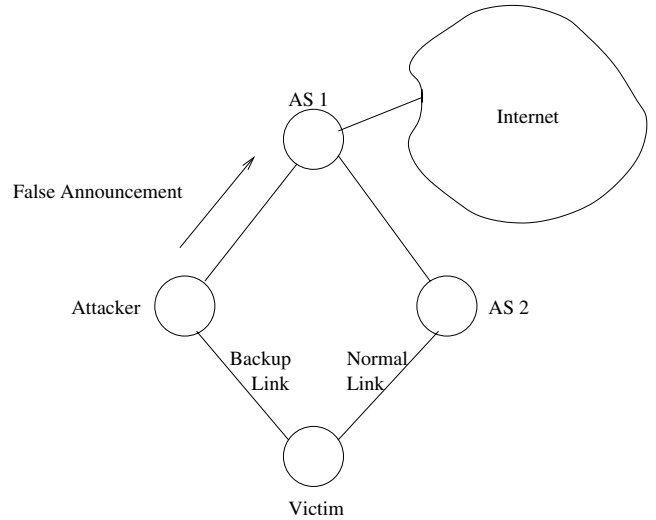


Figure 1: An Attack Scenario for Incoming Traffic

destined to AS 8. It is assumed that the new route is preferred over the original route because of the shorter AS-path. Subsequently, AS 1 further propagates the false announcement resulting in more ASes shifting to the new path. The selection of the fake path over the original path depends on the length of the two paths and the configuration of the routers. Typically, shorter AS-paths are selected, although in some special cases administrator configuration may impose higher local preference for certain types of paths, e.g. paths through a cheaper neighbor, resulting in a longer path being preferred. In this example, assuming that shorter AS-paths are preferred, ASes 1, 2 and 3 switch to the new spurious path.

This example suggests a general principle under which an attacker can intercept a victim's traffic, but still forwarded that traffic to reach the victim after copying or modifying it. Consider a graph where nodes represent ASes, and an edge exists between customers and providers that peer, and between siblings that peer. This graph fairly represents possible paths, and BGP will select the shortest path, so long as the "weight" comparison doesn't matter, and local preference definitions are limited to customer ASes selecting providers.

Let A_v be an intended victim router, A_a be an attacking router, and A_t be the AS through which A_a normally routes traffic toward A_v (Figure 3). Define $d_{t,v}$ to be the shortest number of hops in the graph from A_t to A_v , and define $d_{t,a}$ to be the shortest number of hops in the graph from A_t to A_a on paths that do not use the $A_a - A_t$ link. If A_a makes announcements (to peers other than A_t) that indicate that A_a is within k hops of A_t , then all ASes whose distance to A_a , plus k , is smaller than their distance to A_v in the graph will be fooled into directing their traffic toward the impersonator. However, A_a can still cause A_v 's traffic




Figure 2: An Attack Scenario on More General Sample




Figure 3: An Attack Scenario on Interception

to be delivered simply by forwarding it to A_t , provided that A_t is not itself affected by A_a 's false announcements. A_t is immune from the false announcement provided that $d_{t,a} + k > d_{t,v}$. Parameter k here governs a sort of trade-off—the smaller it is, the larger the set of ASes that may be fooled, but the more stringent is the requirement on A_t that permits traffic to be delivered to A_v after being snooped. In practice, a compromised router would identify potential victims (and hence applicable ASes to play the role of A_t), try to determine $d_{t,v}$, choose $k = d_{t,v} - d_{t,a}$, and craft false announcements which claim that A_a is within k hops of A_v .

This attack can be viewed as a variation of the well-known man in the middle (MITM) attack, in which players are ASes and messages are intercepted in one direction instead of both directions. Furthermore, it is more powerful than the MITM attack in the sense that it can affect traffic not just between two players, e.g. Alice and Bob, but between a number of sender ASes and one receiver AS, where each of the involved ASes bears a large number of end users. The impact of the false announcements made by the compromised AS depends on the topological properties of the compromised AS and the victim AS. Intuitively, if the compromised AS is located near the core of the AS topology it will affect more ASes. Also if the victim AS is located at the periphery of the AS topology it is more vulnerable to an attack.

4.2.2 Outgoing Traffic

Figure 4 illustrates a topology we use to discuss attacks on a victim's outbound traffic. Assume that the victim wants to send its outgoing traffic through the normal link to AS 2 instead of the backup link to attacker 1, again because

of economic reasons or performance reasons. In this case, what the victim usually does is to set the local preference of the normal link larger than that of the backup link, so that all outgoing traffic would go through the normal link as in Section 4.1.2. Assume that there is an attacker who wants to force all outgoing traffic of the victim to go through the backup link. To achieve this goal, the attacker compromises two BGP speakers: one in a neighboring AS, and the other in a gateway AS. The attacker knows the AS path from the neighboring AS, so it can withdraw all the other paths from the gateway AS. Then the victim has no other choice to select the route from a neighboring attacker and will send its traffic to that route. Assume that attacker 1 has compromised attacker 2 or an attacker has compromised two BGP speakers, attacker 1 and attacker 2. The goal of the attacker is to force all outgoing traffic from the victim to go through the backup link. This can be easily achieved by sending a withdrawal message from attacker 2 to AS 2, stating that attacker 2 now no longer has any connection to the Internet. AS 2 will send a withdrawal message to the victim stating that it does not know any connection to the Internet and the victim has to use the backup link.

