IEICE TRANS. INF. & SYST., VOL.E95–D, NO.1 JANUARY 2012

INVITED PAPER Special Section on Trust, Security and Privacy in Computing and Communication Systems

A Systematic Approach to Evaluating the Trustworthiness of the Internet Inter-Domain Routing Information

Peidong ZHU[†a)], **Huayang CAO**[†], **Wenping DENG**[†], **Kan CHEN**[†], ***and*** **Xiaoqiang WANG**[†], ***Nonmembers***

SUMMARY Various incidents expose the vulnerability and fragility of the Internet inter-domain routing, and highlight the need for further efforts in developing new approaches to evaluating the trustworthiness of routing information. Based on collections of BGP routing information, we disclose a variety of anomalies and malicious attacks and demonstrate their potential impacts on the Internet security. This paper proposes a systematic approach to detecting the anomalies in inter-domain routing, combining effectively spatial-temporal multiple-view method, knowledge-based method, and cooperative verification method, and illustrates how it helps in alleviating the routing threats by taking advantage of various measures. The main contribution of our approach lies on critical techniques including the construction of routing information sets, the design of detection engines, the anomaly verification and the encouragement mechanism for collaboration among ASs. Our approach has been well verified by our Internet Service Provider (ISP) partners and has been shown to be effective in detecting anomalies and attacks in inter-domain routing.

key words: *BGP, inter-domain routing, trustworthiness, security*

1. Introduction

Today's Internet is composed of more than 30,000 independently administered Autonomous Systems (AS), including Internet Service Providers (ISP), universities, and enterprise networks. ASs are coupled by the Border Gateway Protocol (BGP) [1] so that all ASs together are formed into a single globe-spanning entity. Routing in the Internet is inherently complex. It is controlled by diverse policies, decided locally by each Autonomous System (AS), but acting globally across the entire system. Furthermore, it depends on protocols for routing between and within individual ASs, on the router-level topology inside Internet domains, and on the peering structure between ASs. A BGP route includes a list of ASs, namely an AS-PATH, followed by an IP prefix reachable through that AS-PATH. The last AS in the list is commonly referred as the origin AS. For example, if prefix p is associated with an AS-PATH path = $(n_1, n_2, \ldots, n_M)$, then AS n_M should be the origin AS of prefix p.

Unfortunately, various incidents reveal that the inter-domain routing is vulnerable to a variety of attacks due to the lack of validation to the inter-domain routing information [2]. For example, as illustrated by the AS 7007 incident [3] in 1997, a small ISP with AS number 7007, which accidentally de-aggregated and advertised a large portion of the Internet address space, attracting traffic away from the real destinations, disrupted much of the Internet's connectivity for over two hours. Similarly, on February 24, 2008 [4], the Pakistan Telecom (AS 17557) advertised a sub-prefix of YouTube's (AS 36561) assigned network 208.65.153.0/24 to PCCW Global (AS 3491). PCCW then forwarded this announcement to the rest of the Internet, which resulted in the hijacking of YouTube's traffic on a global scale. In another recent incident on April 8, 2010 [5], China Telecom (AS 23724) hijacked more than 37,000 prefixes of the world's routes, where AS 23724 is only a small AS of China Telecom. These incidents demonstrate that Internet routing system is inevitably suffering from bogus advertisement and prefix hijacking [6], [7]. This raises the question *whether an AS can trust the route entry learned from its peering neighbors*.

It is commonly believed that the false routing information was originated from either misconfiguration or malicious attacks. The critical reason for routing failures lies on the lack of effective ways to evaluate the trustworthiness of routing information. Regarding the trustworthiness of routing information, the following problems are coming up: 1) Are the routing elements appearing in a route entry valid or invalid? For example, is the prefix (or AS number) a private or an unallocated identifier? 2) Is the prefix originated by a legitimate AS? In other words, if the AS is authorized to announce the prefix. 3) Is the AS-PATH corresponding to the route entry legitimate, i.e., it is not forged by an AS contained in the path?

There are a plethora of proposed countermeasures [8]–[11] addressing inter-domain routing security problems. However, these mechanisms are limited by their further deployment since they need to change the routing protocol. Anomaly detection [12]–[16] provides an off-line help to network operators to diagnose the routing system without any influence and plays a more important role in network operation and diagnosis. Such proposals can be categorized into the following types: (1) RIR/IRR registry based approach, such as [19]. (2) Methods relying on routing history information and/or registry information, such as MyASN [15], IAR (Internet Alert Registry) [14], [24], PHAS [12], [13], Cyclops [16]. (3) Other methods such as [17] and [18]. However, there are several constraints in using these methods. On the one hand, neither RIR/IRR registry information nor routing history information is complete and up-to-date [19], hence it is difficult to maintain a knowledge base for the anomaly detection. On the other hand, using these methods independently can only detect

Manuscript received March 23, 2011.
Manuscript revised June 24, 2011.

[†]The authors are with School of Computer, National University of Defense Technology, Changsha 410073, China.
a) E-mail: pdzhu@nudt.edu.cn

DOI: 10.1587/transinf.E95.D.20

limited kinds of routing anomalies.

