Detecting Selective Dropping Attacks in BGP

M. Chuah
Dept of CSE
Lehigh University
Bethlehem, PA 18015, USA
chuah@cse.lehigh.edu

K. Huang
Dept of CSE
Lehigh University
Bethlehem, PA 18015, USA
kuh205@cse.lehigh.edu

ABSTRACT

Previous studies have shown that current inter-domain routing protocol, Border Gateway Protocol (BGP), is vulnerable to various attacks. Initially, the major concern about BGP security is that malicious BGP routers can arbitrarily falsify BGP routing messages and spread incorrect routing information. Recently, some authors have pointed out the impact of a type of attack, namely selective dropping attack that has not studied before. The authors have shown that such an attack can result in data traffic being blackholed or trapped in a loop. However, the authors did not elaborate on how one can detect selective dropping attacks. In this paper, we present a scheme we designed to detect selective dropping attacks in BGP. We conducted extensive experiments in DETER to evaluate the effectiveness of our scheme using three 30-node AS topologies generated from Brite. Our study shows that our scheme is quite promising.

Categories and Subject Descriptors

C.2.2 [**Network Protocols**]: Routing Protocols

General Terms

Algorithms, Performance, Design.

Keywords

Internet Routing Security, Selective Dropping Attacks, BGP

1. INTRODUCTION

Border Gateway Protocol (BGP) is the de facto interdomain routing protocol [1]. Current Internet consists of many Autonomous Systems (ASes) connected by interdomain (inter-AS) links. Each AS is a set of routers that have the same routing policy within a single administrative domain. BGP is responsible for discovery and maintenance of paths between ASes in the Internet. BGP provides reachability information to ASes and distributes external reachability internally within an AS. With the exponential growth of ASes, BGP has become one of the most critical Internet infrastructures today.

BGP routers exchange routing information via UPDATE messages. UPDATE messages can be classified into two types: namely route withdrawal and route announcement. When a BGP router receives an UPDATE from its neighboring BGP router, this message will be processed stored, and redistributed in accordance with both BGP specification [1], and the routing policies of the local AS.

Previously, the major concern about BGP security is the authenticity and integrity of BGP UPDATES, especially route origin information and AS path information stored in AS_PATH attribute. Incorrect UPDATES, due to either BGP router misconfiguration or malicious attack, may cause serious problems to the global Internet. For example, on May 7, 2005, an AS falsely claimed to originate Google's prefix [2] and parts of the Internet could not reach Google's search engine for roughly an hour as traffic was misdirected to the attacking AS.

Some countermeasures have been proposed to mitigate BGP vulnerabilities. To protect BGP session from spoofed BGP UPDATES sent by outsiders, TCP MD5 signature [3] using shared secret key between two BGP routers was proposed. S-BGP[4] and SoBGP[5] apply cryptography to prevent an attacker (either insider or outsider) from advertising faulty BGP messages or tampering normal messages. However, as noted by Bellovin [6],[7] and UC Davis researchers [8], cryptography-based security mechanisms, cannot protect routing protocols against some kind of attacks. In [8], the authors describe selective dropping attacks which can cause data traffic blackhole and persistent traffic loop. However, the authors do not present any technique to detect such attacks.

In this paper, we present a scheme called Instability Analysis with Neighbor Probing (IANP) for detecting selective dropping attacks in BGP. Our IANP scheme is inspired by the Locating Internet Routing Instabilities paper [9]. In our scheme, an instability analysis will be triggered at an observation point (referred to as the monitor) when the number of received BGP UPDATES within a burst exceeds a certain threshold. When the monitor cannot identify the instability source, it will probe its neighboring routers to see if they can identify the source of instability. Once the source of instability is identified, the monitor will check its stable route database to see if a selective dropping attack is embedded within this burst of BGP UPDATES. If a monitor suspects that a BGP router is conducting a selective dropping attack, then it will issue a warning message that will be flooded (with limited scope) across the BGP routers in the Internet.

Since some monitors may not be able to locate the source of instability (e.g. not getting any UPDATES even though an instability occurs) and/or detect any router that is conducting selective dropping attacks, we define a damaged cost to quantify the impact of not being able to detect selective dropping attacks. Our damaged cost is defined as the percentage of routes that are using a broken inter-AS link as a result of the selective dropping attack launched by a malicious router.

