An Empirical Study of "Bogon" Route Advertisements

Nick Feamster, Jaeyeon Jung, and Hari Balakrishnan
MIT Computer Science & Artificial Intelligence Laboratory
{feamster,jyjung,hari}@csail.mit.edu

Abstract

An important factor in the robustness of the interdomain routing system is whether the routers in autonomous systems (ASes) filter routes for "bogon" address space—i.e., private address space and address space that has not been allocated by the Internet Assigned Numbers Authority (IANA). This paper presents an empirical study of bogon route announcements, as observed at eight vantage points on the Internet. On average, we observe several bogon routes leaked every few days; a small number of ASes also temporarily leak hundreds of bogon routes. About 40% of these bogon routes are not withdrawn for at least a day. We observed 110 different ASes originating routes for bogon prefixes and a few ASes that were responsible for advertising a disproportionate number of these routes. We also find that some ASes that do filter unallocated prefixes continue to filter them for as long as five months after they have been allocated, mistakenly filtering valid routes. Both of these types of delinquencies have serious implications: the failure to filter valid prefixes can could make nefarious activities such as denial of service attacks difficult to trace; failure to update filters when new prefixes are allocated prevents legitimate routes from being globally visible.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: Security and protection; C.2.6 [**Computer-Communication Networks**]: Internetworking

General Terms

Management, Measurement, Security

Keywords

BGP, anomalies, bogon prefixes

1. Introduction

This paper presents an empirical study of a class of invalid routes that appear on the Internet called "bogons": IP prefixes that are either within private address space [15] or in the address space that IANA [4] has either reserved or has not allocated to any RIRs [1, 2, 6, 7].[1] Because BGP routing involves the exchange of routes between thousands of autonomous systems (ASes), an AS that advertises invalid routes could potentially create widespread instability. Invalid routes fall into two categories: "bogons" and "hijacked" routes (i.e., routes that announce reachability to prefix space that belong to another AS). Detecting hijacked BGP routes is a difficult open problem [17, 24], but bogons are easy to spot because they are invalid *by definition*. The CIDR report provides weekly statistics of bogon routes observed at AS 4637 [9]. This paper provides a complementary investigation of bogon route advertisements, performing a longitudinal study of bogon route advertisements observed at 8 distributed locations over 15 months.

The importance of filtering routes with unallocated or private prefixes (bogons) has been known to the network operations community for some time. One study appears to find that some DDoS attacks originated from bogon prefixes [12], which should encourage ISPs to filter packets with invalid source IP addresses, as well as the routes to these prefixes. It is also believed that spammers may be advertising transient routes for invalid and legacy IP address space, from which they can spam without being traced [22]. To help network operators keep their filters up-to-date, the CIDR report [8] publishes a list of bogon prefixes based on data available from Internet Assigned Numbers Authority (IANA) [4] and Regional Internet Registries (RIRs) [1, 2, 6, 7]. Team Cymru [12] also maintains a list of bogon prefixes.

Updating route filters to ensure that *valid* prefixes are not filtered is as important as filtering invalid routes. Incorrectly filtering a valid route can cause serious reachability problems even if only a small number of ASes do so. ASes can sometimes incorrectly filter valid routes if they fail to update their filters when new prefix space is allocated.

Despite the fact that filtering invalid routes from global routing tables (and not filtering valid routes) is a major component of securing the Internet infrastructure and ensuring reachability, our recent study of router configuration errors [13] suggests that many ASes are delinquent in applying route filters and keeping them up-to-date. Based on our study and others [19], as well as anecdotal evidence, we hypothesized that bogon prefix advertisements would be rather prevalent, and that some prefixes would be incorrectly filtered shortly after they became allocated. In fact, over the course of 15 months, we observed 110 different origin ASes leak more than 13,000 updates for prefixes from bogon IP address space.

In this paper, we characterize bogon route announcements by answering the following questions:

- How often do bogon route announcements appear (*prevalence*), and how long do they last (*persistence*)?
- Are there certain bogon routes (i.e., bogon prefixes and address space) that are leaked by more than one AS?
- How are bogon announcements distributed across the ASes that originate them, and how often does each AS leak bogon routes?
- When an AS leaks bogon routes, how many bogon routes are leaked at once?

[1] Team Cymru [12] and Huston [14] use the same definition of bogon. Appendix A lists the bogon prefixes we used in this study.

We observe ASes leaking invalid routes about once every 1.2 days on average. From 8 vantage points, we observed 403 invalid routes originating from 110 distinct ASes. Roughly half of these events last longer than one hour, and about 40% last longer than one day.