This paper proposes a systematic approach to detecting the anomalies in inter-domain routing, combining effectively spatial-temporal multiple-view method, knowledge-based method, and cooperative verification method. Taking and combining advantage of diverse measures enables us to alleviate the routing threats in various scenarios. Our approach has been well verified by our Internet Service Provider (ISP) partners and has been shown to be effective in detecting anomalies and attacks in inter-domain routing.

The remainder of this paper is structured as follows. Section 2 presents the definition of routing information trustworthiness and typical harms by untrustworthy routing information. Section 3 describes our systematic method for evaluating the trustworthiness of routing information, followed by Sect. 4 with the evaluation for our methods. We finally conclude our work in Sect. 5.

2. Preliminaries

In this section, we start by describing the trustworthiness of routing information. Afterwards we will focus on untrustworthy routing information and potential routing attacks.

2.1 Trustworthiness of Routing Information

Internet routing at global scale is implemented and performed by BGP at the level of IP prefixes announced by Autonomous System (AS). Generally, an IP prefix can be originated by a unique AS-the origin AS which is authorized by the prefix owner. Given a BGP route entry $r = \langle \text{prefix}, \text{path} \rangle$, where prefix is an IP block, consisting of an IP address followed by slash and then the length of the prefix, e.g., 59.42.0.0/16; path = $\langle n_1, n_2, \ldots, n_i \rangle$ is a sequence of ASs, e.g., path = ⟨701, 1239, 4134⟩, and n_i(4134) is the origin AS of the prefix (59.42.0.0/16). Drawing further on the inter-domain route entry, we decompose the trustworthiness of inter-domain routing information as the following aspects:

The validity of the IP prefix p: Let $\mu_1(p)$ denote the validity of a given prefix p. p is considered to be valid only if it is a prefix which is actually in-use publicly, neither a special-use prefix such as private one, nor an unallocated prefix. However, the trustworthiness of a prefix can change to legitimacy from illegitimacy, e.g., just after it is allocated.

The validity of each AS number n in the AS-PATH: Let $\mu_2(n)$ denote the validity of a given AS number n. AS n is considered to be legitimate only if it is an AS which can be used in global Internet routing, i.e., neither a private nor an unallocated AS. Like the prefix, the validity of an AS also can change from illegitimacy to legitimacy if the AS is newly allocated.

The trustworthy degree of the prefix-AS mapping (p, o)**:** Let μ_3(p, o) denote the trustworthy degree of a given prefix-AS mapping (p, o). It indicates the trustworthiness that AS o is the authenticated origin AS of the prefix p, i.e., o is the real origin AS of p.

The trustworthy degree of the AS-AS link (as_1, as_2)**:** Let $\mu_4(\text{as}_1, \text{as}_2)$ denote the trustworthy degree of a given AS-AS link (as_1, as_2). It indicates the trustworthiness that as_1 and as_2 have a real connection due to their commercial relationship.

The trustworthy degree of AS-PATH path = (n_1, $n_2, \ldots, n_i$)**:** Let μ_5(path) denote the trustworthy degree of a given AS-PATH. It indicates the trustworthiness that the path is the real one along which the traffic will be forwarded from AS n_i to n_1, thus each AS hop of the path has the consistent sub path with $(n_1, n_2, \ldots, n_i)$. For example, from the local view of AS $n_j (1 \leq j \leq i)$, n_j has a corresponding route with an AS-PATH $(n_{j+1}, \ldots, n_i)$.

2.2 Untrustworthy Routing Information

Invalid IP prefixes and AS numbers could be intentionally used for spam [20] or phishing attacks [21]. From the control plane of the routing system, this kind of anomaly can be detected and filtered easily according to the allocating records of the Internet resources. However, there could be more sophisticated attacks beyond the untrustworthy routing information.

Untrustworthy (prefix, origin AS) mapping vs. prefix hijacking: Formally, a prefix can only be announced by a unique AS. However, the Classless Inter-Domain Routing (CIDR) allows different ASs to announce covering IP spaces yet these IP prefixes have potential conflicts, e.g., less or more specific prefix announcement. For the CIDR in BGP, IP prefix is an IP address followed by a slash and the prefix length used for the network part. Routers that have routes to such overlapping IP prefixes will choose the unique route according to the Longest Prefix Match Rule.

In Internet routing, a prefix is treated as a different one from its super-prefix or sub-prefix. In consequence, even though an AS has originated a prefix p, e.g., p = 100.100.100.0/24, other ASs can still announce a super-prefix (e.g., 100.100.0.0/16) or a sub-prefix (e.g., 100.100.100.128/25) of p. The less/more specific prefix announcement will also be propagated to the whole Internet and will consequently have a global impact, for incidence, the YouTube hijacking [4]. In such cases, although some filtering mechanisms might block the propagation of prefix hijacking occasionally, the prefix hijacking would be propagated to most of the ASs on a global scale. The real impact of prefix hijacking mainly depends on routing policies such as preference of customer route, preference of shortest path, etc.