Via extensive experiments in DETER [14] using three 30-node AS topologies generated using Brite[10], we have evaluated the effectiveness of our IANP scheme. For each network topology, we simulated a single AS link failure and collected a burst of BGP messages trigged by this failure event. Then, we simulate nodes evaluating these burst of messages using our IANP scheme. Without selective dropping attacks, the percentage of ASes that cannot locate the source of instability ranges from 0 to 20%. With a selective BGP drop attack, an additional of 3 to 13.2% of the ASes fail to locate the source of instability without using the IANP scheme. The IANP scheme allows ASes to identify the source of instability in the presence of selective dropping attack either directly or indirectly via the warning message. Without the warning messages, some ASes may still not be able to detect the selective dropping attack. For such cases, the damage cost is limited to 0.3 to 4.8% in all our experiments. With the warning messages, the damaged cost is reduced to 0.

The rest of the paper is organized as follows. Section 2 summarizes the definition of selective dropping attack and its possible damages. Section 3 describes our IANP attack detection scheme and the corrective actions that can be taken upon attack detection. We also discuss the limitation of the IANP scheme. Next, we present our experimental evaluations of IANP in Section 4. Deployment issues that warrant further investigation are discussed in Section 5. We discuss related work in Section 6. Concluding remarks are provided in Section 7.

2. Definition of Selective Dropping Attack

BGP is a policy routing protocol. According to the inbound, and outbound policies, BGP router may legitimately suppress some UPDATES. The authors in [8] defines two consistency properties for correct BGP operation which we duplicate below:

1. a) If rib(u) $\neq \varepsilon$, there must $\exists$ v $\in$ peers(u) [rib-in(u<=v)=rib(u)].

   b) If rib(u) $=\varepsilon$, rib-in(u<=v) can be arbitrary.
2. a) For any $w \in$ peers(u), if rib-out(u=>w) $\neq \varepsilon$,then rib-out(u=>w)=u $\circ$ rib(u).

   b) It is possible that when rib(u) $\neq \varepsilon$, there exists rib-out(u=>w) $= \varepsilon$ where $w \in$ peers(u)

   where the notations peer(u) denotes the set of peers for node (AS) u, rib-in (u<=w) denote node u's most recently received message from peer w, rib(u) denotes the best path that u adopts and stores in the local-RIB, rib-out(u=>w) denotes the route that u advertises to w.

The two properties listed above are legitimate properties that allow a BGP router at node u to drop BGP UPDATES. Property 1(a) implies node u can select a route from one peer but drop the routes it received from the others. Property 1(b) indicates that node u does not have to use the route announced by node v to reach a particular destination even though node u has no route. Property 2(a) guarantees that no policy allows node u to use one route but announce the other route to its peers. Property 2(b) indicates a policy to authorize node u not to transit the traffic for node v even though node u can reach a particular destination.

Any BGP dropping that are not consistent to these two properties will be classified as malicious dropping. The authors in [8] showed that such malicious dropping can result in traffic blackhole and persistent traffic loop.

3. Overview of Instability Analysis with Neighbor Probing (IANP) Scheme

In this section, we give an overview of the IANP scheme that we propose to detect and mitigate against BGP selective dropping attacks.

3.1 Instablity Analysis

In [10], the authors describe a methodology for identifying the source of instability. Figures 1 and 2 (which are Figures 7 & 8 in [9]) outline their adopted methodology. First, they group updates observed at a given observation point into bursts of updates based on a given prefix (refer to Figure 1). Given a burst, the new stable route is taken to be the last update in the burst. The valid route before the beginning of an update burst is considered to be the old stable route. Then, the procedure associate_event_burst is invoked to group different bursts into the same event if the bursts occur closely in time. Then, the bursts are condensed to identify the instability origins.

Since BGP updates to multiple prefixes are often correlated, the procedure in Figure 2 is used to correlate the located instability origins across multiple prefixes and identify clusters.

