Forensic Analysis of Autonomous System Reachability

DK Lee, Sue Moon
Division of Computer Science, KAIST
{dklee, sbmoon}@an.kaist.ac.kr

Taesang Choi, Taesoo Jeong
ETRI
{choits, tsjeong}@etri.re.kr

ABSTRACT

Security incidents have an adverse impact not only on end systems, but also on Internet routing, resulting in many out-of-reach prefixes. Previous work has looked at performance degradation in the data plane in terms of delay and loss. Also it has been reported that the number of routing updates increased significantly, which could be a reflection of increased routing instability in the control domain. In this paper, we perform a detailed forensic analysis of routing instability during known security incidents and present useful metrics in assessing damage in AS reachability. Any change in AS reachability is a direct indication of whether the AS had fallen victim to the security incident or not.

We choose the Slammer worm attack in January, 2003, as a security incident for closer examination. For our forensic analysis, we use BGP routing data from RouteViews and RIPE. As a way to quantify AS reachability, we propose the following metrics: the prefix count and the address count. The number of unique prefixes in routing tables during the attack fluctuates greatly, but it does not represent the real scope of damage. We define the address count as the cardinality of the set of IP addresses an AS is responsible for either as an origin or transit AS, and observe how address counts changed over time. These two metrics together draw an accurate picture of how reachability to or through the AS had been affected. Though our analysis was done off-line, our methodology can be applied on-line and used in quick real-time assessment of AS reachability.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations—*Network Monitoring*

General Terms

Measurement

Keywords

BGP, security incidents, AS reachability, prefix count, address count

SIGCOMM'06 Workshops September 11-15, 2006, Pisa, Italy.

1. INTRODUCTION

Security incidents such as, system cracking, DDoS attacks, and flash worms, have an adverse impact not only on end systems, but also on Internet routing, resulting in many out-of-reach prefixes. Previous work has looked at packet latency and loss during the security incidents and has shown degrees of deterioration in the Internet data plane [1–3]. Intuitively, security incidents exert a bad effect on the Internet routing, which is crucial to the reliability of the Internet. Certain ASes become partially or totally unreachable during such incidents. However, quantitative analysis on Internet reachability is not well provided yet. The key question then is how to measure and quantify the extent of damage in terms of the Internet reachability, which is the focus of this paper.

In this work, we perform a detailed forensic analysis of routing instability during a known security incident and present useful metrics in assessing the level of damage in AS reachability. We choose the Slammer worm attack in January, 2003, as a security incident for closer examination, for it is the most recent attack that paralyzed the e-commerce infrastructure globally for several hours and is known to be the fastest spreading worm. For our forensic analysis, we use BGP routing table snapshots and updates between Jan. 24th and 27th, 2003, from RouteViews and RIPE NCC RIS. As a way to quantify AS reachability, we propose the following metrics, *the prefix count* and *the address count*.

Previously, routing instability was inferred from changes in the number of routing updates. In the case of the Slammer worm attack, the increased number of route withdrawals was more prominent than that of route advertisements. Consequently, the number of unique prefixes in routing tables for the 3 day time period fluctuated greatly. However, the unique prefix count does not represent the real scope of damage in terms of reachability. We define the address count as the cardinality of the set of IP addresses that an AS is responsible for either as an origin AS or transit AS. We, then, observe how address counts change over time. In our analysis the address count demonstrates how reachability to an AS or through the AS had been affected.

Our contributions are summarized as follows. First, we define the address count as a measure to quantify and track down the extent of damage on each AS during security incidents. We demonstrate our metrics are simple, but effective in identifying ASes that experienced much damage. Second, we show the possibility of forensic examinations on AS reachability by using the readily available BGP data from RouteViews and RIPE NCC RIS. Though our analysis was done off-line, our methodology can be applied on-line and used in quick real-time assessment of AS reachability. We hope NSPs (network service providers) will use our metrics to keep track of their customer ISPs and help them when they see a need.

The paper is organized as follows. Section 2 reviews the related

work regrading the impact of security incidents on BGP changes. Section 3 propose the methods to quantify the change in AS reachability and analyze the extent of damage caused by security incidents. This section also presents the algorithm and the guideline to analyze the result of it. Section 4 examines the impact of security incidents on AS reachability with our methods and Section 5 concludes the paper.

