

Measuring BGP AS Path Looping (BAPL) and Private AS Number Leaking (PANL)

Shenglin Zhang, Ying Liu*, Dan Pei, and Baojun Liu

Abstract: As a path vector protocol, Border Gateway Protocol (BGP) messages contain an entire Autonomous System (AS) path to each destination for breaking arbitrary long AS path loops. However, after observing the global routing data from RouteViews, we find that BGP AS Path Looping (BAPL) behavior does occur and in fact can lead to multi-AS forwarding loops in both IPv4 and IPv6. The number and ratio of BAPLs in IPv4 and IPv6 on a daily basis from August 1, 2011 to August 31, 2015 are analyzed. Moreover, the distribution of BAPLs among duration and loop length in IPv4 and IPv6 are also studied. Several possible explanations for BAPL are discussed in this paper. Private AS Number Leaking (PANL) has contributed to 0.20% of BAPLs in IPv4, and at least 1.76% of BAPLs in IPv4 were attributed to faulty configurations and malicious attacks. Valid explanations, including networks of multinational companies, preventing particular AS from accepting routes, also can lead to BAPLs. Motivated by the large number of PANLs that contribute to BAPLs, we also study the number and the ratio of PANLs per day in the 1492 days. The distribution of the private AS numbers in all of the PANLs is concentrated, and most of them are located in the source of the AS paths. The majority of BAPLs resulted from PANLs endure less than one day, and the number of BAPLs which are caused by two or more leaked private ASes are much larger than that of BAPLs which are caused by one leaked private AS. We explain for this phenomenon and give some advices for the operators of ASes.

Key words: forwarding loops; BGP AS path; private AS number; RouteViews; traceroute

1 Introduction

The Internet consists of thousands of Autonomous Systems (ASes) that are defined as a connected group of one or more IP prefixes governed by a specific and clearly defined routing policy^[1]. At least one intra-domain routing protocol is deployed in an AS to

optimize routing within the domain, such as OSPF^[2], RIP^[3], and IS-IS^[4]. The reachability information among ASes can be exchanged with the inter-domain protocol called Border Gateway Protocol (BGP)^[5]. Each BGP route contains an AS path attribute that lists the path of ASes used to reach the prefix.

BGP is supposed to eliminate path looping. When an AS receives a BGP routing update, it will check whether the AS path attribute contains its own AS number. If so, it will discard this BGP routing message immediately to break the AS path loops. As described in RFC 4271^[5], “this information (the AS path attribute) is sufficient to construct a graph of AS connectivity from which routing loops may be pruned.”

Despite BGP’s design intention of preventing AS path loops, previous research has proved the occurrence of BGP AS Path Looping (BAPL)^[6–8]. A BAPL occurs if there is a loop exists in the AS path attribute.

• Shenglin Zhang, Ying Liu, and Baojun Liu are with Institute for Network Sciences and Cyberspace, Tsinghua University, Beijing 100084, China. E-mail: slzhangsd@gmail.com; liuying@cernet.edu.cn; BJLiu0@gmail.com.

• Dan Pei is with Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China. E-mail: peidan@tsinghua.edu.cn.

* To whom correspondence should be addressed.

Manuscript received: 2016-11-01; revised: 2017-03-06; accepted: 2017-03-07

Specifically, we suppose that there are n ASes in the AS path, and the BGP AS path vector from the source AS p_n to the destination AS p_1 is defined as $\text{asp} = (p_1, p_2, \dots, p_n)$. A BAPL happens if $p_i = p_j, j > i, j - i \neq 1$. Although previous studies have been conducted on BAPL, this topic has not been systematically investigated. Thus, in the current paper, we conduct a systematic study on BAPL using real BGP and traceroute data, covering the following important aspects of BAPL.

Relationship between BAPL and forwarding looping. The BGP AS path denotes the list of ASes through which the BGP update messages propagate, while the forwarding AS path is the list of ASes that actually propagate the data packets.

Previous studies have shown that inter-domain forwarding loops exist in the Internet^[9,10]. In theory, BAPL can potentially cause forwarding loops. If BAPL contributes to multi-AS forwarding loops, then analyzing the distribution of BAPL behavior and studying its possible causes can help us understand how to reduce loop-induced transmission delay and packet loss^[11,12], and to prevent attackers from interrupting the Internet^[13]. However, whether a BAPL can cause real forwarding loops in reality has not been studied; therefore, it is an aspect that we examine in this paper. Our observation verifies that a small fraction of BAPLs (approximately 1%) can cause inter-domain forwarding loops.

Characteristics of BAPL. We observed that there were more than 21 900 BAPL updates occurring for IPv4 per day and more than 3800 for IPv6 on average. The majority (more than 74%) of the loops in IPv6 lasted shorter than one day, while a non-trivial number of BAPL updates lasted longer than a month. Two-AS loops and three-AS loops dominated the loop length distribution.

Potential causes of BAPL. We show that BAPL may occur for a few valid reasons, such as networks of multinational companies and preventing particular AS from accepting routes. In addition, Private AS Number Leaking (PANL) contributed to 0.20% of BAPLs in IPv4. BAPLs caused by invalid reasons (PANL, faulty configurations, and intentional attacks) should be fixed by network operators in case that they lead to inter-domain forwarding loops.

The Internet Assigned Numbers Authority (IANA) has reserved the AS numbers (64 512–65 535) for private use, and thus private AS numbers should not

be propagated on the global Internet^[1]. However, previous studies have demonstrated that PANL exists in the Internet^[10]. A PANL occurs if a private AS number exists in the AS path attribute. Specifically, as described in the definition of asp , if $\exists i \in [1, n], p_i \in [65\,512, 65\,535]$, then a PANL occurs. The large number of PANLs that contribute BAPLs motivates us to investigate PANL in the following aspects.

- (1) *Characteristics of PANL.* On average, more than 5900 PANL updates occur per day in IPv4 and more than 1900 in IPv6. We also observed that private AS numbers are most likely located in the source of the AS path, demonstrating that faulty configuration is the major cause of PANL. In addition, more than 62.3% of the PANLs are contributed by five private AS numbers, which is prone to conflicts when BGP routers select private AS numbers as their next ASes.
- (2) *Relationship between BAPL and PANL.* The majority of the BAPLs brought about by PANLs endure less than a day, and the number of BAPLs that result from the condition in which two or more private ASes do not check the BGP AS path (hereafter, collectively referred to as *type2 BAPL*) are much larger than that of BAPLs that result from one private AS. We illustrate the explanations of *type2 BAPLs* and provide advice to the operators of ASes.