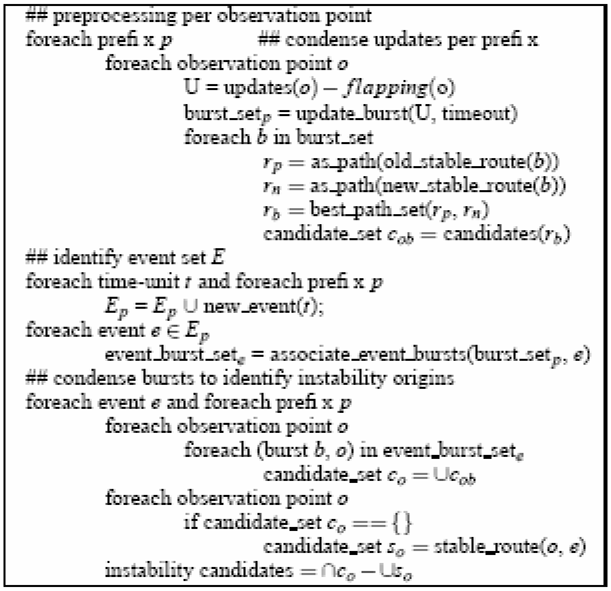

```
## preprocessing per observation point
foreach prefix $p$                ## condense updates per prefix
        foreach observation point $o$
                U = updates($o$) $-$ $flapping$(o)
                burst_set$_p$ = update_burst(U, timeout)
                foreach $b$ in burst_set
                        $r_p$ = as_path(old_stable_route($b$))
                        $r_n$ = as_path(new_stable_route($b$))
                        $r_b$ = best_path_set($r_p$, $r_n$)
                        candidate_set $c_{ob}$ = candidates($r_b$)
## identify event set $E$
foreach time-unit $t$ and foreach prefix $p$
        $E_p = E_p \cup$ new_event($t$);
foreach event $e \in E_p$
        event_burst_set$_e$ = associate_event_bursts(burst_set$_p$, $e$)
## condense bursts to identify instability origins
foreach event $e$ and foreach prefix $p$
        foreach observation point $o$
                foreach (burst $b$, $o$) in event_burst_set$_e$
                        candidate_set $c_o = \cup c_{ob}$
        foreach observation point $o$
                if candidate_set $c_o == \{\}$
                        candidate_set $s_o$ = stable_route($o$, $e$)
        instability candidates $= \cap c_o - \cup s_o$
```

Figure 1: Adopted methodology for locating instabilities per prefix [9]

3.2 Possible detection of Selective Dropping Attack

In the IANP scheme we design, we use the same methodology described in [9] to group BGP updates and locate the source of instability. The difference is when the locating instability procedure fails to locate a unique source of instability, we allow that observation point to probe its neighbors to see if they can identify the source of instability. Based on these additional information provided by its neighbors, an observation point may then be able to identify the source of instability. Next, this observation point will check its current routing table to see if the troubled inter-AS link is used in any current routing table. If this troubled link is used in any best-path route, then the detection module returns a true (indicating that there is a possible selective dropping attack). Otherwise, the detection module returns a false (no indication of selective dropping attack). The pseudo code of the selective dropping attack detection module of our IANP scheme is shown in Figure 3.

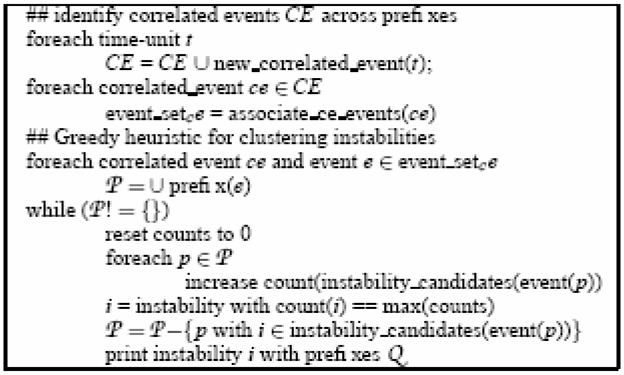

```
## identify correlated events $CE$ across prefixes
foreach time-unit $t$
        $CE = CE \cup$ new_correlated_event($t$);
foreach correlated_event $ce \in CE$
        event_set$_{ce}$ = associate_ce_events($ce$)
## Greedy heuristic for clustering instabilities
foreach correlated event $ce$ and event $e \in$ event_set$_{ce}$
        $\mathcal{P} = \cup$ prefix($e$)
while ($\mathcal{P} != \{\}$)
        reset counts to 0
        foreach $p \in \mathcal{P}$
                increase count(instability_candidates(event($p$))
        $i$ = instability with count($i$) == max(counts)
        $\mathcal{P} = \mathcal{P} - \{p$ with $i \in$ instability_candidates(event($p$))$\}$
        print instability $i$ with prefixes $Q$
```

Figure 2: Adopted methodology to correlate events across prefixes [9].