2. RELATED WORK

Previous work analyze the BGP update surges during Slammer worm impact [2,4,5], and studies the BGP behavior under the stress of large-scale accidents such as large-scale power outages [6]. This work is different from ours in that they did not deal with the problem of quantifying the extent of damage. They count BGP updates and analyze the possible reasons for surging points during the impact of security incidents.

Wang *et al*. [7] showed how BGP actually performs under stressful conditions by classifying the BGP updates into classes and then inferring leading causes for each class. Similarly, Li *et al*. used BGP updates to detect Internet anomalies [8]. The objective of their work is to tag a BGP event (i.e., a series of BGP updates) with one of the following labels: normal, blackout, worm, or misconfiguration. They also introduce an Internet Routing Forensics framework to process BGP routing data systematically.

Xie *et al*. [9] defines static reachability of IP networks as a set of packets routable from one point to another. They use packet filters, routing information, and packet transformations to calculate static reachability. As mentioned in their future work, their definition is different from ours in that we consider AS reachability as a whole per AS and analyze its time-varying behavior.

Our work is also related to the problem of locating an origin of routing instabilities. Feldmann *et al*. shows the method to locate the Internet instabilities with the insights that if there is an AS path change, then some instability has to have occurred on one of two AS paths, the previous best path or the new best path [10]. Our work focuses on the eventual impact of such instabilities on reachability.

3. METHODOLOGY

An autonomous system (AS) is a collection of subnetworks under the control of one administrative entity that presents a common routing policy. Typically, an AS owns a set of prefixes, and those prefixes are not always reachable from everywhere in the Internet. If the interdomain routing changes due to network maintenance or unexpected failures, the reachable portion from a vantage point in the Internet to the AS could be changed. By analyzing how often and long certain prefixes of an AS are unreachable, we can infer if the AS is experiencing a problem.

Our objective is to analyze the change in the reachable prefixes of ASes during a known security problem, and evaluate the usefulness of our methodology in assessing the damage from the attack. In our forensic analysis, we examine how the available forwarding paths to the AS disappear and reappear during the attack period. In doing so, we use the following two metrics, the prefix count and the address count. We show that these two metrics are simple, but effective ways to assess damage to any AS in the Internet. As calculation of prefix and address counts is straightforward from any BGP table, ISPs can monitor their BGP tables continuously and even perform online analysis, as attacks take place and damage occur.

3.1 BGP Data

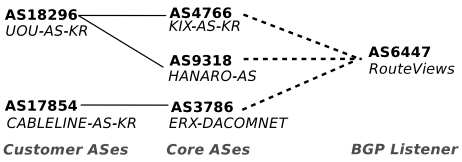

Figure 1: An example network with two stub ASes and three transit ASes based on RouteViews data. A solid line is the direct link between two ASes. A dotted line means a multihop link in which two or more ASes may be involved. AS18296 is multihomed and has two exit points, AS4766 and AS9318. On the other hands AS17854 has only one exit point AS3786.

We use two sets of BGP routing table snapshots: one from RouteViews' BGP listener, `route-views2.oregon-ix.net`, taken at 2 hour intervals [11], and the other from RIPE NCC Routing Information Service (RIS) taken at 8 hour intervals [12]. RIPE NCC operates eight monitoring points (rrc00 - rrc07) and each monitoring BGP listener peers with different ASes. Our analysis is based on the data from the rrc00 monitoring point at RIPE NCC in Amsterdam.

Routing table snapshots are stored in MRTD (Multi-threaded Routing Toolkit Daemon) format [13]. A snapshot taken at every two hour interval from the RouteViews is about 380 MB in MRTD machine readable ASCII format and contains more than 3 million entries. A route in a BGP table snapshot contains the AS-PATH, the next-hop AS number, and the next-hop IP address per destination prefix. The AS-PATH attribute is a list of ASes that describes the path to the prefix. It tells who owns and originally announces the prefix (origin AS) and who provides transit service for the prefix (transit AS). In other words, by examining AS-PATH attributes of a BGP table, we can list origin prefixes that belong to an AS and transit prefixes that use transit service provided by this AS.

There are some challenges to consider when using this BGP data for our network reachability study. First, we cannot examine a single best path to an AS. Peers of a BGP listener have their own best BGP paths to the destined IP prefix, and inform the listener what they know through the external BGP (eBGP) multihop peering sessions. Eventually, a BGP listener archives multiple best BGP paths for a given prefix. With the BGP routing table from a passive listener of an eBGP session, ranking BGP paths to find the best one is pointless. Second, according to the BGP selective export rule, a peer AS of a BGP listener usually does not export routes learned from its provider and peers [14]. Therefore, listeners only hear limited routing information from their peers, and its routing table collects an incomplete view of the Internet.