Untrustworthy AS-PATH vs. path forgery: There are sophisticated AS path forgery tricks as described in [24]. In-

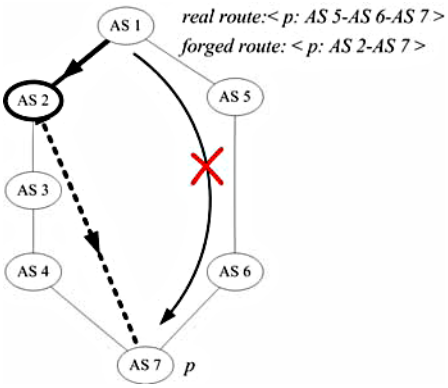

Fig. 1 An exmaple for AS-PATH forgery.

stead of originating another AS's prefix directly, the attacker announces a route with a shorter path in order to attract the traffic from others. As shown in Fig. 1, AS2 has a reachable route ⟨AS3-AS4-AS7⟩ to the destination p, but it is not the best one because AS 5 has a shorter route ⟨AS6-AS7⟩. AS2 forges a new route to p with a shorter AS path ⟨AS7⟩ and advertises this route to AS1. After that, AS1 might choose the forged route as the new forwarding path, so the traffic will switch to the route via AS2. As an intermediate node, AS2 can eavesdrop or modify the traffic from AS1 destined to p. Compared to prefix hijacking, AS path forgery is even more difficult to detect.

3. Approach

In this section, we will describe our systematic approach to evaluating the trustworthiness of routing information. The systematic approach consists of the following methods: the method based on registry information, the method based on history routing information, the spatial multiple-view method, the source-based method, and the reputation-based method. Each of these methods has its own advantage and disadvantage, however, our systematic approach integrates all methods together so that it can be better used in evaluating the trustworthiness of routing information and detecting routing anomalies.

3.1 Based on Registry Information

IP addresses and AS numbers are basic resources for Internet routing. They are allocated by the Internet Corporation for Assigned Names and Numbers (ICANN) to the Regional Internet Registries (RIR), including ARIN, RIPE NCC, APNIC, LACNIC, and AfricNIC. RIRs then the network resources to Local Internet Registries (LIR) and organizations such as Internet Service Providers.

As for the allocation, some special IP addresses (e.g., the private addresses, multicast addresses) and special AS numbers (e.g., the private AS numbers) are not allocated to any organization. If such elements appear in a route entry, then the route entry is considered to be untrustworthy for the invalid element. Moreover, regarding the normal IP addresses and AS numbers, only a part of them have been allocated. According to the information from RIR, we can confirm whether a given AS or an IP prefix has been allocated or not.

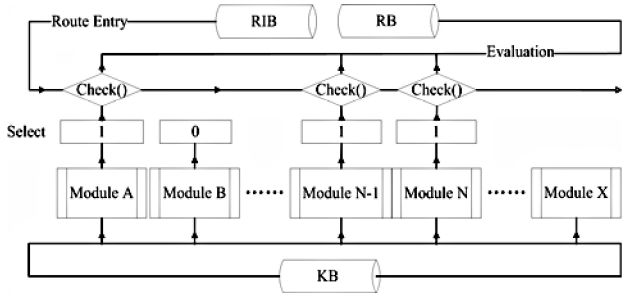

Fig. 2 The detection on routing information.

Based on the allocating information and registry information, we can construct the basic knowledge base. Drawing further on this, routing anomaly is defined as the violation to the routing knowledge base. The detection is proposed for diagnosing the validity of route entries. As described above, the unallocated and special-use prefixes and AS numbers are considered to be invalid. Besides, according to the knowledge extracted from RIRs, part of the untrustworthy prefix-AS mappings can also be verified. Moreover, operators can set additional criteria for valid and invalid route entries. In the later process, these valid entries will be chosen as right routing information, and invalid ones will be discarded and reported. It should be noticed that, due to the restriction of monitor scope and lack of AS-AS link knowledge, we are not able to verify the validity of AS-PATH and detect the conflicts between AS-PATHs. The detection process is as follows.

In Fig. 2, route entries flow out from RIB (Routing Information Base), and receive checks from detection modules. The Select module controls the detection function, and KB (Knowledge Base) provides necessary knowledge for supporting such detection. Finally, the evaluation results of route entries are stored in the RB (Result Base).

However, the current global registry of prefix ownership and routing policies maintained by RIR has been widely considered to be outdated, incomplete, and inaccurate. For instance, RIPE NCC is considered to be the best maintained RIR, however, only 73% of its prefix-AS registry information can be strongly validated in 2004 [19].

3.2 Based on History Routing Information

In Sect. 3.1, we described our knowledge-based anomaly detection model. The exactness of this detection excessively depends on the accuracy and completeness of the knowledge base. As described above, we have little access to the prefix-AS mapping and AS-AS links information, which is important for the detection of resource ownership and validity of AS-PATH. In this subsection, we will address this problem and propose a temporal-based trustworthiness evaluation on prefix-AS mappings, AS-AS links and AS-PATHs.

In Internet routing, origin ASs of prefixes and AS links change for a variety of reasons. Some of them are legal and normal, such as routing aggregation and de-aggregation, traffic engineering, reallocation from providers to customers, new commercial contract, ISP reorganization or bankruptcy, and so on But there are also some reasons illegal or abnormal, such as equipment failure, misconfiguration, routing instability, and malicious attacks. Previous work has revealed the two following characteristics of the Internet routing announcements:

(1) Most of the Internet routing announcements are very stable and the vast majority of BGP instability stems from a small number of unpopular destinations [22]. We analyze the stability of the prefix-AS mappings from Route Views [23] from 01/01/2008 to 31/12/2008, and find that less than 1% of the route entries have changes within two days, yet more than 90% of them are absolutely stable for longer than 15 days. Authors of [24] also pointed out that the majority of BGP route entries are long lasting.