About 70% of the invalid announcements and nearly half of the events that caused invalid routes to be leaked involved three portions of private address space: 172.16.0.0/12, 192.0.2.0/24, and 10.0.0.0/8.

Routes from the space 0.0.0.0/7 were leaked by 71 different origin ASes (i.e., almost 75% of ASes that leaked *any* invalid route). Many of these appeared to be routes that are commonly used for testing or internal network addressing.

30 ASes originated invalid prefixes more than once, and two tier-1 providers originated invalid prefixes more than 10 times. Invalid prefixes originated by tier-1 ASes were typically observed at more of our monitoring points than those originated at smaller ASes.

The majority of events only leaked a single prefix, and two-thirds leaked two prefixes or fewer. We observed 14 events where a single AS originated more than 100 invalid prefixes.

Some ASes that *do* filter unallocated prefix spaces do not update their filters until three to five months after prefixes are allocated and may not update their filters until nearly two months after the first advertisement from the newly allocated space before updating their filters.

Table 1: Summary of preliminary findings from our study of "bogon" route advertisements at eight vantage points.

- Do ASes update their route filters when IP address space is allocated from previously unallocated space?

Our study offers a preliminary look into the characteristics of bogon routes on the Internet, as seen from eight topologically diverse vantage points. At each vantage point, we collect BGP update data via iBGP, and cluster BGP updates into distinct events;[2] this clustering allows us to roughly estimate the magnitude of any particular event according to its size (i.e., how many prefixes were leaked in the event) and visibility (i.e., how many of our monitoring points saw the event). Some findings we present about specific ASes leaking routes depend highly on the location of our vantage points, since most of our vantage points are at the network edge and many edge ASes (it is thought) perform some amount of filtering, but we attempt to draw general conclusions where appropriate.

Table 1 summarizes our findings. We observe that some ASes leak bogon prefixes to the Internet slightly more than once per day on average. Roughly half of these routes are not withdrawn within an hour, and several of these events leak more than 100 prefixes at a time. Moreover, when an address range is allocated, ASes often fail to update their filters to allow these prefixes to be advertised, adversely affecting the reachability of *valid* prefixes. Certain ASes appear to filter routes belonging to legitimate prefix space for as long as three to five months after the space was allocated.

The observations in this paper have implications for routing configuration and security. The majority of events that leaked invalid routes involve private address space (e.g., 10.0.0.0/8); because private address space designations change infrequently, routers should filter routes from private address space by default (network operators could override such a default if desired). Our results suggest that a bogon prefix originated by a tier-1 ISP will typically be more widely visible than an invalid route advertised from a regional or stub AS: thus, an attacker is likely to cause more disruption by injecting routes from a tier-1 ISP. Finally, half of the (likely accidental) invalid route advertisements last longer than an hour, and 40% last longer than a single day, which suggests that a significant fraction of misconfiguration events in general may be long-lived; this suggests that many misconfigurations are not quickly noticed, and that operators need auxiliary tools to help them quickly find certain types of misconfigurations.

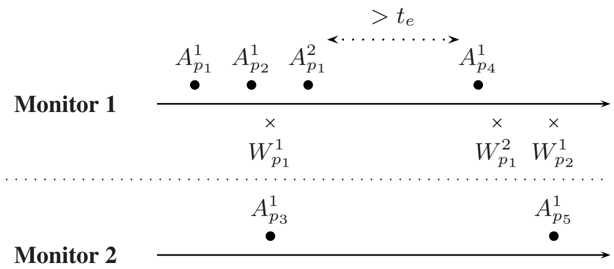

Figure 1: Example of how update streams are grouped into origin AS-based events. All updates shown in this figure are categorized into the same event, except for $A^1_{p_4}$ and $A^1_{p_5}$, which are categorized into a new event. A corresponds to an announcement and W corresponds to a withdrawal.

2. Background and Definitions

In this section, we briefly overview BGP routing and the mechanics of route filtering. We then define bogon events for the purposes of our analysis.

2.1 *Overview of BGP and Filtering*

BGP4 is the Internet's interdomain routing protocol [21]; the Internet comprises about 17,000 independently operated networks, or autonomous systems (ASes), that exchange reachability information using BGP. To advertise reachability to some Internet destination, a router in an AS "injects" a route for an IP prefix into BGP, and that router advertises that route to other ASes in the form of a BGP update. Many IP prefixes are designated for addressing *private* networks or for testing [15]; these prefixes, as well as those that have not been allocated to any regional Internet registry, should not be advertised globally.