For this work, we focus on the data collected during the Slammer worm released on January 25th, 2003. Figure 1 shows the connectivity graph of the ASes which we will refer later to examine the time-varying aspect of network reachability. We also use the Internet hierarchy data [15], which identifies the position of each AS in the Internet on January 9th, 2003. This data contains 14695 ASes and their corresponding hierarchical levels: dense core, transit core, outer core, small regional ISP, and customer. For the convenience of analysis, we divide the Internet hierarchy data into three regional lists: ARIN, RIPE, and APNIC, based on the IANA ASN range allocation information [16]. ARIN has 19 dense cores, 60 transit cores, 612 outer cores, 682 small regional ISPs, and 8056 customers. We provide analysis results for ASes in ARIN only. For full results, refer to our technical report [17].

3.2 Origin and Transit Prefix Counts

A route in a BGP table specifies the AS path to a destination prefix. The last AS in the AS path is the one that originates the prefix. Other ASs in the AS path only transit traffic destined to the prefix. We refer to the prefix an AS originates as an origin prefix, and that an AS transits as a transit prefix.

For any given AS, we count the number of origin and transit prefixes found in BGP table snapshots. For each of RouteViews and RIPE NCC snapshots, we calculate the origin and transit prefixes separately, for the two snapshots from RouteViews and RIPE NCC are collected at different times over different intervals and cannot be combined easily. As a BGP listener hears more than one best paths to any given prefix, we only count once the same origin and transit prefixes from a BGP table snapshot.

Any decrease or increase in the origin and transit prefix counts signals change in the reachability to or through the specific AS. Though intuitive, these two metrics are limited in the following ways. First, it does not take the length of a prefix into consideration. Prefixes with different lengths are counted all equally once, and thus their magnitudes are ignored. Therefore, any two ASes, with the same origin and transit counts, are always regarded as having the same level of reachability. However, AS X with $1.0.0.0/8$ has a larger number of reachable hosts than AS Y with $1.2.3.4/28$ which has only 16 hosts. Second, inclusive relationships among prefixes are not reflected in this counting method. For example, let AS X has the following four prefixes: $1.0.0.0/8$, $1.1.0.0/16$, $1.2.0.0/16$, and $1.255.0.0/16$. In this case, $1.0.0.0/8$ includes all other prefixes, but the prefix count is 4, not 1. Similarly, prefix aggregation and disaggregation over time cause change in AS reachability, if we only consider prefix counts. An aggregation of two prefixes into one decreases the prefix count by one. However, it does not decrease the actual number of addresses an AS originates or transits.

3.3 Origin and Transit Address Counts

To deal with the above limitations of prefix counts, we choose to count the number of unique addresses as a way to quantify the level of AS reachability. The number of unique addresses is the cardinality of a set of IP addresses, and we term it as *address count*. Our address counting algorithm works with a set of IP prefixes, and this prefix set is constructed from a BGP routing table snapshot. Following notations are used to define prefix sets.

$$origin(p) : \text{origin AS of prefix } p$$

The function $origin()$ returns the origin AS of a destination prefix p from a BGP routing table snapshot.

$$transit(p, X) = \begin{cases} True, & origin(p) \text{ is not } X \\ & \text{and } X \text{ is in AS-PATH of } p \\ False, & \text{otherwise} \end{cases}$$

The function $transit()$ takes two parameters, a prefix p and an AS number X, and returns true, if the origin AS of the prefix p is not X and X is in the AS-PATH of p; otherwise, it returns false. The function $transit()$ indicates if an AS X appears in a path to a destination prefix p in a given BGP table.

From a BGP routing table snapshot, we consider the following three sets of prefixes, $\mathbf{U}$, $\mathbf{O}(X)$, and $\mathbf{T}(X)$; the latter two sets are parameterized by an AS X.

$$\begin{aligned} \mathbf{U} &: \text{ all the prefixes in a BGP routing table snapshot} \\ \mathbf{O}(X) &: \text{ a set of origin prefixes of an AS X} \\ &\quad = \{p \mid origin(p) = X\} \\ \mathbf{T}(X) &: \text{ a set of transit prefixes of an AS X} \\ &\quad = \{p \mid transit(p, X) = True\} \end{aligned}$$

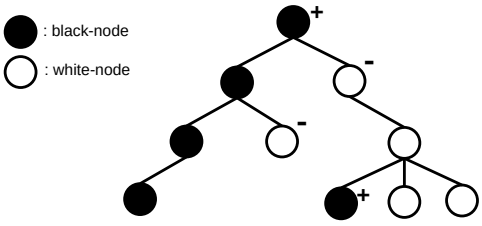

Figure 2: Prefix tree. For any node x, $\forall y$ in the subtree of x satisfies $incl(x, y)$.