4.2.3 Discussion

In Section 4.2.1 and 4.2.2, we show how an attacker can completely control a victim's traffic, on a small topology. We explained how a slightly more complex topology is still vulnerable to attacks on inbound traffic; however, the attack described on outbound traffic requires a compromised router for every link to the victim on which a withdrawal message is to be generated. While these constraints are severe, it




Figure 4: An Attack Scenario for Outgoing Traffic

turns out that these sub-topologies appear quite often in the Internet, as we next see.

4.3 Topology Analysis

Since the attack scheme that we described in Section 4.2 is topology-dependent, it is important to analyze the Internet topology. However, it is impractical to analyze the whole Internet topology at the level of individual devices. However there are fewer than 20,000 ASes, and so a study of AS topology is quite feasible. A historical record of the Internet’s AS topology is freely available on the Internet, through the Route Views project (University of Oregon 2001). We study the topology sampled on Nov. 19, 2003. This topology has 16362 nodes, including 13752 “pure customers” and 13 “pure providers”. By “pure customer”, we mean that an AS which is only a customer to one or more ASes, but not a provider, a peer, nor a sibling to other ASes. Similarly, by “pure provider”, we mean that an AS which is only a provider to one or more ASes, but not a customer of any AS. A pure customer can have multiple providers; in our topology, most of them are single- or multi-homed ASes (a single-homed AS has only one provider and a multi-homed AS has two providers). Specifically, 5573 pure customers are single-homed and 6603 pure customers are multi-homed. Roughly speaking, about a half of all pure customers are multi-homed. In our analysis, we focus on multi-homed ASes because they are the ones in the example topology in Figure 1 and Figure 4.

In the topology that we use, there are 9810 occurrences of the example topology. To get this result, we first take each

multi-homed pure customer, and see if the two providers of the pure customer (there are exactly two providers) have a common higher provider. For the sake of our discussion, we call pure customers of the example topology as victims, two providers of a victim as intermediate nodes, and common higher providers as gateways since they provide a connection to the rest of the Internet. Note that there is a possibility that the two intermediate nodes have more than one common gateways and that different victims share a common pair of intermediate nodes. Thus, this result does not indicate that there are 9810 pure customers that are potential victims (actually, this does not make sense because there are only 6603 pure customers in the topology). In fact, for a given pair of intermediate nodes, we can identify 2.5330 common gateways and 2.3557 common victims on average. An implication of this information is that if an attacker could compromise routers in an intermediate node and a gateway, 2.3557 ASes on average would be under the influence of the attacker.

The number of potential victims are 3854 out of 6603 multi-homed pure customers. To get this result, we count all victims in 9810 occurrences of the example topology only once. However, it is important to consider the number of outgoing connections of the intermediate nodes together with the number of potential victims in order to see if the attack scheme actually has an impact. As we discuss in Section 4.2.3, the attack has much weaker or no impact if there are extra outgoing connections to the rest of the Internet among the intermediate nodes. Thus we count the number of “extra” outgoing connections of each intermediate node excluding connections through the common gateways. Figure 6 shows the distribution. As the distribution shows, it is heavy-tailed; most of the intermediate nodes have few extra connections. Specifically, 111 intermediate nodes out of 916 intermediate nodes in total have no extra connection to the rest of the Internet. 224 intermediate nodes have only one extra connection and 144 intermediate nodes have two extra connections. 63.3% of all intermediate nodes have less than 3 extra connections and 90.72% of all intermediate nodes have less than 10 extra connections. This result tells us that if an attacker who already compromised an intermediate node and a gateway only needs to compromise a few more routers to attack a victim effectively. Figure 5 shows how this attack can be accomplished. It gives us a more comprehensive view of the attack scheme. In the figure, two common gateways and victims exist for a given pair of intermediate nodes (as discussed before, there are 2.5330 common gateways and 2.3557 common victims on average). Also, intermediate node 1 has an outgoing connection to the Internet. In this case, an attacker needs to compromise 4 points to be able to attack both victims.

In summary, there are many instances of the example topology and the intermediate nodes in those instances only have few extra connections to the rest of the Internet. Thus,




Figure 5: A More Comprehensive View

an attacker needs to compromise a few more routers to achieve his goal. In addition, it is possible for an attacker to attack multiple ASes at the same time with the same set of compromised routers.