A preliminary version of this work was previously published^[14]. The current paper presents several new results, including (1) extension of the study on PANL and (2) discussion of the relationship between BAPLs and PANLs.

The rest of this paper is organized as follows. Previous studies related to BAPL are summarized in Section 2. Section 3 describes the data sets and methodology used in this study. Section 4 discusses the relationship between BAPL and forwarding loops. Section 5 presents the BAPL characteristics. Section 6 provides the explanations of BAPL. Section 7 shows the measurement of PANL, and Section 8 discusses the relationship between PANL and BAPL in depth. Section 9 concludes our work and discusses the future research.

2 Related Work

Several studies have been conducted on routing loops, but few of them have focused on BAPL behavior or on the relationship between BAPL and forwarding looping.

Some studies have focused on *forwarding* AS path loops or *forwarding* routing loops. For example, Paxson has studied routing loops using end-to-end traceroute measurements collected in 1994 and 1995^[15]. Although this paper focused on persistent loops, it found a few transient loops and conjectured that such loops were due to link failure information. Mao et al.^[10] believed that some ASes did not broadcast their infrastructure addresses and others could announce the addresses of shared equipment at border points between ASes, which led to some forwarding AS path loops in traceroute. Xia et al.^[9] presented a measurement study on persistent forwarding loops, and analyzed the possibility of flooding attacks that exploited persistent forwarding loops. They performed extensive measurements to study persistent forwarding loops, and found that persistent loops across multiple ASes did exist on the Internet. Traceroute was also used for measurement in the study, which showed that 0.2% of routable addresses were found to experience persistent forwarding loops. Nevertheless, loop detection only using end-to-end tools such as traceroute is error-prone and cannot successfully detect transient loops^[11].

Pei et al.^[16] investigated *transient* BGP path vector route looping behavior. They analyzed the cause of transient BAPL behavior theoretically and explained how AS path loops would form and resolve and how long they would last. This paper believes that routing updates are slowed down by delays because of physical constraints and protocol mechanisms. Therefore, the inconsistent routing information on different nodes leads to AS path loops during convergence, which depends on the ability of each node to choose an alternative path without loops. The Minimum Route Advertisement Interval (MRAI) is the main factor in the duration of transient AS path loops. Furthermore, Weitz et al.^[17] proposed a bagpipe system to verify the BGP configuration. Methods of detection of BGP AS path loops have been studied by Refs. [18, 19]. Moreover, the private AS number leaking in BGP AS path information has also been mentioned by Refs. [20–23].

Mahajan et al.^[7] presented an example of BAPL: a key AS of Internet introduced BAPL intentionally to achieve several strategies, while this behavior was unnecessary for most operators of BGP routers. Similarly, Shi et al.^[6], Javed et al.^[24], and Katz-Bossett et al.^[25] also introduced an instance for BAPL caused by intentional configuration: the University of Washington and Georgia Institute of Technology conducted out a

rerouting experiment that applied to AS 47065 and led to BAPLs.

Different from previous studies, the current paper focuses on the distribution and causes of both *transient* and *persistent* BAPL behaviors. To avoid the disadvantage of measurement using only traceroute, we have conducted out our measurement study on BAPL using both *RouteViews*^[26] and *traceroute*^[27].

3 Data Sets and Methodology

The *forwarding* AS path denotes the list of ASes traversed by data packets. BGP AS path, also called signaling AS path, represents the list of ASes that propagate the BGP update messages. For example, as Fig. 1 shows, ASa, which has a destination p , propagates the BGP update message to ASd through ASb and ASc, and then (ASd, ASc, ASb, ASa) is ASd's BGP signaling AS path to destination p . Meanwhile, ASd forwards packets to destination p in ASa along the path of ASd, ASc, ASb, ASa, and then (ASd, ASc, ASb, ASa) is ASd's forwarding AS path to destination p . However, these two types of AS paths are not always identical due to various reasons, such as route aggregation/filtering and forwarding anomalies^[10, 28].

To collect the forwarding AS paths, we employ traceroute^[29], which is widely used to observe routing problems and discover the underlying network topology. In traceroute, the interfaces on a forwarding path are identified and the round-trip time statistics for each hop along the way are reported. This approach is considered as the only effective way to observe how packets pass through the Internet under the circumstance of no access to private routing data. To improve our understanding of the relationship between BAPL and forwarding looping behavior, we follow the methodology presented in Ref. [10] and measure the

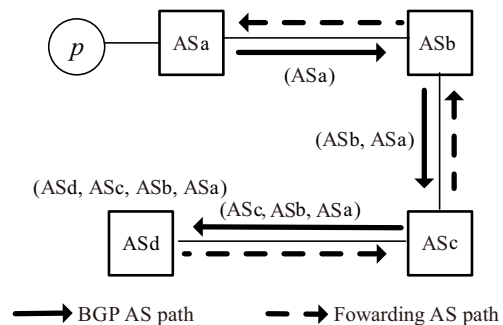


Fig. 1 An example of BGP AS path and forwarding AS path.

signaling AS path and forwarding AS path at the same time. When a forwarding AS loop is identical to the signaling AS path loop, we consider that the forwarding AS loop is attributed to BAPL.

Each BGP node announces its best paths to all destinations to its neighbors, and records the most recent paths received from all of its neighbors. BGP advertises the route to each destination only once, and sends subsequent updates only upon route changes. Consecutive updates for the same destination are spaced out by M seconds (default value 30) using an MRAL. When a current path to a destination is no longer available, the BGP router will attempt to find an alternative path by checking all the saved paths it learned from its neighbors previously. If no alternative path exists, the router will send an explicit path withdrawal message to its neighbors.

To obtain the signaling AS path used in this study, we collect data from the publicly available Oregon RouteViews route-views4 collector^[26], which gathers BGP data from its geographically distributed AS peers (sometimes also called monitors)) for both IPv4 and IPv6.

The route-views4 collector dumps snapshots of the BGP routing table (RIB) for each of its peers every two hours in the Multi-threaded Routing Toolkit (MRT)^[29] format. In addition, the collector receives BGP routing updates from its peers, and writes the collected BGP routing updates into files every 15 minutes in the MRT format^[30]. BGP RIB and updates both contain attributes such as timestamp, peer IP, peer AS, prefix, AS path, and origin AS. Among these attributes, the AS path attribute is the signaling AS path, and we use it to analyze BAPLs and PANLs. The timestamp in the RIB is the time when the snapshot is dumped, while the timestamp in the update is the time when the update is received from a peer.