For any AS X of which we aim to analyze the reachability, we consider two different kinds of address counts: origin and transit address counts. Simply speaking, the origin address count is the number of unique IP addresses in $\mathbf{O}(X)$ and the transit address count is the number of unique IP addresses in $\mathbf{T}(X)$. Unfortunately, calculating the origin count is not so simple. Let us consider the following example. AS X decides to allocate a block of its address space to a customer, and then this block of addresses or a prefix is owned and administrated by the customer, not by AS X. In this case, we claim that such prefix no longer belongs to AS X. It punches a hole in the original block of addresses that AS X owns. Moreover, some portion of addresses given out to the customer could be reclaimed by AS X. We should take inclusive relations between ASes into consideration when calculating the origin address count. On the other hand, punch holing is of minor concern when calculating the transit address count.

In order to compute the address count, we first need a better representation of a set of IP prefixes which is useful to recursive calculation of an address count. In this paper, we suggest to use a tree representation, *prefix tree*, for representing prefix sets and calculating the address count. We represent the punch-holed relations between the prefixes as parent-child relations in a tree. First, we define the function $incl()$ on prefix p and q as follows.

$$incl(p, q) = \begin{cases} True, & \text{if } len(p) < len(q) \text{ and} \\ & len(p) \text{ bits of } q \text{ are the same as } p \\ False, & \text{otherwise} \end{cases}$$

Given $\mathbf{U}$ and AS X of our interest, we calculate the origin address count (OAC) as follows. In order to calculate the origin address counts while reflecting the inclusive relations, we build trees of prefixes from $\mathbf{O}(X)$. The easiest way to construct prefix trees from $\mathbf{O}(X)$ is to sort prefixes in $\mathbf{O}(X)$ in an ascending order of prefix lengths, make the first prefix a root, and insert remaining prefixes into the tree. A prefix is inserted as a child of a node with the longest matching prefix. That is, a parent-child relation in a prefix tree reflects $incl(x, y)$, where x is a parent and y is a child. If no root satisfies the $incl(root, p)$ for a yet-to-be-inserted node p, then create a new tree with a root p.

Once the prefix trees are built from $\mathbf{O}(X)$, we label all the nodes in the prefix trees as black nodes. Then we insert all the remaining prefixes of $\mathbf{U} - \mathbf{O}(X)$ into the prefix trees, while satisfying the same parent-child $incl()$ relation as above, and label them as white nodes. We do not create a prefix tree with a white-node as a root, and simply discard such white nodes.

Once prefix trees are constructed, calculating the origin address count is straightforward. Visit all the nodes in a tree once, and check if the node is the same type as its parent. If yes, then do nothing. Otherwise, add the number of addresses of the node's prefix to the total origin address count, if the node is a black node.

Table 1: Changes in the sum of origin address counts before and after the Slammer worm attack

	ARIN	RIPE	APNIC
Total # of ASes	9429	4175	1770
Sum of OAC at 04:18	972,550,686	154,140,342	101,158,364
Sum of OAC at 06:19	934,193,454	153,016,886	98,735,452
Difference	-38,357,232	-1,123,456	-2,422,912
# of decreased ASes	274 (-38,806,832)	96 (-1,165,696)	97 (-2,511,872)
# of increased ASes	62 (+449,600)	26 (+42,240)	9 (+88,960)

Subtract if the node is a white-node. When a child node is of the same type as a parent node, then the prefix of a parent includes all the addresses of the child and the child node's prefix need not be considered in address counting.

Figure 2 shows an example of a prefix tree. Nodes are colored according to their types. Nodes marked either with a plus or minus sign are those included in the address count calculation as described above. Those not marked are of the same type as their parents and not considered in the address count calculation.

The origin address count of an AS X changes if there are newly announced or withdrawn prefixes by X. Moreover, the origin address count also depends on the number of punched holes and the size of a hole. That is there could be an administrative change on its customer or neighbor ASes.