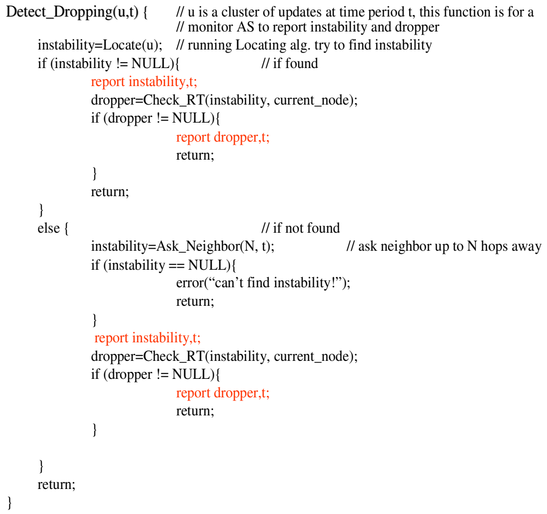

```
Detect_Dropping(u,t) {      // u is a cluster of updates at time period t, this function is for a
                            // monitor AS to report instability and dropper
    instability=Locate(u);  // running Locating alg. try to find instability
    if (instability != NULL){               // if found
            report instability,t;
            dropper=Check_RT(instability, current_node);
            if (dropper != NULL){
                    report dropper,t;
                    return;
            }
            return;
    }
    else {                                  // if not found
            instability=Ask_Neighbor(N, t);            // ask neighbor up to N hops away
            if (instability == NULL){
                    error("can't find instability!");
                    return;
            }
             report instability,t;
            dropper=Check_RT(instability, current_node);
            if (dropper != NULL){
                    report dropper,t;
                    return;
            }

    }
    return;
}
```

Figure 3: IANP selective dropping attack detection procedure

We illustrate our attack detection method using an example based on the network topology shown in Figure 4. This flat AS topology is generated using BRITE[10]. In Figure 4, we assume that each node represents an AS, there is a BGP router associated with each AS. When there is a link between two nodes, we assume that a BGP peer session is set up between the BGP routers associated with these two ASes.

Assume that link 30-15 is broken and this results in a burst of BGP update messages. If there is no selective dropping attack, all the nodes can locate the source of instability. However, if the link 30-15 is broken and at the same time, the router in AS15 launches a selective dropping attack towards AS14, then ASes 14, 13, 25, 26 cannot locate the source of the instability. AS14, however, can verify the source of instability after it has asked its neighboring routers i.e. ASes 16 & 3 for information on the source of instability that they may have identified for the same burst. Upon receiving the information from AS16, AS14 was able to check its routing table and discover a discrepancy about the status of the inter-AS link 14-15. Thus, AS14 was able to detect a potential selective dropping attack. AS14 can then issue a warning message which is propagated across the network. The warning message contains information about the identifier of the malicious router and any suspected broken links that are not reported. Such warning messages are authenticated so that attackers will not issue forged warning messages to confuse neighbors. BGP router identity authentication approach proposed in S-BGP[4] can be used for authenticating warning messages.

Without using the IANP scheme, five of the 870 (=30*29) AS routes will use an AS path that goes through the broken link as a result of the selective dropping attack. With the IANP scheme, only three of the 870 AS routes will use an AS path that goes through the broken link if no warning message is issued. If the warning message is flooded across the whole network, then no AS will utilize a path that goes through the broken inter-AS link and hence the damaged cost is reduced to 0.

One potential limitation of the IANP scheme is that unless a router has peers with at least 2 or 3 other routers, this router will not be able to get additional information to help it locate the source of instability as well as detect any potential malicious router that selectively drops BGP updates. One possible extension is to allow a router with more connectivity to propagate the information with limited scope. The warning message mentioned earlier plays such a role of alerting other routers with low BGP connectivity about the presence of potential malicious routers.