When an AS learns a route to a destination, its *import policy* can filter (i.e., ignore) it or modify certain route attributes (e.g., assign a "local preference" value to the route). That router will then select a single best route for each destination and readvertise it over every BGP session for which the *export policy* permits re-announcement. A router will readvertise at least one route for every prefix it learns as long as: (1) the import policy does not discard the route, (2) the export policy does not prevent that prefix from being advertised, and (3) the export policy does not discard the best route based on other route attributes. Thus, if BGP learns a route to a destination via BGP, we can conclude that one or more routers belonging in each of the ASes in the route's AS path fails to filter that prefix in its import and export policy.

2.2 *Prefix-Based and Origin AS-Based Events*

We observe BGP updates at 8 distinct *monitors*; each monitor receives a stream of iBGP update messages from a border router in the AS where it is deployed (we discuss our data collection techniques further in Section 3). Because (1) a single prefix announcement may be visible at multiple monitors, (2) a single prefix withdrawal can cause a flurry of updates [18], and (3) a single configuration fault can cause many distinct prefixes to be advertised,

[2] We define these "events" in Section 2.2.

a simple count of BGP updates is not a good indicator of magnitude. Instead, we define two types of events—*prefix-based events* and *origin AS-based events*—that cluster BGP updates in different ways.

A *prefix-based event* is defined by a period of time whereby a single bogon prefix is being announced by some origin AS. A prefix-based event *begins* when a monitor receives a new announcement for a bogon prefix from some origin AS. Every announcement for the same prefix and origin AS before a withdrawal for that prefix is grouped into the same prefix-based event. A prefix based event *ends* when the route for that corresponding bogon prefix is withdrawn. If a BGP route propagates to more than one monitor, we consider each of those route announcements as separate prefix-based events. Prefix-based events indicate how many times we witnessed a bogon prefix leaked by some origin AS; it is a more accurate reflection of route leaks than a simple tally of announcements because instability or path exploration may artificially amplify the number of announcements.

An *origin AS-based event* attributes multiple bogon route announcements that occur in close succession to the same cause, as might happen with a single configuration change, a sequence of related configuration changes, or other operational incident (e.g., a router reboot). We define origin AS-based events to help us quantify how often these types of incidents occur, as well as how many distinct bogon routes might be leaked as the result of a single incident. A new origin AS-based event *begins* when any monitor receives a bogon route from some origin AS, and that origin AS has not announced any bogon routes within the previous t_e minutes. An origin AS-based event *ends* when no monitor learns any bogon route announcements from that origin AS for at least t_e minutes. A withdrawal of a bogon prefix is associated with the origin AS-based event for the corresponding route announcement. In our analysis, we set t_e to be 60 minutes (previous studies have used a similar time interval to separate distinct BGP events [10]). The withdrawal of a bogon is associated with the event that most recently advertised that route, regardless of the monitor at which it was received.

Figure 1 shows how we group BGP messages into distinct origin AS-based events. Updates $A^1_{p_1}$, $A^1_{p_2}$, $W^1_{p_1}$, $A^2_{p_1}$, and $A^1_{p_3}$ are clustered into the same origin AS-based event even though they were received at two separate monitors, because they were all received from the same origin AS within a time interval that included no quiescent periods of more than t_e minutes. The withdrawals $W^2_{p_1}$ and $W^1_{p_2}$ that correspond to $A^2_{p_1}$, $A^1_{p_2}$ are classified as part of the same event, even though they were learned after more than t_e minutes of quiescence. $A^1_{p_4}$ and $A^1_{p_5}$ are classified as a separate origin AS-based event because $A^1_{p_4}$ occurs more than t_e minutes after the last announcement in the preceding event for that origin AS, as observed at any monitor.

3. Data Collection

Table 2 shows the hosts at which we collected BGP messages. These hosts ran Zebra 0.92a, an open source software router [23], configured to log all BGP updates. The clocks of the monitors are synchronized to within 10 milliseconds. Table 2 also shows the number of BGP updates collected at each site between July 1, 2003 and October 9, 2004, the period of time over which we performed our analysis.

Figure 2 shows where the MIT collection host sits in relation to the border router of the hosting network and the rest of the Internet; other monitors sit in similar positions relative to their border routers. MIT's border router has four upstream feeds: a commercial feed via Genuity/Level3 (AS 1), Cogent (AS 174), Comcast (AS 7015), and to Internet2 via the Northeast Exchange (AS 10578).