Finally, the transit address count is simply the cardinality of the set of IP addresses that belong to prefixes in $\mathbf{T}(X)$.

Clearly, our address count is limited in other fundamental ways. First, without knowledge on how many addresses are used actually in the advertised prefixes, it is hard to know whether the metric does quantify changes in reachability accurately or not. Moreover, a transit address count shows changes in the routing path, but in an extreme case, it does not distinguish between a damage of security incidents and an artifact of normal routing policy changes such as load balancing.

4. ANALYSIS RESULTS

In this section, we provide the results of our forensic analysis on the autonomous system reachability during the chosen security incident, the Slammer worm. We investigate if there was any globally observable change in the origin address count before and after the worm attack. We observe that the worm attack had different impact on prefix and address counts. They did not increase or decrease in synchrony. Then we perform a microscopic analysis of origin and transit address counts on 5 ASes in Figure 1.

4.1 Changes in Address Counts

The Slammer worm began to infect hosts at 05:30 UTC on Saturday, January 25, 2003. We choose two routing table snapshots, rib.20030124.2118 and rib.20030124.2319 from RouteViews. They were taken at 04:18 UTC and 06:19 UTC, just before and 49 minutes after the attack, respectively. Then, we calculate the origin address counts (OAC) for 14,695 ASes that are listed on Agarwal *et al.*'s Internet hierarchy data from the same day.

Table 1 shows the amount of changes in the sums of OACs of all 14,695 ASes before and after the worm attack. We obtain the total sums for ARIN, RIPE, and APNIC ASes separately. From the table, we can see that the total number of reachable addresses decreased during the Slammer worm period. Interestingly enough, several ASes actually experienced increases in their address counts. We will give plausible explanations for this appearance later in this section.

Figure 3 is the result of origin address counts and origin prefix counts for ARIN ASes. ASes on the x-axis are sorted by their AS hierarchical levels. ASes in the same AS level are randomly distributed and they have the same order in all graphs of Figure 3.

Figures 3(a) and 3(b) plot the amount of address counts, just before the attack. ASes with the same number of prefix counts could show different level of reachability with the address count because of the different size of IP prefixes they have. Interestingly, the result of address counts does not correspond to the Internet hierarchy data which are based on relationships among ASes. In Figure 3(b), ASes of HP, MIT, Exchange Point Block, Apple, and General Electrics are classified as the customer ASes, but they have mostly the same amount of address counts as the core or regional ISPs have.

Figure 3(c) and 3(d) are the amount of differences between the snapshot rib.20030124.2118 and rib.20030124.2319. Customer ASes have far smaller number of prefixes then core ASes and thus their decrease in the origin prefix counts after the attack is not readily noticeable in Figure 3(c). OACs of the majority of customer ASes are also smaller than those of core ASes. Worse yet, two customer ASes that have relatively large OACs saw a serious decrease in their OACs and overshadowed other ASes' decrease in OACs in Figure 3(d).

From absolute differences of prefix and address counts before and after the attack as in Figures 3(c) and 3(d), we cannot tell if there was any significant change in AS reachability. To help better visualize change due to the attack, we normalize differences of prefix and address counts as *(before - after)/before* and plot them in Figures 3(e) and 3(f). Now we see that many of the ASes actually did not have any transit prefix or originated any prefix: 100% loss of reachability. We could infer the extend of damage from the decrease in the address and prefix counts.

Moreover, we can compare the amount of decreases of one AS with others to figure out the relative amount of damage. With the result of the prefix count, we can say that University of Puerto Rico suffered the most serious problem in their reachability. This is because University of Puerto Rico initially had 2 $/16$ prefixes, 163 $/24$ prefixes, 2 $/25$ prefixes, and 4 $/26$ prefixes, but during the worm attack, it showed only 1 $/16$ prefix and 80 $/24$ prefixes. On the other hand, the amount of damage that University of Puerto Rico suffers is not serious in the result of the address count. With the result of the address count, HP and Exchange Point Block seem to suffer the most serious amount of damages.

4.2 Per-AS Microscopic Analysis

In this section, we apply our address count to the several routing table snapshots during the Slammer worm period, and present the corresponding results. We then try to identify the possible reasons for changes in reachability. We calculate the address count for two distinct types of ASes, customer AS and transit core AS. A customer, especially stub AS is always the last AS in any AS path in which it appears, and it could be regarded as a simple extension of the other AS. A transit core AS has connections to more than one AS and allows itself to be used as a conduit for traffic between other ASes. First, we examine the case of customer ASes that are relatively straightforward to analyze, then we investigate the transit core ones.