We collect the RIB data at 00:00:00 on August 1, 2011 and the BGP update data from RouteViews in 1492 days from August 1, 2011 to August 31, 2015. Based on the RIB data and the update data, we obtained the routing table at anytime during the period. When a new update appears, a corresponding record will be added to the routing table. A record may be removed from the routing table as a result of a withdrawal or a different update.

4 BAPL May Lead to Forwarding Loops

In general, a packet from the source traverses a sequence of routers to reach the destination. A packet

experiences a forwarding loop if it traverses a set of routers more than once. Studies have shown that forwarding loops can cause packets in the loops with higher loss rate and longer delay. Other packets that traverse one or more links in the loop, could have longer delay and higher jitters due to the resource consumption caused by the looping packets^[11,12]. Such a vulnerability can be exploited by attackers to overload the shared links to disrupt the Internet connectivity to certain victim destination addresses or prefixes^[13].

Xia et al.^[9] and Mao et al.^[10] have shown that multi-AS forwarding loops existed in the Internet. Will BAPL contribute to multi-AS forwarding loops? We suppose that BAPL may lead to inter-domain forwarding loops, and then analyzing the distributions and explanations of BAPL behavior will help to prevent part of forwarding looping, thereby reducing packet loss rate, preventing attackers from disrupting the Internet, and decreasing link utilization and corresponding delay.

We conducted case studies to analyze whether the BAPLs we have observed can actually cause forwarding AS path loops or not. We tried to find a RouterView peer AS among the observed BAPLs who had a looking glass router that allowed us to run traceroute toward the destination prefix. For example, on September 8, 2013, using RouteViews^[26], we observed a signaling AS path (AS1299, AS6453, AS577, AS7788, AS6407, AS7788) destined for prefix 64.26.148.0/24 in the RIB entry for the monitor 80.91.255.62 (from AS1299), and this BAPL lasted more than a few days. The traceroute^[27] resulted from 80.91.255.62 (which happened to be a looking glass router) to 64.26.148.28 (an IP address in the destination prefix) witnessed a forwarding loop as shown in Table 1. Using the method introduced in Ref. [10], we converted the router-level forwarding path into a forwarding AS path, which turned out to be identical to the signaling AS path. In particular, the forwarding AS loop was identical to the BGP AS path loop.

Moreover, we repeated the preceding experiments and found that only 1% of the signaling loop accounted for forwarding loops. This percentage might be biased because we only sampled the signaling loops in which we could use the looking glass router to run traceroute. Nevertheless, our findings show that BAPL behavior could indeed cause inter-domain forwarding loops. This observation motivates us to conduct further in-depth investigation of BAPL behavior in the rest of this paper.

Table 1 An example of traces that contains forwarding loops.

Hop	Router address	AS number
1	213.155.133.147	1299
2	213.155.133.142	1299
3	213.155.130.51	1299
4	80.91.249.29	1299
5	213.155.131.139	1299
6	213.248.100.178	1299
7	63.243.128.42	6453
8	64.86.85.1	6453
9	216.6.87.9	6453
10	216.6.98.58	6453
11	64.86.85.1	6453
12	216.6.98.58	6453
13	67.69.218.3	577
14	209.217.64.37	7788
15	206.191.0.89	7788
16	67.230.128.70	6407
17	209.217.64.37	7788
18	206.191.0.89	7788
19	67.230.128.70	6407
20	209.217.64.37	7788
21	206.191.0.89	7788
22	67.230.128.70	6407
...


5 Measurement of BAPL

Studies have shown that BAPLs exist on the Internet^[6, 7, 24, 25], but the scale of BAPLs in IPv4 and IPv6 remains unclear. We used to believe that all of BAPLs were caused by misconfigurations. However, Xia et al.^[9] observed a great many persistent forwarding AS path loops that may be caused by persistent BAPLs. Moreover, we want to know the distribution among the duration of BAPLs and the explanations for persistent BAPLs. As BAPLs may lead to forwarding AS path looping, and the loop length is important for attackers to amplify the traffic in the forwarding links, the distribution of BAPL loop length is also studied in this paper.

5.1 Total number and ratio of BAPLs

We define a BAPL as a BGP update, the BGP AS path of which includes an AS path loop. With the daily BGP update data described in Section 3, the number of BAPLs per day is counted.

Figure 2a shows the number of BAPLs and the ratio of the number for BAPLs to the number of all of the BGP updates collected by RouteViews in IPv4 on a

**Fig. 2** The number and the ratio of BAPLs.

daily basis from August 1, 2011 to August 31, 2015. Figure 2b shows the number and ratio of BAPLs in IPv6. Overall, 32 712 387 BAPLs have been observed in IPv4 and 5 563 527 in IPv6 during 1492 days.

The medians of the number and ratio of BAPLs for each year in IPv4 and IPv6 are listed in Table 2. The median of the number for each year is the median number of the set of BAPL numbers per day, and the median of the ratio for a certain year is the median ratio of the set of BAPL ratios on a daily basis. The number of BAPLs increased dramatically from 2011 to 2014 in IPv4 and decreased in 2015. Due to the explosion

Table 2 Medians of BAPLs per year.

Year	Number of IPv4	Ratio of IPv4	Number of IPv6	Ratio of IPv6
2011	4866	9.28×10^{-4}	13	3.55×10^{-5}
2012	5431	1.06×10^{-3}	17	5.99×10^{-5}
2013	18 249	2.70×10^{-3}	27	5.78×10^{-5}
2014	28 810	2.84×10^{-3}	983	8.76×10^{-4}
2015	16 405	1.26×10^{-3}	189	9.33×10^{-5}

of the global BGP routing table, the ratio of BAPLs remained stable in 2012 and 2014, and witnessed a rapid growth in 2013 and a dramatic decrease in 2015. In IPv6, the number of BAPLs remained stable in 2012 and 2013, increased rapidly in 2014, and decreased sharply in 2015, as well as the ratio. The rapid increase in 2014 may have resulted from faulty configurations, with the BAPLs alive for a long time (shown in Fig. 3 and Fig. 2b). As a result, in 2014, the medians of BAPL quantity in IPv4 and IPv6 are higher than those in 2013 and 2015.