Host	*BGP Peers*	*Updates*
MIT (AS 3)	Genuity, Cogent, Comcast, Internet2	84,011,988
PSG (AS 3130)	Genuity, Verio	76,317,493
Vineyard (AS 10781)	Qwest, Savvis	66,584,023
Nortel (AS 14177)	AT&T Canada	21,325,982
Aros (AS 6521)	UUNet, Electric Lightwave	13,786,220
PWH (AS 6549)	7 ISPs	196,587,735
GBLX-JFK (AS 3549)	Many ISPs	110,064,852
GBLX-LON (AS 3549)	Many ISPs	148,344,497

Table 2: Information about BGP data collected from networks where our monitoring hosts are located. We analyze all updates since July 1, 2003.

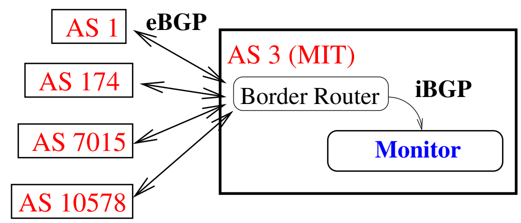

Figure 2: At each collection host, we collect BGP messages from the network's border router. The figure shows the configuration for MIT, which obtains upstream connectivity from Genuity (AS 1) Cogent (AS 174), Comcast (AS 7015), and the Northeast Exchange (via AS 10578).

The monitor receives BGP updates from the border router. Because of the configuration, the monitors will not see all BGP messages heard by the border router; they see only BGP messages that cause a change in the border router's choice of *best* route to a prefix.

Despite not observing all BGP updates, the monitors observe all messages relevant to invalid prefixes that were received at their respective ASes. Because iBGP readvertises one best route for *every* prefix, a monitor will always advertise a route to a bogon prefix if its border router receives such a route.

4. Results

In this section, we address the questions from Section 1. We observed 13,411 BGP updates (373 origin AS-based events and 4,770 prefix-based events, as defined in Section 2.2) for 403 different bogon prefixes in 36 distinct regions of bogon IP address space. On average, an origin AS-based event leaks bogon prefixes roughly every one and a half days. Figure 3 shows the number of prefix-based events and origin AS-based events observed per week observed at eight vantage points.

4.1 Prevalence and Persistence

Prevalence: Table 3 shows the bogon *address space* from which bogon prefix announcements were most commonly leaked, in terms of both prefix-based events and origin AS-based events. The table confirms one hypothesis that we had: many routes with bogon prefixes come from *private* or *test* IP address space (as defined by RFC 3330 [15]), rather than unallocated address space. This observation suggests that many bogon route leaks are probably accidental, resulting from leaks of routes that were most likely intended to be routes for infrastructure inside of a single AS. Most of the leaks of private IP address space were visible at only one or two monitoring points, but many announcements for bogon prefixes in reserved IP address space were more widely visible.

Figure 4 shows the number of BGP routing *prefix-based events* that we observed for specific bogon prefixes. Prefix-based events

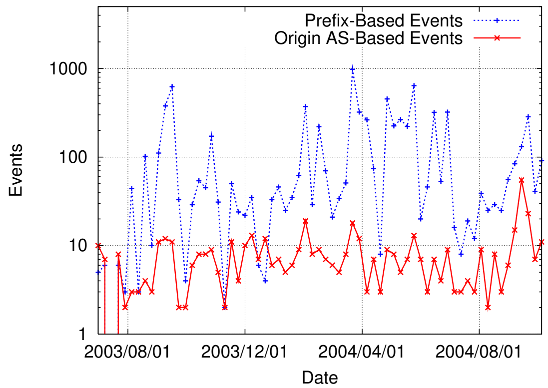

Figure 3: Bogon route announcement events observed at eight vantage points.

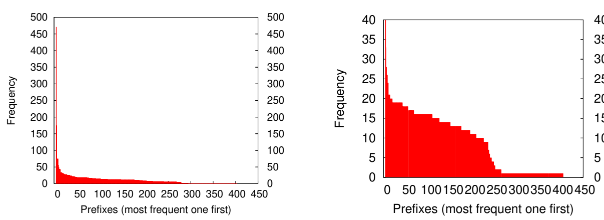

(a) Prefix-based events.

(b) Origin AS-based events.