Reachability to a Customer AS

Here we give forensic analysis results of two customer ASes, AS17854 and AS18296.

We use our address count as follows. Since we are dealing with the case of customer ASes which do not provide a transit service to other ASes, we only consider the variation in origin address counts.

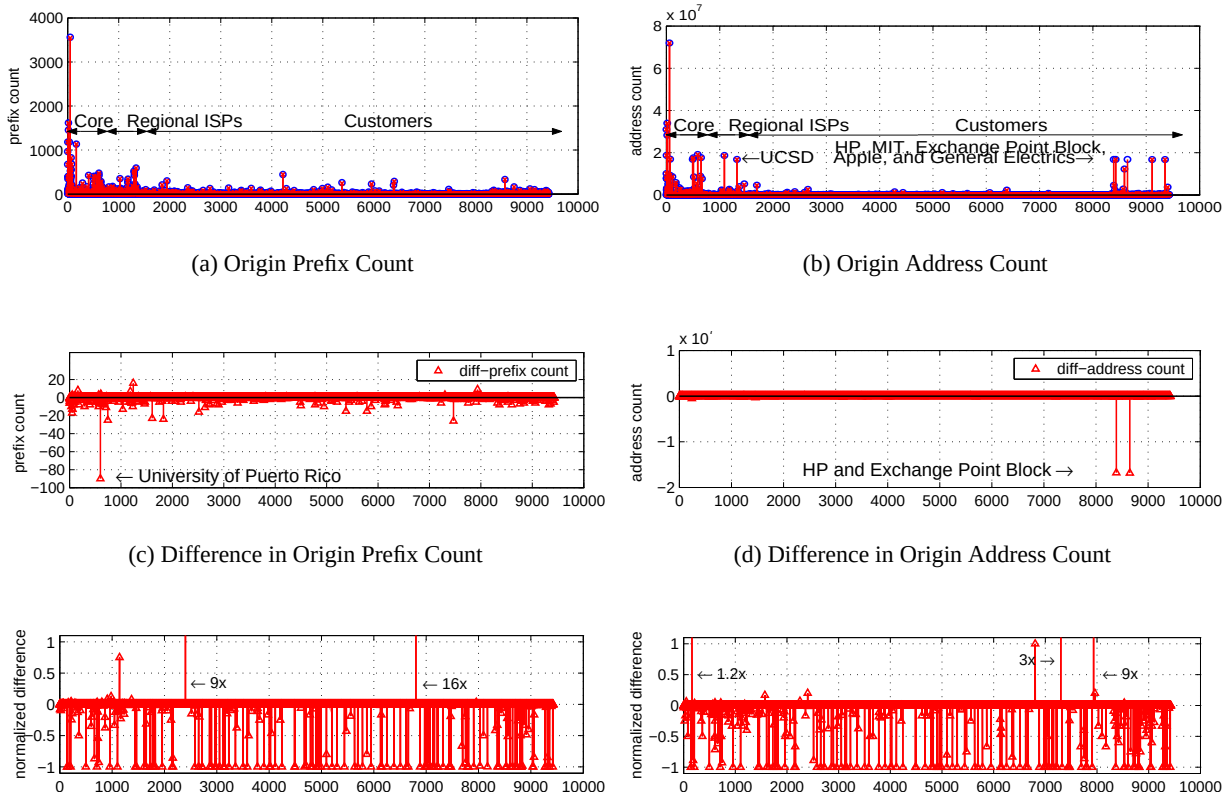

(a) Origin Prefix Count

(b) Origin Address Count

(c) Difference in Origin Prefix Count

(d) Difference in Origin Address Count

(e) Normalized Difference in Origin Prefix Count

(f) Normalized Difference in Origin Address Count

Figure 3: ARIN ASes

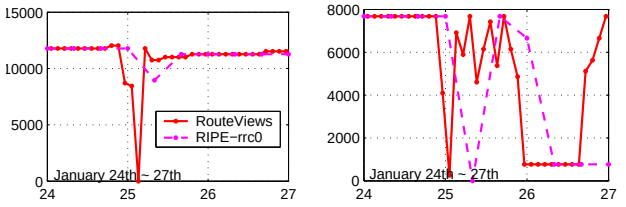

(a) AS17854

(b) AS18296

Figure 4: Result of origin address count (OAC) on Customer AS examples. A solid line and a dotted line is a result of RouteViews and RIPE, respectively.