The deployment scale of IPv6 is much smaller than that of IPv4, and most of the facilities in IPv6 are deployed later than those in IPv4. Incidents such as faulty configurations, malicious attacks, and other potential causes discussed in Section 6 occur much less frequently in IPv6 than those in IPv4. As a result, the number and ratio of BAPLs in IPv6 are much smaller than those in IPv4.

5.2 Duration of BAPLs

We also studied the duration of BAPL with the BGP update data (defined in Section 3). As described in the preceding sections, a BGP entry can be removed from the routing table due to a withdrawal or a new update. We define the duration of a BAPL as the time interval between its announcement and its withdrawal or a new replacement announcement without the same AS path. For a BGP RIB entry, the duration is the period from 00:00:00 on August 1, 2011 to the time when the entry is withdrawn or replaced with a different update.

The persistence of the BAPLs is studied. Figure 3 shows the Cumulative Distribution Functions (CDF) of the distribution of duration for BAPLs in IPv4 and IPv6. That is, 32 712 364 out of 32 275 791 (98.64%) BAPLs in IPv4 and 4 146 789 out of 5 563 518 (74.53%) BAPLs in IPv6 last shorter than one day. These short-lived BAPLs could be attributed to configuration faults or

malicious attacks. In IPv4, the longest duration is 669 days (until 24:00 on August 31 2015), while in IPv6, the duration is 1260 days (until 24:00 on August 31, 2015).

The result is beyond our expectations. We previously expected that fault configuration was the only factor that led to BAPL behavior. Were it true, BAPL should last shorter than what we have observed. We discuss this problem in Section 6.

5.3 Loop length of BAPLs

As discussed, BAPLs may lead to multi-AS forwarding loops. AS path loops and forwarding paths may share one or more links to the destination prefixes or addresses. An attacker can use BAPLs to overload the shared links to interrupt the connectivity with those reachable prefixes or addresses^[13].

The length of an AS path loop is important for the traffic amplification in the links. When a packet enters an AS path loop, the packet may traverse the links in the loop several times before its TTL expires. Obviously, the shorter the loop length is, the more times the packet will spend to traverse the links in the loop. Since the BGP AS path vector from p_n to p_1 is $asp = (p_n, p_{n-1}, \dots, p_1)$, by definition, $p_i = p_j, j > i, j - i \neq 1$ for a BAPL, and the loop length of asp is $j - i$. Figure 4 shows loop length distribution of BAPLs in IPv4 and IPv6. The number of the BAPLs of AS path loop length l for a certain year denotes the number of different RIB entries or updates that contain looped AS path with the loop length l in the year. Obviously, the bulk of BAPLs had 2-hop or 3-hop loops for both IPv4 and IPv6, which facilitates amplification of the amount of traffic remarkably to destination addresses in the links that appear in the loops.

As RFC 4271^[5] describes, BAPL should not occur in any case, but our observation shows the large scale of BAPLs in both IPv4 and IPv6. We previously




Fig. 3 CDF of BAPLs duration.




Fig. 4 Loop length of BAPL.

believed that misconfiguration was the only explanation for BAPL, and all BAPLs should be transient, while our observation shows that a large number of BAPLs last longer than one day. Furthermore, most BAPLs have a 2-hop or 3-hop loop, which is easily exploited by attackers who overload the links.

6 Explanations of BAPL

A few possible causes can lead to BAPL such as deployment of routing policies, route experiments, PANL, networks of multinational companies, faulty configuration, and intentional attacks.

6.1 Multinational companies

Some multinational companies have exchange points all over the world, and several exchange points may share the same AS number, where operators configure their Customer Edge (CE) routers to accept routes of which the AS path attributes contain their own AS number^[31]. When BGP routing updates pass through exchange points with the same AS number but locate in different countries, they may also go through one or more ASes among the exchange points. It appears as if the BGP updates the loop in the AS path from the BGP perspective. For example, NTT Communications Corporation^[32] has exchange points in Frankfurt, Tokyo, and several cities in the United States, which share the same AS number, 2914. When the prefix in AS_x propagates the BGP updates to AS_y (AS_x and AS_y represent independent AS.), as Fig. 5 illustrates, the message passes through the exchange point of NTT in Chicago and Frankfurt. If the BGP router of the exchange point in Frankfurt computes the degree of preference of the route based on preconfigured policy information, and does not discard the routing updates of which the AS path attributes contain AS2914, a BAPL (AS_x, AS2914, AS_b, ..., AS_a, AS2914, AS_y) from the BGP perspective occurs.




Fig. 5 Multinational companies.

6.2 Preventing particular AS from accepting routes

Some BGP operators prepend to an AS path to keep other providers from picking up the routes. For example, the BGP operator of AS 3066 wanted to send routes to Sprint (AS1239)^[33] that were not to be picked up by UUnet/Verizon Business (AS701)^[34], so the path (AS3066, AS701) was prepended to the AS path^[8]. When the BGP routing updates containing the AS path vector (AS1239, AS3066, AS701, AS3066) arrived to AS701, AS701 discards the message immediately and no relative traffic traverses AS701. As a result, a BAPL (AS1239, AS3066, AS701, AS3066) was propagated on the Internet. The BAPL was artificially injected by the operator of AS3066, which would not lead to any forwarding loop.

Similarly, on August 18, 2011, the University of Washington and the Georgia Institute of Technology conducted a rerouting experiment that applied AS47065^[6,24,25]. In this experiment, a looped AS path (47065, x , 47065) for prefix 184.164.255.0/24 was announced, so that AS_x could not accept this route later, and related traffic would not pass through AS_x. Obviously, the prepense configuration of the network operator on BGP routers can lead to BAPLs that do not account for any forwarding loop.

Zebra^[35] is a well-known tool used for conducting BGP route experiments. The operators can turn a server into a full-powered router with Zebra, and change the configurations dynamically using the terminal interface. Specifically, with the help of Zebra, the operators can filter AS paths, and modify any attribute of the BGP update including the BGP AS path.

6.3 Private AS number leaking

The IANA has reserved the AS numbers (65 512–65 535) for private use and private AS numbers should not to be advertised on the global Internet^[1]. However, we have observed a large number of AS paths that contain private AS numbers. As explained by Mao et al.^[10], when a customer who uses a private AS number mistakenly leaks BGP routes learned from one upstream provider to another, an AS path containing the private AS number may arise. Some private AS number leaking events even account for BAPLs, which is quite beyond our expectations. As described in the definition of asp, the context where BAPLs are caused by private AS number leaking can be described as $p_i = p_j, j >$

$i, j - i \neq 1, \forall m \in (i, j), p_m \in [65\ 512, 65\ 535]$.