(2) As demonstrated in [25] and [26], most of the misconfigured and malicious announcements are short-lived because they may disturb the running Internet and will be soon discovered.

Let τ be the current day to evaluate, then the sliding history window is defined as a time period from a starting day (s) to the evaluating day (τ). Here we mark the sliding history window as $[s, \tau]$. According to (1) and (2), we define proposition 1 as follows:

Proposition 1: *A prefix-AS mapping (an AS-AS link, or an AS-PATH, in the following, we take prefix-AS mapping as our example) which is stable during the sliding history window $[s, \tau]$ in the Internet routing is considered to be trustworthy at the evaluating time τ.*

According to proposition 1, the trustworthy degree of a prefix-AS mapping depends on its stability in a sliding history window. In this paper, we choose 10 days as the window size. For each prefix-AS mapping (AS link) appearing in the history window, we keep an active sequence (v) to state its stability. If a prefix-AS mapping (an AS-PATH) is active at time τ, then the τ-th element of its active sequence is $v(\tau) = 1$, otherwise, $v(\tau) = 0$ (inactive). For example, given a mapping $\langle p, n\rangle$, where p = 207.67.209.0/24, n = 174, s = 20080101, τ = 20080110, if it has an active sequence "10111 01111", where the 2nd and 6th elements are '0', then we can say that $\langle p, n\rangle$ did not appear in the BGP RIBs on January 2 and January 6 in 2008.

Proposition 1 states that the more stably a prefix-AS mapping (AS link) exists, the more trustworthy it should be. Since we use a 0-1 sequence, the K-bit active sequence $v(v = v(1)v(2) \ldots\ldots v(K))$, to record the continuous appearing status of each prefix-AS mapping (AS link) in the snapshots during the history window, it is the best indicator of the stability of the associated prefix-AS mapping (AS link).

If v = "111 . . . 11", then the prefix-AS mapping is completely stable during the history window, if v = "100 . . . 00", then the prefix-AS mapping only appeared at the first day of the window, and if v = "1010 . . . 10", then the prefix-AS mapping is flapping and instable in the history window.

Table 1 Examples for calculating μ according to $f(v)$.

v (K=10)	μ computation	μ
1000000000	1 / 10	0.1
0000011111	5 /10	0.5
1111111111	10 / 10	1

In light of Proposition 1, we seek to quantify the trustworthiness of prefix-AS mappings from their stability. Intuitively, we quantify the trustworthiness of the mapping (p, n) by counting the frequency of its appearance in the window. Thus, we propose the basic stability-trustworthiness model as follows.

$$f(v) = \sum_{i=1}^{K} v(i) \tag{1}$$

where $v(i)$ is the i-th element of the active sequence v of a given prefix-AS in the history window. $f(v)$ obtains its maximum (minimum) value $K(0)$ when all the bits of v are '1' ('0'). We normalize the trustworthiness value by mapping the value to range of [0, 1] as follows.

$$\mu(p, n) = f(v)/K \tag{2}$$

According to the basic stability-trustworthiness model, $f(v) = 1$ if v = "11 . . . 11", and $f(v) = 0$ if v = "00 . . . 00". Hence the more stable the prefix-AS mapping is, the higher $f(v)$ will be. Table 1 gives examples for calculating $f(v)$.

It can be inferred from proposition 1 that μ indicates the trustworthiness value of route entries. We extract prefix-AS mapping and AS-AS links from BGP RIBs, and calculate trustworthiness value of them. Therefore, we can decide whether a prefix-AS mapping (p, n) is trustworthy or not according to a given threshold α, i.e., if $\mu(p, n) > \alpha$, then (p, n) is trustworthy, otherwise not.

However, the view captured from limited vantage points is inherently incomplete [27]. This clearly points out that this method is inevitably limited by the completeness of the collected routing information from BGP RIBs.

3.3 Spatial Multiple-View Method

The preceding methods are able to diagnose the routing information for anomalies, and to evaluate the trustworthiness of routing information. However, as mentioned in previous sections, we have already known that both of the two proposed methods have their own limitations. In this section, we propose a method that can verify the trustworthiness of routing information (especially for AS-PATHs) from multiple vantage points, i.e., it is a multiple-view method.

It's a natural way that we obtain RIBs from relevant ASs and check the consistency of AS-PATH. Yet some ISPs are not willing to provide their routing information, due to confidential and commercial reasons. Here we propose a spatial multiple-view detection measure to verify untrustworthy routing information. As shown in Fig. 3, we set a

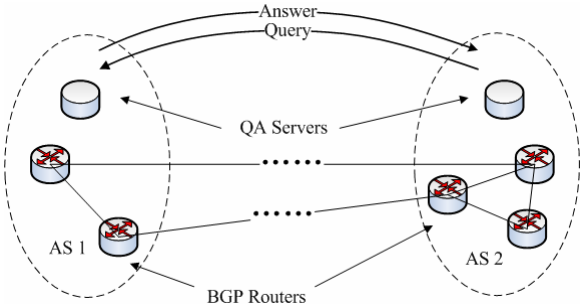

Fig. 3 Deployment of QA servers.