Figure 4 shows the origin address counts for the two customer ASes during the Slammer worm attack. For most of the time, origin address counts remain the same. However, there is a sudden dip (possibly owing to the worm impact), resulting in address counts of 0 and 3.3% from RouteViews data and 50.9 and 0% from RIPE data. This shows that these customers were totally or partially unreachable during the worm attack. It is interesting to note that two distinct vantage points, RouteViews and RIPE, had different views of the network.

Reachability to a Transit Core AS

For forensic analysis on transit core ASes, we choose three ASes: AS9318, AS4766, and AS3786 that are geographically located close and act as major transit ASes in South Korea. We are interested to see if geographically close ASes show similar behaviors against the same worm attack. Because these ASes provide transit service to their customers, we investigate if there was any significant change in their transit address counts.

Figure 5(a) shows the changes in the level of reachability of AS9318. Contrary to the result of stub ASes, we observe increases in the origin address count. As we mentioned in Section 3.3, an AS can decide to give away some parts of its address space to its customer or neighbor ASes, then these assigned addresses punch holes in the origin address count of a provider. The origin address count successfully reflects such punched holes, and increases in the amount of origin address count reveals the disappearance or shrink of such holes. There could be administrative problems on its customer or neighbor ASes or on its BGP peer links to the neighbor ASes.

The transit address count of AS9318 is reduced to 6.42% of its initial value. AS9318 transits total 805 prefixes on the day prior to the worm advent, however during the attack, only 55 of initial 805 prefixes are transited by AS9318. Interestingly, 708 of 805 prefixes still maintain their reachability through other alternative ASes in different countries from AS9318. On the other hand, the remaining 97 of 805 prefixes completely disappeared from routing table snapshots and lost their reachability as viewed from the vantage point

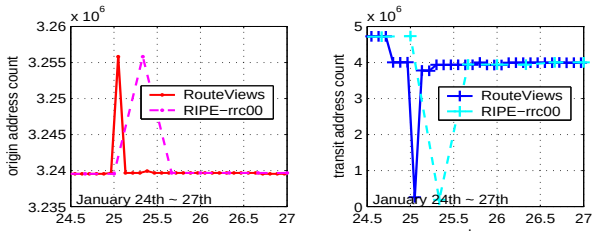

(a) AS9318

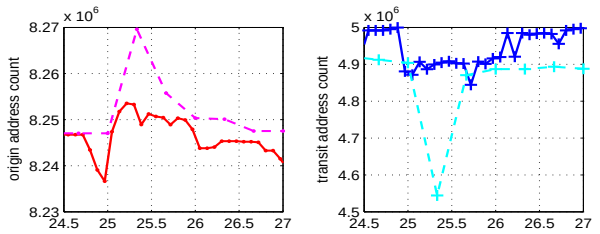

(b) AS4766

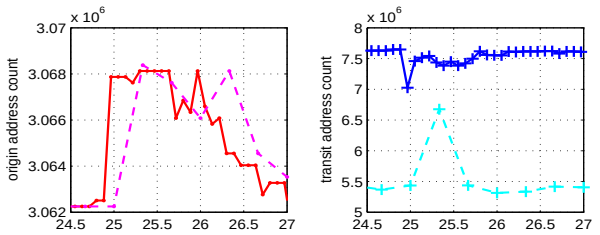

(c) AS3786

Figure 5: Time-series plots of origin and transit address counts

of RouteViews.

From the results of AS9318, we can conclude that the reachability of prefixes owned and announced by AS9318 are preserved during the worm attack. On the other hand, almost all prefixes faced the difficulties to transit through AS9318 as its transit service provider. During the worm attack, AS9318 seems to have lost its role as a transit point and sink to a valley depending on other ASes to transit its prefixes. Also the increases in the amount of reduced-origin address count show that AS9618 experience the loss of reachability to several portion of its customer or neighbor ASes.

The analysis result of AS4766 is presented in Figure 5(b). We observe following differences in comparison to AS9318. The origin address count first decreased and then increased. It means that the number of punched holes or the size of a hole on AS4766 increased owing to the newly and temporarily announced prefixes. There could be an address hijacking caused by BGP misconfiguration or human error faced on the Slammer worm attack.