Specifically, in our observation, at least 66 156 out of 32 712 387 (0.20%) BAPLs are definitely caused by AS number leaking in IPv4.

If an AS is requested to communicate with a single provider using BGP, it can use a private AS number, which is not used unless the routing policy between the provider and the AS is not visible on the Internet. As Fig. 6 shows, when the prefix in AS_x propagates BGP updates to AS_y , the message passes through AS_a and AS_p . AS_p communicates with its single provider AS_a using a private AS number, and normally, the private AS number should not be advertised on the Internet. However, when the private AS number is leaked on the Internet, the usual reason is misconfiguration, and then an AS path loop ($AS_x, AS_a, AS_p, AS_a, AS_y$) from the BGP perspective forms.

6.4 Faulty configurations or malicious attacks

BAPLs can also arise when a BGP router incorrectly reserves routing updates of which the AS path attributes contain the local AS number. This condition could occur because of configuration errors or even malicious attacks.

Argus^[36] is an agile system to detect prefix hijacking and other anomalies caused by misconfigurations or malicious attacks; this system has been collecting data since June 1, 2011. After cross checking the RouteViews data described in Section 3 and the data collected from Argus, we found that in IPv4, at least 170 036 out of 5 973 568 (2.85%) BAPLs were associated with prefix hijacking or other routing anomalies from June 1, 2011 to December 31, 2013. These prefix hijacking and routing anomalies could be attributed to faulty configurations or intentional attacks, which means that at least 2.85% of BAPLs were caused by misconfigurations or malicious attacks.

As we have stated, several valid factors can contribute to BAPLs, such as multinational cooperation and preventing particular AS from accepting routes, while other BAPLs can be attributed to invalid reasons, such as misconfigurations and intentional attacks, which contributed to at least 2.85% of BAPLs in IPv4.




Fig. 6 Private AS number leaking.


7 Measurement of PANL

As we have explained, private AS numbers have been reserved for private use and should not be advertised on the Internet. However, we found a large number of AS paths that contain private AS numbers in the measurement study of BAPL. PANLs may not only leak the private information of ASes, but also fluctuate the length of the BGP AS path. In addition, PANLs may affect the BGP decisions. To the best of our knowledge, no previous work has investigated PANL systematically. Therefore, we discuss the characteristics of PANL in this section.

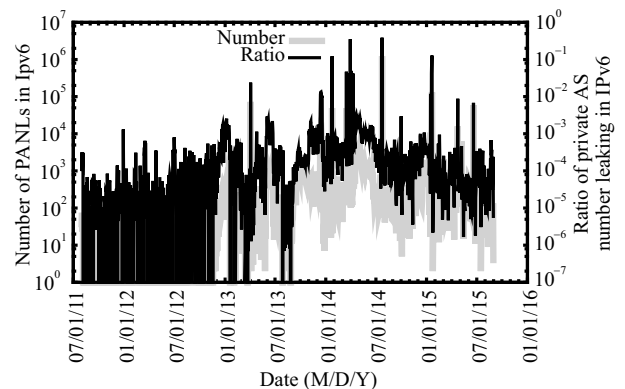
7.1 Total number and ratio of PANL

We define a PANL as a BGP update whose AS path includes a private AS number. We count the number of PANLs per day based on the BGP update data described in Section 3. Figure 7a shows the number and ratio of the number of PANLs to the number of all BGP updates collected by RouteViews in IPv4 on a daily basis from August 1, 2011 to August 31, 2015. Figure 7b shows the number and ratio of PANLs in IPv6.

According to Fig. 7, PANLs have been in the global network for a long time. The spikes in Fig. 7 indicate



(a) IPv4



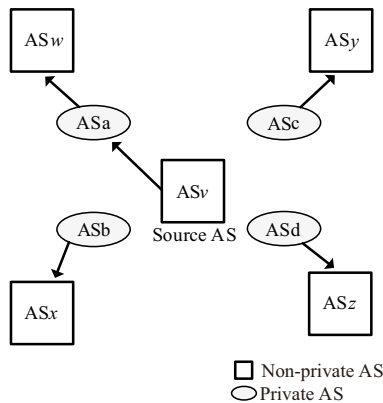
(b) IPv6

Fig. 7 Number and the ratio of private AS number leaking.

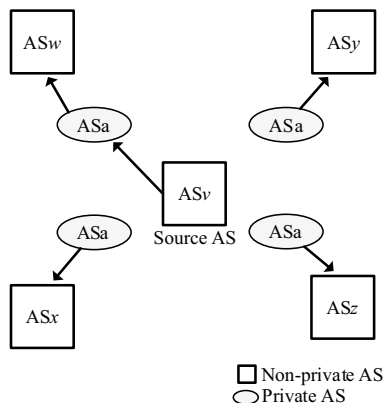
that the number of PANLs sometimes increase sharply due to misconfigurations. We observed more than 5900 PANLs in IPv4 and more than 1900 PANLs in IPv6 per day on average. The number of PANLs in IPv6 are small compared with those in IPv4 in 2011 and 2012, as well as the ratio of the number of PANLs to the number of all BGP updates. The reason is that compared with IPv4, the deployment scale of IPv6 is small and the facilities of IPv6 are new. Misconfigurations and other potential causes of PANLs occur more frequently in IPv4 than in IPv6. As the scale of deployment of IPv6 increases and the IPv6 facilities continue to age, the number of PANLs grew rapidly in 2014.

7.2 Distribution of private AS numbers in PANLs

Since the AS number is the only identifier to distinguish one AS from others, a BGP router broadcasts the BGP AS path by identifying the next hop AS number. If the distribution of private AS numbers in PANLs is concentrated, then PANL can result in BGP broadcast disorders between ASes. As the example in Fig. 8a



(a) Private AS numbers in PANLs are dispersed.



(b) Private AS numbers in PANLs are concentrated.

Fig. 8 Effect of the distribution of private AS numbers in PANLs.

shows, the non-private AS v is connected to 4 private ASes, i.e., ASa, ASb, ASc, ASd. Although ASa is a private AS, AS v is able to broadcast the AS path to AS w because ASa is unique among the neighbor ASes of AS v . However, if the private AS numbers of the neighbor ASes of AS v are the same, as Fig. 8b shows, AS v may not broadcast the BGP AS path to AS w correctly because it may select the next hop incorrectly.