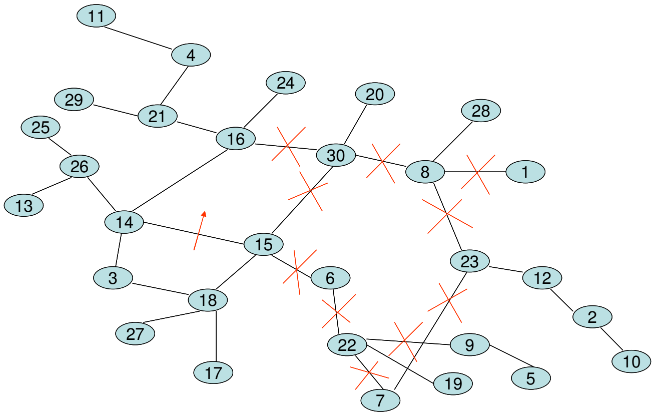

Figure 4: Network Topology 1

3.3 Corrective Action Upon Attack Detection

Any router that discovers that another router is potentially a malicious router can issue a warning message. If a router receives K warning messages implicating the same router as a potential malicious router, then this router can reduce the reputation score of this malicious router. When the reputation score of a router drops below a threshold, other ASes can include a routing policy that excludes any BGP messages sent by this malicious router. Another corrective action we can take is to remove all stable routes in the AS forwarding table that traverse the identified broken inter-AS link (which is the source of instability). The warning messages need to be authenticated so that no malicious attacks can be launched against the routing infrastructure using such messages.

4. Experimental Evaluation using DETER

We conduct extensive experiments using DETER [14] to evaluate the effectiveness of our IANP scheme. Three flat AS topologies are generated using Brite [10]. For each topology, we introduce 10 inter-AS link failures with and without selective dropping attacks. For each link failure, we evaluate the percentage of ASes that can identify the source of instability as well as detecting the presence of selective dropping attacks. For the cases where the ASes cannot detect the presence of selective dropping attacks, we also evaluate the damage cost. We define the damage cost as the ratio of the number of stable-paths that utilize the broken inter-AS link over the number of total stable-paths. In our experiments, since there is only one prefix per AS, we have at most n*(n-1) stable paths where n is the number of ASes (nodes) in each network. When there is a fault, the number of possible paths may be smaller than n*(n-1). For each experiment we conduct, we show the damage cost with and without taking the corrective action of removing all stable routes that traverse the identified broken inter-AS link.

4.1 Test Scenario 1

We conduct 10 experiments using Network Topology 1 illustrated in Figure 4. In each experiment, we fail one of the inter-AS links shown as red crosses in Figure 4. The burst of BGP UPDATE messages that results from this failed link are collected and fed to the IANP analyzer running at each BGP router. Then, we repeat the same experiment (i.e. failing one of the marked inter-AS links) but also let AS15 launch a selective dropping attack towards AS14 by dropping all BGP updates from AS15 to AS14.

Our results are tabulated in Table 1. In the table, w1 is defined as the percentage of nodes that cannot locate the source of instability and DC refers to the damage cost that is defined earlier. Without the selective dropping attack, the percentage of ASes that cannot locate the source of instability ranges from 0 to 20%. With the selective dropping attack, the percentage of ASes that cannot locate the source of instability increases by 3.3 to 13.3%. Note that we count absolute increase not relative increase i.e. without selective dropping attack, if 6 nodes (out of 30 notes) cannot locate the source of instability and with selective dropping attack, 10 nodes cannot locate the source of instability, then the increase is (10-6)/30*100%=13.3%. All the cases of not being able to locate the source of instability is due to the fact that these ASes do not receive any BGP updates due to the instability. When the IANP scheme is used, the percentage of ASes that cannot locate the source of instability decreases slightly by 3.3 to 6.7%.

With selective dropping attack, the damaged cost without deploying IANP for these 10 instabilities range from 0 to 7.8%. If IANP without warning message is deployed, then the damage cost reduces (range from 0 to 3.1%). With the additional warning messages, the damage cost is reduced to 0.