Finally, Figure 5(c) presents the result for AS3786. It follows the typical pattern of previous results: the origin address count increases and the transit address count decreases. However, the transit address count from RIPE data shows the opposite pattern, increase of 22.8%, to the one as a view of RouteViews. This increase observed in RIPE data is by 195 newly announced prefixes that are not shown on the day prior to the worm attack. As a view of RouteViews, we can see 179 prefixes of these 195 new ones in RIPE has been transited by AS3786 before the worm period. This conflict shows that address counts calculated from different vantage points could have different values.

5. CONCLUDING REMARKS

In this paper, we perform a detailed forensic analysis of routing instability during known security incidents and assess the level of damage in AS reachability. As a way to quantify the damages, we propose the following metrics: the prefix count and the address count. We choose the Slammer worm attack in January, 2003, as a security incident for closer examination, and use BGP data readily available from RouteViews and RIPE to evaluate our methods. We show our metrics are simple, but effective in identifying ASes that experienced much damage.

We hope NSPs (network service providers) will use our metrics to keep track of their customer ISPs and help them when they see a need faced on the security incidents. We believe that our work can show the possibilities of forensic examination on AS reachability with the public BGP routing table snapshots. We also hope that more forensic analysis will follow to quantify damage from cyber attacks better and more accurately.

6. REFERENCES

[1] David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, and Nicholas Weaver, "The spread of the Sapphire/Slammer worm," Tech. Rep., CAIDA, ICSI, Silicon Defense, UC Berkeley EECS and UC San Diego CSE, 2003.

[2] James Aldridge, Daniel Karrenberg, Henk Uijterwaal, and Rene Wilhelm, "Sapphire/Slammer worm impact on internet performance," Tech. Rep., New Project Group/ RIPE NCC, February 2003.

[3] RIPE NCC, "Test Traffic Measurements Service," `http://www.ripe.net/ttm`.

[4] Mohit Lad, Xiaoliang Zhao, Beichuan Zhang, Dan Massey, and Lixia Zhang, "Analysis of BGP update surge during slammer worm attack," in *Proceedings of 6th International Workshop on Distributed Computing (IWDC)*, December 2004.

[5] James H. Cowie, Andy T. Ogielski, BJ premore, Eric A. Smith, and Todd Underwood, "Imapact of the 2003 blackouts on internet communications," Tech. Rep., Renesys Corporation, March 2004.

[6] Zhen Wu, Eric Purpus, and Jun Li, "BGP behavior analysis during the August 2003 blackout," in *Proceedings of IEEE IM*, March 2005.

[7] Lan Wang, Xiaoliang Zhao, Dan Pei, Randy Bush, Daniel Massey, Allison Mankin, S. Felix Wu, and Lixia Zhang, "Observation and analysis of BGP behavior under stress," in *Proceedings of ACM SIGCOMM IMW*, November 2002.

[8] Jun Li, Dejing Dou, Zhen Wu, Shiwoong Kim, and Vikash Agarwal, "An internet routing forensics framework for discovering rules of abnormal bgp events," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, pp. 57–66, 2005.

[9] Geoffrey G. Xie, Jibin Zhan, David A. Maltz, Albert Greenberg, Gili Hjalmtysson, Hui Zhang, and Jennifer Rexford, "On static reachability analysis of ip networks," in *Proceedings of IEEE INFOCOM 2005*, March 2005.

[10] Anja Feldmann, Olaf Maennel, Z. Morley Mao, Arthur Berger, and Bruce Maggas, "Locating Internet routing instrabilities," in *Proceedings of ACM SIGCOMM*, August 2004.

[11] Advanced Network Technology Center and University of Oregon, "The RouteViews project," `http://www.routeviews.org`.

[12] RIPE NCC, "The Routing Information Service," `http://www.ripe.net/ris`.

[13] MRTD, "Multi-threaded Routing Toolkit," `http://www.mrtd.net`.

[14] Lixin Gao, "On inferring autonomous system relationships in the internet," in *Proceedings of IEEE Global Internet*, November 2000.

[15] Lakshminarayanan Subramanian, Sharad Agarwal, Jennifer Rexford, and Randy H. Katz, "Characterizing the internet hierarchy from multiple vantage points," in *Proceedings of IEEE INFOCOM 2002*, June 2002.

[16] IANA, "Autonomous system numbers," `http://www.iana.org/assignments/as-numbers`.

[17] DK Lee and Sue Moon, "Forensic analysis of autonomous system reachability," Tech. Rep. CS-TR-2006-257, Division of Computer Science, Department of EECS, Korea Advanced Institute of Science and Technology, May 2006.