As Table 3 shows, five private AS numbers constitute the majority of all the private AS numbers in PANLs, which can lead to BGP AS path broadcast disorders.

Two main explanations account for the concentrated distribution of private AS numbers in PANLs as follows:

- Popular private AS numbers are frequently used by operators. For example, AS65001, AS65000, AS64777, AS65534, and AS65535 are likely used because operators can easily remember these private numbers. Therefore, these private AS numbers appear frequently in PANLs.
- Intentional configurations of network operators. Sometimes, the operators of AS x do not want to forward packets freely, and the operators of AS x can insert several same private AS numbers into the AS path around AS x . Then, the AS path becomes a low priority when the source AS tries to select an AS path to forward packets. Therefore, fewer packets are forwarded by the AS x . For example, the private AS number AS65332 is inserted into multiple AS paths and each AS path contains several AS65332. This type of PANLs is due to operators intentional configurations for business benefits.

7.3 Location of private AS numbers in PANLs

As explained in the configuration recommendations by Cisco^[37], BGP routes should prohibit anomalous AS paths that contain private AS numbers from being announced to the Internet. In addition, according to RFC 4271^[5], BGP routers should filter private AS numbers in BGP updates that are announced from intra-domains. As a result, a private AS number should never be located at the source of an AS path. However, our

Table 3 Distribution of the private AS numbers in PANLs.

Amount	Private AS number	Percentage (%)
2 407 697	65 332	21.2
1 938 457	65 001	17.1
1 687 932	65 000	14.9
734 153	64 777	6.5
294 950	655 34	2.6

experiment results contradict this inference.

The distribution of the locations of private AS numbers in PANLs is shown in Fig. 9. The location of a private AS number in PANL indicates the number of ASes between the first private AS number (counted from the source AS number to the destination AS number) and the source AS number of the BGP AS path. As Fig. 9 shows, whether in IPv4 or in IPv6, the majority of private AS numbers in PANLs are located at the source of BGP AS paths.

As we have mentioned, the private AS number should not appear in the source of a BGP AS path if all of the configurations of the BGP routers are correct. Therefore, we can conclude that the fault configuration of the BGP routers is the major cause of PANL.

Overall, numerous PANLs occur every day, both in IPv4 and in IPv6. In addition, the private AS numbers in PANLs are concentrated, which can result in BGP broadcast disorders. Most private AS numbers appear in the source AS of the BGP AS path, demonstrating that faulty configurations contribute to the majority of PANLs.

8 Relationship Between PANL and BAPL

As mentioned, 66 156 BAPLs are caused by PANLs from August 1, 2011 to August 31, 2015. The large number motivates us to further study the relationship between PANL and BAPL, and the characteristics of BAPLs resulted from PANL.

8.1 Total number of BAPLs that result from PANLs

The number of BAPLs caused by PANLs per day is shown in Fig. 10. The large number of spikes in the figure indicates that the BAPLs that result from PANLs usually last for less than a day. We can infer that faulty configurations or malicious attacks led to these PANLs and then the PANLs led to BAPLs, and the operators

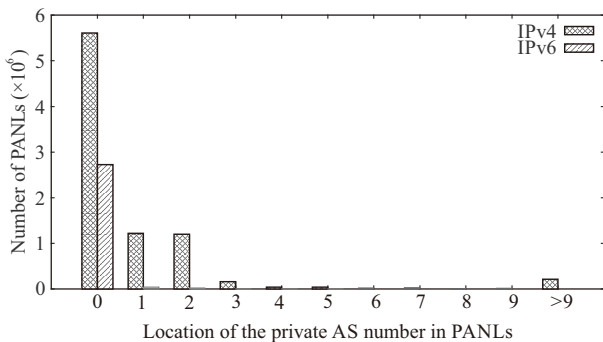


Fig. 9 Distribution of the locations of private AS numbers in PANLs.




Fig. 10 Daily number of BAPLs resulted from PANLs.

mitigated the misconfigurations or attacks in a short time.

Furthermore, to provide advice for network operators to fix the aforementioned problem to avoid PANLs and BAPLs, we investigate the possible fault configurations that may result in PANLs and BAPLs.

8.2 Types of looping

We analyze the characteristics of the BAPLs that result from PANLs. We find that two types of BAPLs result from PANLs, and the complex type of BAPLs caused by PANLs occur much more frequently.

The first type of BAPLs caused by PANLs (type1 BAPL) is shown in Fig. 11a. When the AS path (AS_x, AS_a) is delivered to a private AS_a, the routers of the private AS_a do not check whether AS_a is in the AS path. The routers add AS_a to the AS path and then a looped AS path (AS_a, AS_x, AS_a) occurs.

Two or more private ASes do not check whether they are in the AS path and can lead to the second type of BAPLs that result from PANLs (type2 BAPL). As Fig. 11b shows, AS_a and AS_b are both private ASes. When the AS path (AS_b, AS_a) comes to private AS_a, the routers of the private AS_a will not check whether




Fig. 11 Two types of BAPLs that are resulted from PANLs.

ASa is in the AS path, and then the looped AS path (ASa, ASb, ASa) will occur. When the looped AS path (ASa, ASb, ASa) is delivered to private ASb, ASb will not check whether ASb is in the AS path either, and a more complex looped AS path (ASb, ASa, ASb, ASa) will occur.

Figure 12 shows a more complicated version of *type2 BAPL*. When the BGP AS path (ASb, ASc, ASa) is delivered to private ASa, the routers of private ASa will not determine whether the AS path includes ASa. The routers will add ASa to the AS path and deliver the looped AS path (ASa, ASb, ASc, ASa) to ASb. Similarly, the routers of ASb will deliver the looped AS path (ASb, ASa, ASb, ASc, ASa) to ASd. In the end, a looped BGP AS path (ASa, ASd, ASb, ASa, ASb, ASc, ASa) occurs.

We can observe that the *type2 BAPL* usually include multiple BGP AS links. If these BAPLs lead to forwarding loops, and attackers use these loops to overload the links to interrupt the connectivity, multiple links will be affected in the *type2 BAPL*. In addition, the *type2 BAPL* can generate a long AS path, thereby wasting the storage resources of BGP routing tables. From August 1, 2011 to August 31, 2015, we observed 63 755 *type2 BAPLs*. That is, 95.85% of BAPLs that are induced by PANLs are *type2 BAPLs*. Therefore, we intend to find how the *type2 BAPL* forms and how to resolve it with a typical example.