	Without IANP		IANP Without Warning		IANP with Warning	
	w1	DC	w1	DC	w1	DC
Expt 1	20%	6.20%	13.30%	0.60%	0%	0%
Expt 2	10%	2.80%	10%	1.40%	0%	0%
Expt 3	10%	7.80%	10%	0.00%	0%	0%
Expt 4	13.30%	7.10%	10%	3.10%	0%	0%
Expt 5	33.30%	0%	33.30%	0%	0%	0%
Expt 6	0%	3.30%	0%	0%	0%	0%
Expt 7	20%	4.30%	16.70%	1.70%	0%	0%
Expt 8	33.30%	0.60%	33.30%	0.30%	0%	0%
Expt 9	10%	6.40%	10%	0.70%	0%	0%
Expt 10	23.30%	0.00%	23.30%	0%	0%	0%

Table 1: Results of Test 1

4.2 Test Scenario 2

Figure 5 shows the network topology used for Experiment 2. Again, we introduce 10 inter-AS link failures and for each link failure, we collect burst of BGP updates with and without the BGP selective dropping attack launched by AS16 towards AS23. Our results are tabulated in Table 2.

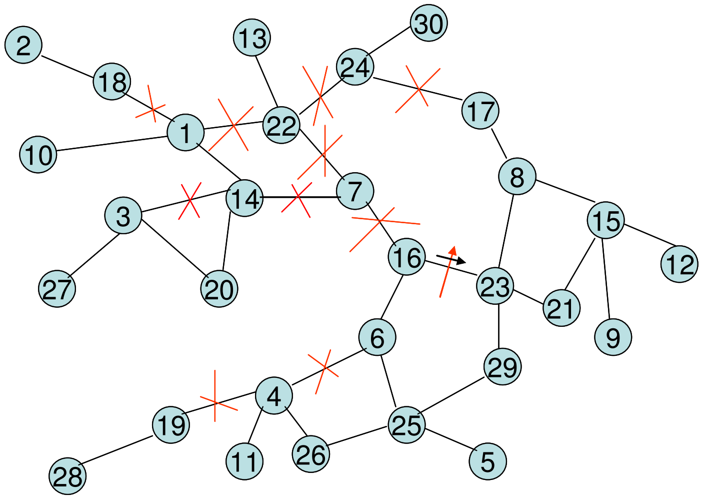

Figure 5: Network Topology 2

	Without IANP		IANP Without Warning		IANP with Warning	
	w1	DC	w1	DC	w1	DC
Expt 1	0.0%	6.9%	0.0%	0.0%	0%	0%
Expt 2	0.0%	3.2%	0.0%	0.0%	0%	0%
Expt 3	0.0%	0.0%	0.0%	0.0%	0%	0%
Expt 4	13.3%	0.0%	13.3%	0.0%	0%	0%
Expt 5	30.0%	1.4%	26.7%	0.5%	0%	0%
Expt 6	26.7%	5.3%	23.3%	2.0%	0%	0%
Expt 7	13.3%	1.6%	13.3%	1.1%	0%	0%
Expt 8	23.3%	17.5%	20.0%	4.8%	0%	0%
Expt 9	16.7%	6.9%	16.7%	1.2%	0%	0%
Expt 10	23.3%	4.1%	20.0%	2%	0%	0%

Table 2: Results of Experiment 2

From Table 2, we see that without using any detection scheme, the damage cost caused by the selective dropping attack ranges from 0 to 17.5%. A maximum of 30% of the nodes cannot detect the source of instability when selective dropping attack is present. With the IANP scheme deployed but without the warning message, the damage cost has been reduced to a maximum of 4.8%. The maximum number of nodes that cannot detect the source of instability is reduced to a maximum of 26.7%. When combined with the warning message, the IANP scheme reduce the damage cost to 0. With the warning message, all nodes can be alerted of the potential problems and take appropriate corrective actions.

4.3 Test Scenario 3

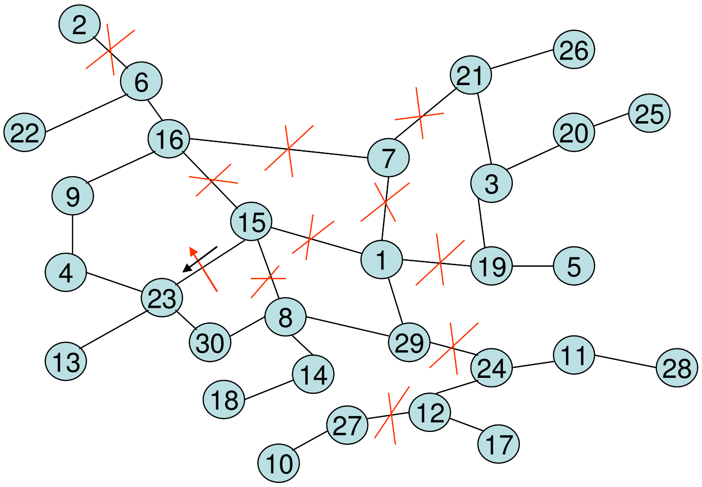

Figure 6: Network Topology 3

Figure 6 shows the network topology used for Experiment 3. Again 10 inter-AS link failures are introduced and for each link failure, we collect the burst of BGP updates with and without the selective dropping attack launched by AS 15 towards AS23.