5 EXPERIMENTAL RESULTS

We compare the network performance before and after attack, so as to evaluate any effect caused by our attack methods, using simulation. Afterwards, we verify whether this attack indeed succeeds in redirecting the traffic by observing any change of routing tables.

5.1 Traffic Comparison

We evaluate the effect of the attack by comparing the incoming traffic load before and after the attack. BGP model in SSFNet (Cowie et al. 1999) on the simple example topology (Figure 1) is used as the comparison baseline. We assume that the delay and the data rate of normal links are set to 0.2s and 2Gbps, while those of backup links are set to 0.5s and 100Mbps, respectively. Figure 7 shows that the average transfer rate decreases almost linearly as the traffic passing through the gateway increases. When the gateway is responsible for 33.3% of the whole traffic, aver-

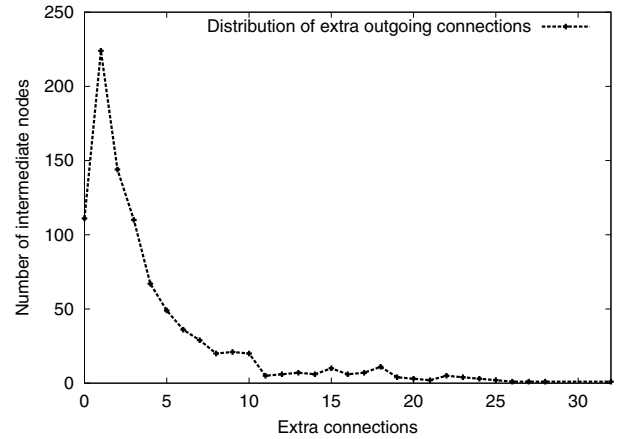


Figure 6: Distribution of Extra Outgoing Connections




Figure 7: Comparison of Incoming Traffic Load

age transfer rate decreases by 22.18%. When the gateway is responsible for 90% of the whole traffic, average transfer rate decreases by 47.63%. This result indicates that our attack can be effective when a victim is peripherally located and the providers of the victim are linked with a single large ISP network. We showed that this condition is met fairly often, by analyzing the AS-level partial topology.

5.2 Attack Verification

The attack against padding technique is reproduced in a simulation environment for verification. We use a BGP implementation for the ns-2 (McCanne and Floyd 1997) simulator, called BGP++ (Dimitropoulos and Riley 2003). For the attack, we recreate the generalized sample topology (Figure 2) as follows: ASes 1, 2 and 5 have 4 BGP routers, ASes 3, 6 and 7 have 3 BGP routers, AS 4 has 2 BGP routers and AS 8 one. Routers within an AS are connected

in a full-mesh topology, as required by RFC 1771. The delay and the data rate of all links are set to 10ms and 10Mbps, respectively. The victim AS 8, is originating the prefix 8.0.0.0/24. One of the two routers in AS 4 is considered compromised. We configure the latter to make a false announcement to AS 1, conveying that it originates the prefix 8.0.0.0/24 at the simulation second t . After running the simulation we observe that after the t second the routers of AS 1, 2 and 3 switch to using the spurious path, in which AS 4 appears as originator of the 8.0.0.0/24 prefix. This is shown in Figure 8(a) that depicts a snapshot of the routing table of a router in AS 1 before t , and Figure 8(b) that depicts a snapshot of the same routing table after t .

6 CONCLUSIONS

We have presented attack mechanisms that can redirect both incoming and outgoing traffic to the paths that the attackers specify. By compromising appropriate BGP routers, the attacker can disable the traffic engineering of a specific AS, regardless of its configuration settings. Even though the effectiveness of this attack depends on the topology, we showed that there are fairly many instances in the AS-level topology where our methods can successfully attack a victim. The experiments indicate that our attack scheme can be successfully exploited by propagating the false topology, and that it can also degrade the network performance of a victim.

This paper does not address the potential global impact that the proposed attack can generate. To examine how many ASes an attack has to compromise in order to degrade the global network performance is an interesting problem. Our future work also includes exploring the possibilities that the BGP attack can generally cause damage on the Internet. Considering that even a misconfiguration could cause a crash of large networks, we believe that malicious attacks can make more severe impact. In that sense, developing such an attack method and estimating the possible effect by large-scale simulation will remain as a valuable future work.

ACKNOWLEDGMENTS

This research was supported in part by DARPA Contracts N66001-96-C-8530 and N66002-00-1-8934, NSF Grants ANI-9977544, ANI-0136969, ANI-0240477, CCR-0209144, ECS-0225417, and Department of Energy contract DE-AC05-00OR22725. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes. In addition this project was supported under Award No. 2000-DT-CX-K001 from the Office for Domestic Preparedness, U.S. Department of Homeland Security. Points of view in this

document are those of the author(s) and do not necessarily represent the official position of the U.S. Department of Homeland Security.