8.3 Explanations of *type2 BAPLs*

The large number of *type2 BAPLs* motivates us to investigate the causes behind them. The customer ASes of a non-private AS typically use private AS numbers to communicate with the provider AS, and the routing policy between the provider AS and customer AS should not be advertised to the Internet. However, if the private AS numbers of two or more customer ASes are visible within and outside the provider AS, BAPLs are likely to happen.

For example, as Fig. 13 shows, ASa and ASb are




Fig. 12 An example of complicated *type2 BAPLs*.

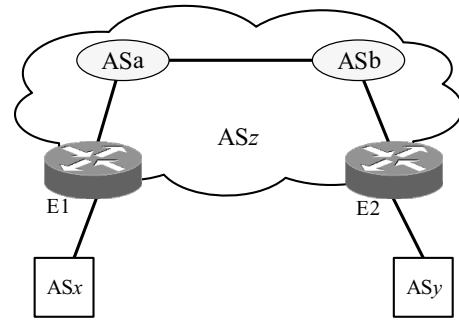


Fig. 13 Illustration of how *type2 BAPLs* form.

customer ASes whose provider AS are ASz. ASa and ASb use private AS numbers to communicate with the non-private AS ASz via non-BGP protocols such as OSPF, and ASx and ASy are non-private ASes connected with ASz via the BGP protocol. ASa and ASb are advertised to the Internet via the BGP router E1 and E2, respectively. After ASz advertises the BGP AS path (ASz, ASx) to ASa via E1, ASa adds ASa to the AS path and advertise the BGP AS path (ASa, ASz, ASx) to ASb and ASz. Then ASb adds ASb to the AS path and advertise the BGP AS path (ASb, ASa, ASz, ASx) to ASa and ASz, and then ASa adds ASa to the AS path and advertises the BGP AS path (ASa, ASb, ASa, ASz, ASx) back to ASb because ASa does not run the BGP protocol and check whether it is in the AS path. Similarly, ASb will also advertise the BGP AS path (ASb, ASa, ASb, ASa, ASz, ASx) to ASz and ASa. Then, a *type2 BAPL* occurs.

The example shows that PANLs can easily result in *type2 BAPLs*, which is consistent with our observations. If the operators of ASes can strictly prohibit the private AS number from being advertised to the Internet, or make customers of ASes that use the private AS number check whether they are in the BGP AS path, *type2 BAPLs* will be eliminated from the BGP AS path table. Therefore, the size of the BGP AS table will be reduced, and no link will be congested because of attacks induced by BAPLs.

9 Conclusion

The BAPLs studied in this paper can be helpful in understanding the operational behavior of BGP in both IPv4 and IPv6. Motivated by this idea, we initially tried to explore the relationship between BAPL behavior and forwarding looping, but we found that only a small part (approximately 1%) of BAPLs can lead to forwarding loops. We have studied the global BGP routing data in

1456 days and analyzed the number and ratio of BAPLs in IPv4 and IPv6. In addition, the duration of BAPLs and the distribution of loop length are also discussed in this paper. Different from our initial expectations, a nontrivial number of BAPLs lasts for more than one month.

Furthermore, BAPL can be attributed to various factors. Reasonable explanations, including routing experiments, multinational cooperation, and preventing particular AS from accepting routes, have also contributed to BAPLs. BAPLs caused by invalid reasons such as the deployment of routing policies, PANL, and misconfigurations should be fixed.

The large number of PANLs that result in BAPLs motivate us to conduct in-depth research on PANLs. We have investigated the number and ratio of PANLs per day from August 1, 2011 to August 31, 2015, and found that the private AS numbers in all of the PANLs are concentrated, which can lead to BGP AS path advertising disorders. The fact that most private AS numbers are located at the source AS of the BGP AS path shows that misconfigurations are the main reasons for PANLs.

We have studied the number of BAPLs that are caused by PANLs per day, and classified them into two types, *type1 BAPL* in which a single private AS leads to BAPL, and *type2 BAPL* in which two or more private ASes result in BAPL. We explain how *type2 BAPL* forms, and provide suggestions to operators of ASes.

In the future, we plan to focus on the correlation between BAPL behaviors and other BGP anomalies. We are also interested in the effects of BAPL on Internet routing instability.

Acknowledgment

This study was supported by the National Natural Science Foundation of China (Nos. 61772307 and 61161140454) and the National Key Basic Research and Development (973) Program of China (Nos. 2013CB329105 and 2009CB320500).