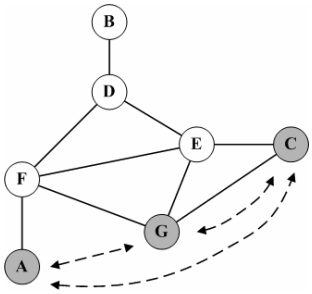

Fig. 4 Cooperative monitoring technique.

Q (Query) A (Answer) server in each AS. When receiving an update, the QA server will send a query to other relevant QAs to request confirmation, and then the latter ones check their local RIBs and send their verifying reply to the query.

Considering a route entry $(p, \langle n_1, n_2, \ldots, n_M \rangle)$, QA in AS n_l will construct a query as (query: $p, \langle n_1, n_2, \ldots, n_M \rangle$) to QA in AS n_i. The latter QA only needs to answer "Y" if there is an exported route entry with $\langle p, \langle n_{i+1}, \ldots, n_M \rangle \rangle$ in the RIB of AS n_i, otherwise "N". If QA in $\mathrm{ISP_A}$ answers "Y" for a query, it seems that the QA reveals some privacy routing information to others. However, it should be noticed that such a route entry has already been advertised, and it is no longer privacy information. Then, if this QA answers "N", $\mathrm{ISP_A}$ leaks nothing but helps another QA to confirm the inconsistency. Obviously, such mechanism can be implemented in an incremental way, but the verification effectiveness surely correlates with the amount of QAs. Besides the detection of AS-PATH consistency, QA servers can also manage various queries without or just with a little leaking of the privacy information from ISPs.

3.4 Source-Based Method

When an ISP receives a route, sometimes it is difficult to identify the untrustworthy mapping of prefix to the claimed origin AS from the registry information. Therefore, the ability of traditional receiver-based methods is inherently limited. In fact, real origin AS, or the owner of the prefix, has complete knowledge for diagnosing the validity of routes which are relevant to themselves, and will take necessary actions to solve the hijacking problem if its prefix is forged. However, it is difficult for them to monitor all of the routing information, because the bogus routes can be propagated beyond its view scope. In this subsection, we propose a cooperative and deployable method, called Co-Monitor, to evaluate the trustworthiness of prefix-AS mapping, namely to identify prefix hijacking. Since the validity of prefix-AS mapping is finally decided by the origin AS of the prefix, we call this method as

The source-based verification technique emphasizes that each AS takes charge of its prefixes and distributes the monitoring responsibility of its prefixes among all of the participants: [28].

Co-Monitor consists of four steps. First, each participating AS constructs a *prefix-to-origin* table mapping prefixes to ASs. Generally, a participant's table contains the mapping of its own prefixes to itself at least. Second, these ASs exchange respective mappings with each other. Third, these ASs locally keep an eye on routes announced by neighboring ASs. Fourth, if a participating AS detects an event that the prefix origin of a route is inconsistent with the previously received mappings, a notification is delivered to the mapping originator AS. By this means, ASs can extend the capabilities of monitoring their prefixes from different vantages on the Internet, and can determine whether their prefixes are hijacked by other in real-time. As shown in Fig. 4, there are three ASs, *A*, *C* and *G*, coordinated to monitoring respective prefixes with the help of the other ASs. If AS C detects a disagreement mapping from AS A, it will notify AS *A* of this anomaly.

3.5 Reputation-Based Method

We have discussed four approaches for evaluating and verifying the trustworthiness of routing information. These approaches can be used to detect abnormal route entries. However, they have no further countermeasures in preventing malicious ASs to generate untrustworthy route again, i.e., none of them has proposed punishment mechanism for the ASs with bad behaviors. In this section, we propose a reputation mechanism which encourages the collaboration among ASs to enhance the trustworthy of routing information.

Our proposal supports incremental deployment and does not need modification on BGP protocol. Every AS has a reputation agent (RA) and a routing monitor (M). Reputation agent calculates and stores reputation of BGP neighbors. A routing monitor creates an iBGP session with other border routers in the same AS to collect and analyze the BGP update message. Reputation agent and routing monitor run on an independent reputation server with sufficient CPU and memory resources to process all BGP update messages in real time. When routing monitor detects new bogus route using the methods described in above subsections, it triggers RA to recalculate the direct evaluation. Figure 5 gives an illustration of reputation system deployment.

The reputation value can be used in two respects. First, it can be used in routing decision process. An AS prefers the routes from neighbors with highest reputation evalua-

Table 2 Capabilities of our methods.

detection capability / method	Invalid identifiers (Prefix or AS)	Untrustworthy AS-PATHs	Untrustworthy prefix-AS mappings	Untrustworthy AS-AS links
knowledge-based	**Y**	**N**	**Y** (partially)	**Y** (partially)
temporal-based	N	**Y** (partially)	**Y** (partially)	**Y** (partially)
s- multi-view	N	**Y** (partially)	**Y** (partially)	N
source-based	N	N	**Y** (partially)	N

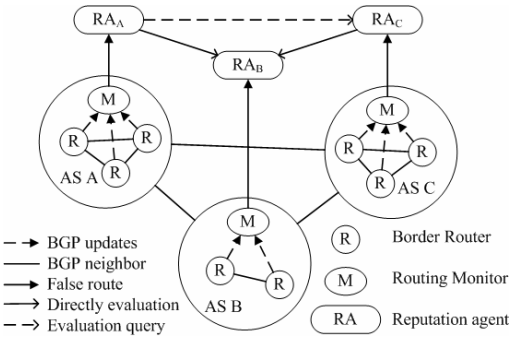

Fig. 5 An illustration of reputation system deployment.

tion. Second, it can also be utilized as an metric to evaluate the trustworthiness of the routing information. When other evaluation methods fail to identify the forged one from multiple claimed origins, the one from the AS with smaller reputation value is most suspectable.