	Without IANP		IANP Without Warning		IANP with Warning	
	w1	DC	w1	DC	w1	DC
Expt 1	3.3%	3.4%	3.3%	0.1%	0.0%	0.0%
Expt 2	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Expt 3	13.3%	0.0%	13.3%	0.0%	0.0%	0.0%
Expt 4	6.7%	1.1%	6.7%	0.2%	0.0%	0.0%
Expt 5	13.3%	9.0%	3.3%	1.7%	0.0%	0.0%
Expt 6	13.3%	2.6%	10.0%	0.2%	0.0%	0.0%
Expt 7	10.0%	2.6%	7.0%	0.6%	0.0%	0.0%
Expt 8	23.3%	0.0%	23.3%	0.0%	0.0%	0.0%
Expt 9	3.3%	22.7%	3.3%	0.1%	0.0%	0.0%
Expt 10	0.0%	6.9%	0.0%	0.0%	0.0%	0.0%

Table 3: Results of Experiment 3

The results in Table 3 shows that without the IANP scheme, the damaged cost as a result of the selective drop attack ranges from 0 to 22.7%. With the IANP scheme but without the warning message, the damaged cost already reduces significantly to a maximum of 1.7%. With the warning message, the damaged cost is reduced to 0.

Note that in all these 3 experiments, we only allow the AS to probe its immediate neighbors. We anticipate that if we allow the AS to probe its K-hop neighbors, the percentage of nodes that cannot locate the source of instability will decrease since all the cases where the ASes cannot locate are due to having no BGP updates.

5. Discussions

There are several issues that warrant further investigation before one can deploy the IANP scheme. First, it is interesting to investigate the minimum percentage of ASes that need to deploy IANP in a large AS topology to reduce the damaged cost to less than a certain threshold (say 2% or 5%). Second, the AS topologies we use for our experiments in DETER are relatively small compared to the real Internet AS topologies. The effectiveness of the IANP scheme in bigger Internet AS topologies needs to be evaluated. Furthermore, our experiment also assumes a flat AS topology. It will be interesting to investigate the effectiveness of the IANP in a large tiered AS topology. We intend to use a home-grown BGP simulator or SSFNet [15] for such evaluations.

In addition, there are several enhancements that we can introduce to the IANP scheme. For example, we anticipate that a hierarchical based IANP scheme can be deployed. For example, within a big AS, only a limited number of BGP routers need to deploy the IANP scheme e.g. the BGP routers that have the highest number of peers for they see the largest number of BGP updates resulting from any instability. Upon detection of potential malicious routers, warning messages can be issued to all BGP routers within the same AS. Such warning messages need to be authenticated to avoid being used as bogus warning messages to disrupt the Internet routing infrastructure.

There are several deployment issues that need to be addressed too. First, we need to propose adding two new BGP message types for neighbor probing and warning purposes. Next, one needs to carefully think about the scope of the warning message distribution. For scalability, we anticipate that the warning message can be flooded hierarchically based on Internet tiers and/or setting a maximum TTL limit to the warning message being relayed. In addition, one can add intelligence to the processing of the warning message to decide if a router needs to propagate the warning message. For example, if the impact of the malicious attack is localized, the warning message needs not be propagated beyond a certain tier. Another issue is related to limiting the number of neighbor probing messages that can be issued per BGP peer so that this will not be used as a resource-consumption attack towards a BGP router. We intend to investigate all these interesting issues in the near future.

6. Related Work

Packet dropping attack has been studied in other research work. X. Zhang et al [12] explored the negative impact of packet dropping attacks in TCP. K. Bradley et al [13] presented WATCHERS protocol to detect and react to routers that drop or misroute packets. Using multiple decentralized counters, WATCHERS track the traffic flow and detect routers which violate the conservation principle. Comparing to malicious dropping of data packets, dropping routing packets may cause more damages.