References

- [1] J. Hawkinson, Guidelines for creation, selection, and registration of an autonomous system (AS), <http://tools.ietf.org/html/rfc1930>, April 4, 2015.
- [2] J. Moy, OSPF version 2, <http://www.ietf.org/rfc/rfc2328>, April 5, 2015.
- [3] G. Malkin, RIP version 2, <http://www.ietf.org/rfc/rfc2453.txt>, April 7, 2015.
- [4] R. Callon, Use of OSI IS-IS for routing in TCP/IP and dual environments, <http://tools.ietf.org/rfc/rfc1195.txt>, April 8, 2015.
- [5] Y. Rekhter and T. Li, A border gateway protocol 4 (BGP-4), <http://www.rfc-editor.org/rfc/rfc4271.txt>, April 4, 2015.
- [6] X. G. Shi, Y. Xiang, Z. L. Wang, X. Yin, and J. P. Wu, Detecting prefix hijackings in the internet with argus, in *Proc. 2012 Internet Measurement Conf.*, Boston, MA, USA, 2012, pp. 15–28.
- [7] R. Mahajan, D. Wetherall, and T. Anderson, Understanding BGP misconfiguration, *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4, pp. 3–16, 2002.
- [8] R. E. Seastrom, Re: As path loops in practice? https://www.nanog.org/maillinglist/maillarchives/old_archive/2003-12/msg00260.html, April 2, 2015.
- [9] J. H. Xia, L. X. Gao, and T. Fei, A measurement study of persistent forwarding loops on the internet, *Computer Networks*, vol. 51, no. 17, pp. 4780–4796, 2007.
- [10] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz, Towards an accurate AS-level traceroute tool, in *Proc. 2003 Conf. Applications, Technologies, Architectures, and Protocols for Computer Communications*, Karlsruhe, Germany, 2003, pp. 365–378.
- [11] U. Hengartner, S. Moon, R. Mortier, and C. Diot, Detection and analysis of routing loops in packet traces, in *Proc. 2nd ACM SIGCOMM Workshop on Internet Measurement*, Marseille, France, 2002, pp. 107–112.
- [12] D. Pei, L. Wang, D. Massey, S. F. Wu, and L. X. Zhang, A study of packet delivery performance during routing convergence, in *Proc. 2003 Int. Conf. Dependable Systems and Networks*, San Francisco, CA, USA, 2003, pp. 183–192.
- [13] J. H. Xia, L. X. Gao, and T. Fei, Flooding attacks by exploiting persistent forwarding loops, in *Proc. 5th ACM SIGCOMM Conf. Internet Measurement*, Berkeley, CA, USA, 2005, p. 36.
- [14] S. L. Zhang, Y. Liu, and D. Pei, A measurement study on BGP AS path looping (BAPL) behavior, in *2014 23rd Int. Conf. Computer Communication and Networks (ICCCN)*, Shanghai, China, 2014, pp. 1–7.
- [15] V. Paxson, End-to-end routing behavior in the internet, *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, pp. 601–615, 1997.
- [16] D. Pei, X. G. Zhao, D. Massey, and L. X. Zhang, A study of BGP path vector route looping behavior, in *Proc. 24th Int. Conf. Distributed Computing Systems*, Tokyo, Japan, 2004, pp. 720–729.
- [17] K. Weitz, D. Woos, E. Torlak, M. D. Ernst, A. Krishnamurthy, and Z. Tatlock, Bagpipe: Verified BGP configuration checking, in *ACM SIGCOMM Workshop on Networking and Programming Languages (NetPL 2016)*, Florianópolis, Brazil, 2016.
- [18] B. Al-Musawi, P. Branch, and G. Armitage, BGP anomaly detection techniques: A survey, *IEEE Communications Surveys & Tutorials*, vol. 19, no. 1, pp. 377–396, 2017.
- [19] D. Walton, A. Retana, E. Chen, and J. Scudder, Solutions for BGP persistent route oscillation, Technical Report, IETF, 2016.
- [20] Q. Jacquemart, G. Urvoy-Keller, and E. Biersack, Behind

- IP prefix overlaps in the BGP routing table, in *Proc. 17th Int. Conf. Passive and Active Network Measurement*, Heraklion, Greece, 2016, pp. 289–301.
- [21] M. Čsovi, S. Obradovio, and L. Trajkovi, Classifying anomalous events in BGP datasets, in *2016 IEEE Canadian Conf. Electrical and Computer Engineering (CCECE)*, Vancouver, Canada, 2016, pp. 1–4.
- [22] J. Schutrup and B. ter Borch, BGP hijack alert system, University of Amsterdam, the Netherlands, 2016.
- [23] Y. Y. Wang, J. Bi, K. Y. Zhang, and Y. C. Wu, A framework for fine-grained inter-domain routing diversity via SDN, in *2016 8th Int. Conf. Ubiquitous and Future Networks (ICUFN)*, Vienna, Austria, 2016, pp. 751–756.
- [24] U. Javed, I. Cunha, D. Choffnes, E. Katz-Bassett, T. Anderson, and A. Krishnamurthy, PoiRoot: Investigating the root cause of interdomain path changes, *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, pp. 183–194, 2013.
- [25] E. Katz-Bassett, C. Scott, D. R. Choffnes, Í Cunha, V. Valancius, N. Feamster, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy, Lifeguard: Practical repair of persistent route failures, *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 395–406, 2012.
- [26] University of Oregon route views project, <http://www.routeviews.org/>, 2017.
- [27] T. Kernen, Traceroute, <http://www.traceroute.org>, April 1, 2015.
- [28] T. G. Griffin and G. Wilfong, On the correctness of IBGP configuration, *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4, pp. 17–29, 2002.
- [29] L. Blunk and M. Karir, Multi-threaded routing toolkit (MRT) routing information export format, <http://tools.ietf.org/html/rfc6396>, April 6, 2015.
- [30] P. C. Cheng, X. Zhao, B. C. Zhang, and L. X. Zhang, Longitudinal study of BGP monitor session failures, *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 2, pp. 34–42, 2010.
- [31] P. Marques, Internal BGP as the provider/customer edge protocol for BGP/MPLS IP virtual private networks (VPNs), <http://www.rfc-editor.org/rfc/rfc6368.txt>, April 5, 2015.
- [32] <http://www.peeringdb.com/net/1045>, 2015.
- [33] Sprint, <http://www.sprint.com/>, April 2, 2015.
- [34] Verizon, <http://www.verizonenterprise.com/>, April 2, 2015.
- [35] Zebra, <https://www.gnu.org/software/zebra/>, April 5, 2015.
- [36] Y. Xiang, Z. L. Wang, X. Yin, and J. P. Wu, Argus: An accurate and agile system to detecting IP prefix hijacking, in *2011 19th IEEE Int. Conf. Network Protocols (ICNP)*, Vancouver, Canada, 2011, pp. 43–48.
- [37] S. Hogg and E. Vyncke, *IPv6 Security*. Amsterdam, the Netherlands: Pearson Education, 2008.



Shenglin Zhang received the BE degree from Xidian University in 2012. He is currently working towards the PhD degree in the Institute of Network Sciences and Cyberspace, Tsinghua University. His research interests include failure detection and prediction in datacenter networks, BGP AS path behavior, and P2P

optimization.



Ying Liu received the BS degree in information engineering, MS degree in computer science, and PhD degree in applied mathematics from Xidian University in 1995, 1998, and 2001, respectively. She made postdoctoral research in the Department of Computer Science and Technology, Tsinghua

University in 2001–2003. She is currently an associate professor in the Institute for Network Sciences and Cyberspace, Tsinghua University. Her research interests include multicast routing, network architecture, and router design and implementation. She is an IEEE member.



Dan Pei received the BE and MS degrees in computer science from Tsinghua University in 1997 and 2000, respectively, and the PhD degree in computer science from University of California, Los Angeles (UCLA) in 2005. He is currently an assistant professor in the Department of Computer Science and Technology,

Tsinghua University. His research interests include network and service management in general. He is an IEEE senior member and an ACM senior member.



Baojun Liu received the BE degree from Xidian University in 2015. He is currently working towards the PhD degree in the Institute of Network Sciences and Cyberspace, Tsinghua University. His research interests include BGP AS path behavior, internet routing instability, and IPv6 security.