4. Evaluation

In this section, we will describe how to combine our proposed technologies as a whole to address the trust problem in routing system.

4.1 Systematic Using of the Proposed Methods

The preceding methods can evaluate the trustworthiness and diagnose anomalies of routing information in diverse scenarios. Here, we give a sum-up of these methods.

As for the validity (see Sect. 1) of routing elements, especially the evaluation of the validity of the IP prefix, AS number and prefix-AS mapping, the knowledge-based anomaly detection method can work well. The RIRs are often used to confirm the legality of resources usage. However, the allocation records of Internet resources in RIRs are not accurate, thus we may face some degree of false ratio in using this method to evaluate the validity of route entries.

To distinguish whether there is an unwelcome AS in an AS-PATH, we can also use the knowledge-based approach. As we can extract accurate Country-AS mapping and Country-prefix mapping from RIRs, those ASs belonging to hostile areas can be easily distinguished. Also due to the accuracy limitation of prefix-AS and AS-organization mappings in RIRs, we are not able to tell all unwelcome ASs from normal ones, but we can improve the accuracy by digging out these relationships in RIRs.

Regarding the validity of AS-PATH and prefix announcement, sometimes we don't know whether there exist the AS-AS links in Internet topology and prefix-AS mappings. Here, the temporal-based trustworthiness evaluation model is proposed. Previous studies have shown that the majority of BGP route entries are long lasting in the routing system [22], [24], then we treat the short-lived route entries as bogus ones. Obviously, this method can not be accurate in all situations, but it is relatively accurate in most cases and can be integrated into the monitoring system for routing security.

The verification of routing consistency is implemented by the spatial multiple-view detection method. This method can guarantee the consistency among ASs by utilizing a Query and Answer mechanism. Participants can verify the consistency of received route entries by sending a query to relevant ASs. Since the "answer" message is only either be "Yes" or "No", it reveals little routing information of participants, and it is more likely to be adopted by ISPs.

Besides the above receiver-based evaluation models, we also propose a source-based route verification model as a complement. Because prefix owners have the full knowledge of the usage of their own Internet resources, such a method can obtain an accurate estimation of prefix-AS mappings. However, the main obstacle is the difficulty of wide deployment of route monitors for the prefix owners. We propose a cooperative monitoring mechanism for ASs to overcome the monitoring scope problem. Table 2 shows the applicability of our methods addressing diverse trustworthy related problems.

It can be seen that every problem is addressed by one or more measures, and the "**Y** (partially)" marks are made due to the incomplete knowledge base. When people want to evaluate trustworthiness of routing information in a given aspect, they can choose the proper methods according to Table 2.

Finally, the reputation-based measure is proposed to encourage the collaboration among ASs to enhance the trustworthy of routing information, and sometimes it helps tell the route trustworthiness by the reputations of upstream neighbor and route originator.

4.2 Based on History Routing Information

The temporal-based method is applicable to evaluate the trustworthiness of prefix-AS mappings and AS-AS links.

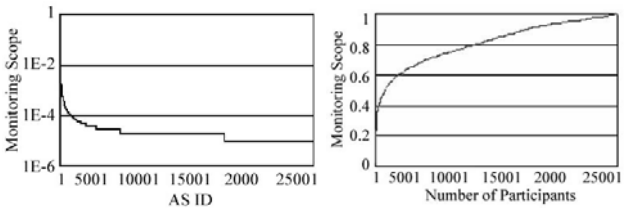

Fig. 6 Monitoring scope without (a) and with (b) Co-Monitor.

Here we will show the effectiveness of our method by applying it to evaluating the trustworthiness of known incidents.

On February 24, 2008, Pakistan Telecom (AS 17557) started an unauthorized announcement of the prefix 208.65.153.0/24 which belongs to YouTube (AS 36561). One of Pakistan Telecom's upstream providers, PCCW Global (AS 3491) forwarded this announcement to the rest of the Internet, which resulted in the hijacking of YouTube traffic on a global scale.

To evaluate the trustworthiness of the prefix-AS mapping (208.65.153.0/24, 17557) on February 24, 2008 (τ = 24/02/2008), the sliding history window W = [15/02/2008, 24/02/2008]. We firstly build the basic trustworthy prefix-AS mapping set from the snapshots of the last ten days, and then get the trustworthiness of the mapping.

The trustworthiness is calculated with Eq. (2), and the value is 0.25. We can say that the mapping of (208.65.153.0/24, 17557) was untrustworthy having the very low trustworthiness value on February 24, 2008. It conflicted with the mapping (208.65.152.0/22, 36561) who has a trustworthy value 1 in the basic trustworthy set.