REFERENCES

- Bono, V. J. 1997. 7007 explanation and apology.
- Convery, S., D. Cook, and M. Franz. 2004. An attack tree for the border gateway protocol. Internet-Draft.
- Dimitropoulos, X., and G. Riley. 2003. Creating realistic BGP models. In *Proceedings of Eleventh International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS'03)*.
- Gao, L. 2000. On inferring autonomous system relationships in the internet. In *Proceedings of IEEE Global Internet Symposium*.
- Mahajan, R., D. Wetherall, and T. Anderson. 2002. Understanding BGP misconfiguration. In *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, 3–16: ACM Press.
- McCanne, S., and S. Floyd. 1997. The LBNL network simulator. [Software on-line]. Available online via <<http://www.isi.edu/nsnam>> [accessed August 20, 2004]. Lawrence Berkeley Laboratory.
- Nordstrom, O., and C. Dovrolis. 2004. Beware of BGP attacks. *SIGCOMM Computer Communications Review* 34 (2): 1–8.
- Quoitin et al. 2003. Interdomain traffic engineering with BGP. *IEEE Communications Magazine*.
- Cowie et al. 1999. Modeling the Global Internet. [Software on-line]. Available online via <<http://ssfnet.org>> [accessed August 20, 2004]. *Computing in Science & Engineering*, 1 (1): 42–50, 1999
- Subramanian et al. 2004. Listen and whisper: Security mechanisms for BGP. In *First Symposium on Networked Systems Design and Implementation (NSDI'04)*.
- University of Oregon 2001. Route views project.
- van Beijnum, I. 2002. *BGP: Building reliable networks with the border gateway protocol*. O'Reilly.

AUTHOR BIOGRAPHIES

JINTAE KIM is a graduate student at Department of Computer Science, University of Illinois at Urbana-Champaign. He has a B.S. in computer science from Seoul National University in Korea (1999). His research interest lies in network security, specifically routing security including issues in BGP. His e-mail address is <kim28@uiuc.edu>.

STEVENY. KO is a Ph.D. candidate at Department of Computer Science, University of Illinois at Urbana-Champaign.


```
BGP table version is 0, local router ID is 0.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 8.0.0.0/24      0.5.1.1              0   5 7 8i

Total number of prefixes 1
```

(a) Routing Table Before Attack Time t

```
BGP table version is 0, local router ID is 0.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 8.0.0.0/24      0.4.1.1              0   4i

Total number of prefixes 1
```

(b) Routing Table After Attack Time t

Figure 8: Routing Table at AS 1 Before and After Attack Time t

He has a B.S. (1999) in mathematics and M.Eng. (2002) in computer engineering from Yonsei University and Seoul National University in Korea, respectively. His research interests are in networking and security. His e-mail address is <sko@cs.uiuc.edu>.

DAVID M. NICOL is a Professor of Electrical and Computer Engineering at the University of Illinois, Urbana-Champaign, and a member of the Coordinated Sciences Laboratory. He is co-author of the textbook Discrete-Event Systems Simulation, and served as Editor-in-Chief at ACM TOMACS from 1997-2003. He will serve as the General Chair of the Winter Simulation Conference in 2006. From 1996-2003 he was Professor of Computer Science at Dartmouth College, where he served as department chair, and at the Institute for Security Technology Studies served as Associate Director for Research and Development, and finally as Acting Director. From 1987-1996 he was on the faculty of the Computer Science department at the College of William and Mary; 1985-1987 he was a staff scientist at the Institute for Computer Applications in Science and Engineering. He has a B.A. in mathematics from Carleton College (1979), an M.S. (1983) and Ph.D. (1985) in computer science from the University of Virginia. His research interests are in high performance computing, performance analysis, simulation and modeling, and network security. He is a Fellow of the IEEE. His e-mail address is <nicol@crhc.uiuc.edu>.

XENOFONTAS A. DIMITROPOULOS is a Ph.D. student in the department of Electrical and Computer Engineering in Georgia Institute of Technology. He received his B.S. in physics (2001) from Aristotle University in

Greece. His research interests span the areas of computer networks and distributed systems. His e-mail address is <fontas@ece.gatech.edu>.

GEORGE F. RILEY is an Assistant Professor of Electrical and Computer Engineering at the Georgia Institute of Technology. He received his Ph.D. in computer science from the Georgia Institute of Technology, College of Computing, in August 2001. His research interests are large-scale simulation using distributed simulation methods. He is the developer of Parallel/Distributed ns2 (pdns), and the Georgia Tech Network Simulator (GTNetS). He can be reached via email at <riley@ece.gatech.edu>.