Bellovin et al [6] studied another attack model called link-cutting. They assume that attackers have a router-level topology map and a list of already-compromised links and routers in advance. The attackers can cut/disable some key links so that the selected traffic will pass through the compromised routers. L. Subramaniam et al [13] propose two security mechanisms for BGP, namely Listen and Whisper. Listen passively probes the data plane and checks whether the underlying routes to different destinations work. Whisper uses cryptographic functions along with routing redundancy to detect bogus route advertisements in the control plane.

The work in [8] is the closest to our work and in fact is the paper that inspires our work. In [8], the authors focus on investigating the impacts of selective dropping attacks but did not elaborate on how to detect and mitigate against such attacks.

7. Conclusions

In this paper, we have described an instability analysis with neighbor probing (IANP) scheme for detecting selective dropping attacks in BGP. We have conducted extensive experiments using DETER to evaluate the effectiveness of our IANP scheme. We show that many observation points can locate the source of instability even in the presence of selective dropping attacks. For those observation points that cannot locate, 90% of the time they will be able to after probing its neighbors. In addition, such nodes issue warning messages after detecting the presence of potential malicious router. The damage cost can be reduced from a maximum of 22.7% to a maximum of 4.8% when our IANP scheme without the warning message is deployed. With the warning message, the damaged cost is reduced to 0.

This work is only preliminary. As indicated in earlier sections, there are several topics that we wish to explore further. We intend to build a simulator that allows us to take the BGP updates from archives like RouteViews [16] and emulate selective dropping attacks to see if our scheme is still effective in large-scale Internet AS topologies. In addition, we intend to investigate if the scheme still works well with limited numbers of deployed IANP-enabled routers. It is also interesting to evaluate the effectiveness of the IANP scheme with different placements of a limited number of IANP-enabled routers in a large-scale network. Last but not least, we intend to build a prototype of our scheme and conduct a more realistic experiment in DETER.

8. REFERENCES

[1] Y. Rekhter and T. Li, “Border Gateway Protocol 4”, RFC 1771, SRI Network Information Center, July, 1995.

[2] T. Wan, “Analysis of BGP prefix origins during Googles May 2005 outage”, 2005.

[3] A. Heffernan, “Protection of BGP sessions via the TCP MD5 signature option”, RFC 2385, SRI Network Information Center, August, 1998

[4] S. Kent et al, “Secure Border Gateway Protocol (S-BGP)”, IEEE Journal on Selected Areas in Communications, Volume18 Issue 4, April, 2000.

[5] C. Kruegel et al, “Topology-based detection of anomalous BGP messages”, Proceedings of 6th Symposium on Recent Advanced in Intrusion Detection (RAID), 2003.

[6] S. M. Belovin and E. R. Gansner, “Using Link Cuts to attack Internet Routing”, draft, May 2003.

[7] S. M. Belovin, “Routing Security”, Talk at British Columbia Institute of Technology, June 2003.

[8] K. Zhang, X. Zhao, F. Wu, “An alaysis of Selective Dropping Attack in BGP“, Proceedings of IEEE IPCCC, April, 2004.

[9] A. Felman, O. Maenel, M. Mao, A. Berger, B. Maggs, ”Locating Internet Routing Instabilities”, Proceedings of ACM Sigcomm, Aug, 2004.

[10] Brite – Internet Topology Generator, http://www.cs.bu.edu/brite/index.html

[11] X. Zhang, S. F. Wu, Z. Fu, T. L. Wu, "Malicious Packet Dropping: How it might impact the TCP performance and how we can detect it", Proceedings of IEEE ICNP, 2000.

[12] K. A. Bradley etc, "Detecting disruptive routers: a distributed network monitoring approach", Proceedings of IEEE Symposium on Security and Privacy, 1998.

[13] L. Subramaniam, V. Roth, I. Stoica, S. Shenker, R. H. Katz, "Listen and Whisper: Security Mechanisms for BGP", Proceedings of NSDI, 2004.

[14] DETER – A Laboratory for Security Research, http://www.isi.edu/deter/

[15] SSFNet, http://www.ssfnet.org

[16] Oregon RouteView http://www.routeviews.org