4.3 Source-Based Method

As we mentioned, the source-based method has the advantage of verifying route announcement with a relatively high accuracy. Here the critical technique is how to extend the view scope of route announcer for verification. We will evaluate the effectiveness of our cooperative mechanism in extending monitor scope.

To demonstrate the benefit of the Co-Monitor in the context of the Internet, we select a BGP snapshot from RouteViews on June 20, 2007 [23]. The Internet topology consists of 25699 ASs. We sort these ASs by their node degrees, and assign an ID (from 1 to 25699) to every AS. Figure 6 (a) shows the monitoring scope (ratio of the amount of monitored ASs to 25699) of every AS without Co-Monitor, mostly less than 10^{-4} (close to zero).

We assume that ASs join the Co-Monitor architecture in turn according to their IDs. The experimental results are depicted in Fig. 6 (b). The results demonstrate clearly that the Co-Monitor extends the monitoring scope of an AS remarkably. For example, if the top 10 ASs joined, their monitoring scope is 12.9% of the Internet. More importantly, because of the power-law property of the Internet, the top 916 ASs (less than 3.6% of 25699) can monitor the 50% range of the Internet. Evidently, the benefit to participants is larger than non-participants. Therefore, an AS should have much incentive to join the Co-Monitor architecture. The incentive mechanism can be described as "I work for everyone and everyone works for me".

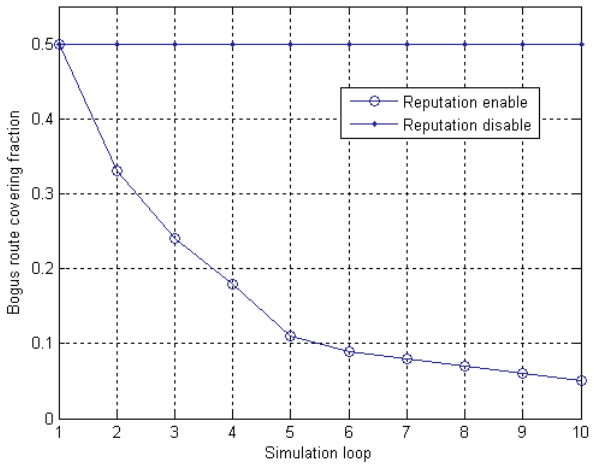

Fig. 7 Suppression of bogus route propagation.

4.4 Reputation-Based Method

The reputation-based method is used not only for verifying the trustworthiness of route information, but also encouraging the announcement of trustworthy routes. Here we evaluate the effect before and after the deployment of our method.

Our simulation assumes that an AS can detect the bogus routes received from neighbor nodes after each loop. We study two scenarios: one with reputation evaluation enabled and another, disabled. It can be seen from Fig. 7 that when reputation evaluation mechanism is not enabled, ASs cannot distinguish which route information is more trustworthy and the bogus route covering fraction is constant. The value of the constant depends on the hijacker selection, AS topology and AS relationships. After the reputation mechanism being enabled, the propagation of bogus route information is effectively restrained.

4.5 Implementation in ISPs' Network

We have integrated all our proposed methods into the "RouSSeau" routing monitoring system [29]. RouSSeau consists of components of routing information gathering module, knowledge base and detection engines, security situation query and visualization modules.

RouSSeau is deployed at backbone networks of several ISPs' in a distributed mode. The data collecting module gathers routing tables and update packets from routers. Notice that the detected routing anomalies do not necessarily occur inside these ISPs' networks since routing events in other part of Internet routing can also spread over these ISPs as well. Table 3 shows the results of a detection of one-day duration.

Table 3 Part of routing evaluation results.

Invalid identifiers	Untrustworthy AS-PATHs	Untrustworthy prefix-AS mappings	Untrustworthy AS-AS links
116	0	18326	51

As illustrated in Table 3, there is a significant amount of "Untrustworthy prefix-AS mappings" routing anomalies. It originates from a large scale BGP hijack incident, since tremendous prefix-AS mappings are evaluated as untrustworthy. Intuitively, the reason for "Untrustworthy AS-AS links" may also be the routing oscillation triggered by this incident. As for the result of "Invalid identifiers", it is due to the occurrence of many private AS numbers in AS-PATHs.

5. Conclusion

In this paper, we have proposed a systematic approach to detecting the anomalies in inter-domain routing, combining effectively spatial-temporal multiple-view method, knowledge-based method, and cooperative verification method, and illustrated how it helps in alleviating the routing threats by taking advantage of comprehensive measures in multiple network planes. Our approach has been well verified by our Internet Service Provider (ISP) partners and has been shown to be effective in detecting anomalies and attacks in inter-domain routing.

Acknowledgements

This work is supported by the project of National Natural Science Foundation of China (NO.60873214).

References

[1] Y. Rekhter, T. Li, and S. Hares, A Border Gateway Protocol 4 (BGP-4), RFC 4271, Jan. 2006.
[2] "An architecture for BGP countermeasures," BBN Rep. 8271, Nov. 1997.
[3] S. Misel, "Wow, as7007," http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html
[4] RIPE NCC, "YouTube Hijacking," http://www.ripe.net/news/study-youtube-hijacking.html, 2008.
[5] BGPmon, "China Telecom Hijacking," http://bgpmon.net/ blog/?p=282, 2010.
[6] H. Ballani, P. Francis, and X. Zhang, "A study of prefix hijacking and interception in the Internet," Proc. SIGCOMM, 2007.
[7] Z. Zhang, Y. Zhang, Y.C. Hu, Z.M. Mao, and R. Bush, "Ispy: Detecting IP prefix hijacking on my own," Proc. SIGCOMM, 2008.
[8] S. Kent, C. Lynn, and K. Seo, "Design and analysis of the secure border gateway protocol (S-BGP)," Proc. DISCEX'00, 2000.
[9] S. Goldberg, S. Halevi, A.D. Jaggard, V. Ramachandran, and R.N. Wright, "Rationality and traffic attraction: Incentives for honest path announcements in BGP," Proc. ACM SIGCOMM, 2008.
[10] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, "How secure are secure inter-domain routing protocols?," Proc. ACM SIGCOMM, 2010.
[11] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin, "Working around BGP: An incremental approach to improving security and accuracy in interdomain routing," Proc. NDSS, vol.3, pp.75–85, 2003.
[12] PHAS: Prefix Hijack Alert System. http://phas.netsec.colostate.edu/stat.html
[13] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A prefix hijack alert system," Proc. USENIX Security Symposium, Aug. 2006.
[14] Internet Alert Registry. http://iar.cs.unm.edu/index.php
[15] Routing Information Service MyASN. http://www.ris.ripe.net/myasn.html
[16] Cyclops. http://cyclops.cs.ucla.edu/
[17] P.D. Zhu, X. Liu, W.P. Deng, and H.Y. Cao, "Cooperative detection of Internet prefix hijacking," J. Internet Technology, vol.11, no.1, pp.33–45, 2010.
[18] N. Hu, P.D. Zhu, and P. Zou, "Reputation mechanism for inter-domain routing security management," Proc. International Conference on Computer and Information Technology, 2009.
[19] G. Siganos and M. Faloutsos, "Neighborhood watch for Internet routing: Can we improve the robustness of Internet routing today?," Proc. IEEE INFOCOM, 2007.
[20] A. Ramachandran and N. Feamster, "Understanding the network-level behavior of spammers," Proc. ACM SIGCOMM, 2006.
[21] O. Nordstrom and C. Dovrolis, "Beware of BGP attacks," ACM SIGCOMM Comput. Commun. Rev, vol.34, 2004.
[22] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "BGP routing stability of popular destinations," Proc. ACM SIGCOMM Internet Measurement Workshop, Nov. 2002.
[23] Route Views. http://www.routeviews.org/
[24] J, Karlin, S, Forrest, and J. Rexford, "Pretty good BGP: Improving BGP by cautiously adopting routes," Proc. ICNP, Nov. 2006.
[25] P. Boothe, J. Hiebert, and R. Bush, "Short-lived prefix hijacking on the Internet," NANOG 36 meeting, 2006.
[26] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," Proc. ACM SIGCOMM, 2002.
[27] R. Bush, O. Maennel, M. Roughan, and S. Uhlig, "Internet optometry: Assessing the broken glasses in Internet reachability," Proc. Internet Measurement Conference, 2009.
[28] C. Prehofer and C. Bettstetter, "Self-organization in communication networks: Principles and design paradigms," IEEE Commun. Mag., pp.78–85, 2005.
[29] W.P. Deng, P.D. Zhu, and X.C. Lu, "ROUSSEAU: A monitoring system for inter-domain routing security," Proc. Communication Networks and Services Research Conference, IEEE Computer Society, pp.255–262, 2008.

Peidong Zhu is a professor with School of Computer Science of National University of Defense Technology (NUDT), China. He received his PhD degree in computer science from NUDT in 1999. From December 2008 to December 2009, he was the visiting professor at St Francis Xavier University, Canada. His research interests include network routing, network security and architecture design of the Internet and various wireless networks. He is a member of the IEEE.

Huayang Cao received his B.S. degree in Network Engieering and M.S. degree in Computer Science from Department of Computer Science, National University of Defense Technology (NUDT), Changsha, Hunan, China, in 2007 and 2009 respectively. He is now a PhD student of NUDT. His research interests include Future Internet and routing security.

Wenping Deng received his B.S. and M.S. degrees in Computer Science from Department of Computer Science, National University of Defense Technology (NUDT), Changsha, Hunan, China, in 2004 and 2006 respectively. He is now a PhD student of NUDT. He was a visiting scholar to the Communication Systems Research Group (CSG) of ETH Zurich, Switzerland, from November 2008 to November 2009. His research interests include Internet routing, routing security, and resilient network.

Kan Chen is a PhD student in School of Computer, National University of Defense Technology (NUDT), Changsha, Hunan, China. He received his B.S. degree and M.S. degree in Computer Network Engineering from NUDT, in 2007 and 2010 respectively. His research interests include Internet routing and Internet security.

Xiaoqiang Wang is a PhD student in School of Computer, National University of Defense Technology (NUDT), China. He received his B.S. degree in Computer Network Engineering from NUDT in 2006. During December 2009 and December 2010, he was a visiting scholar at IP Networking Lab (INL) of University catholique de Louvain (UCL), Belgium. His research interests include future Internet routing and routing security. He is a student member of IEEE.