Figure 4: Prevalence of announcements for *specific bogon prefixes* within bogon IP address space.

are distributed across many different bogon prefixes: more than 85% of bogon prefixes (40% of prefix-based events) were leaked fewer than 15 times, and more than 30% of bogon prefixes were leaked fewer than 5 times. Figure 4 (b) and Table 4 show the prevalence of bogon prefixes in *origin AS-based events*; that is, how many origin AS-based events each bogon prefix appeared in. A single bogon prefix, 1.1.1.0/24, was announced 150 separate times by 9 different origin ASes in 19 distinct origin AS-based events. Announcements for this prefix were observed at all 8 of our monitoring locations, indicating that many ISPs do not filter the bogon space containing this prefix. The North American Network Operators Group (NANOG) runs a mailing list where network operators report operational problems, discuss operational issues, etc. [20]; interestingly, leaks of routes for this exact prefix were a topic of discussion on the NANOG mailing list six years ago [16] (it appears that this prefix is also used for testing and internal addressing). Unfortunately, it appears that the filtering situation for these types of prefixes has not improved. Other specific routes appear to be leaked particularly often: 1.0.0.0/8 was advertised and withdrawn 29 separate times, and many distinct prefixes in bogon address space 172.16.0.0/12 were advertised more than 25 times (more than 3,400 prefix-based events were contained in this IP address space alone, though all of these were leaked by a single origin AS).

Persistence: Bogon prefix announcements can sometimes appear as the result of a configuration problem. For example, to update access control lists, network operators must first "clear" the old filter before installing the new one, possibly resulting in potential temporary leaks of bogon addresses (which can legitimately be used *within* the network) when the filters are cleared. However, bogon prefixes should be noticed and withdrawn shortly if administrators regularly monitor their configurations and BGP tables.

Bogon Space	*Prefix-based Ev.*	*Origin AS-based Ev.*	*Monitors*	*Distinct ASes*
172.16.0.0/12	3439	106	1	4
0.0.0.0/7	276	87	8	71
72.0.0.0/5	169	72	8	10
192.0.2.0/24	17	40	2	3
96.0.0.0/3	240	31	8	12
10.0.0.0/8	34	17	2	5
189.0.0.0/8	29	10	6	2
2.0.0.0/8	61	7	8	3
169.254.0.0/16	10	6	3	5
223.0.0.0/8	350	6	8	2
176.0.0.0/5	10	3	2	3
5.0.0.0/8	9	3	5	3
88.0.0.0/5	24	3	6	3
58.0.0.0/7	14	3	6	3
174.0.0.0/7	11	2	6	2
42.0.0.0/8	26	2	8	2
50.0.0.0/8	6	2	4	2
27.0.0.0/8	2	2	2	2
192.168.0.0/16	12	2	2	2
173.0.0.0/8	2	2	2	2

Table 3: Top 20 most common bogon route announcements from bogon IP address *space*, sorted by the number of origin AS-based events. Many bogon route announcements come from private or test IP address space, shown in bold.

Unfortunately, we observe that about 9.3% of bogon prefix announcements were never withdrawn. Upon further examination, we found that some of these networks were still in the routing tables and reachable by traceroute. Figure 5 (a) shows the cumulative distribution of each event duration for *prefix-based events*. Over 47% of these last longer than 1 hour; among these, over 57% lasted longer than a day. Of the prefix-based events that lasted longer than an hour, the vast majority were from 172.16.0.0/12. We also saw many long-lived announcements from 0.0.0.0/7. The fact that many leaks involving private IP prefixes (i.e., likely misconfigurations) last longer than an hour suggests that many accidental events are not corrected immediately.

We also study the persistence of *origin AS-based events*. Figure 5 (b) shows the persistence of origin AS-based events. As shown by the "event time" line, nearly 30% of all origin AS-based events consisted of a single bogon prefix announcement without a corresponding withdrawal (hence, an event duration of zero seconds). Some origin AS-based events do not include a withdrawal because they "end" (i.e., they are followed by at least an hour of quiescence and followed by another sequence of announcements). The "time to *first* withdrawal" line indicates that only 50% of origin AS-based events contained any withdrawal messages; the median amount of time to the arrival of the first withdrawal for the other 50% of origin AS-based events that had a withdrawal is about an hour. The "time to *final* withdrawal" message indicates that only 30% of the origin AS-based events ended in a withdrawal; about half of the origin AS-based events that did end in a withdrawal lasted longer than one day. Both graphs in Figure 5 indicate that many bogon route announcements persist in the routing table for longer than an hour or even a day. The median length of all origin AS-based events was longer than an hour, and roughly a quarter of these events last longer than a week.

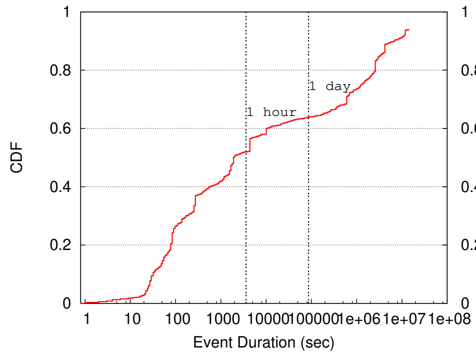

(a) Prefix-based events.

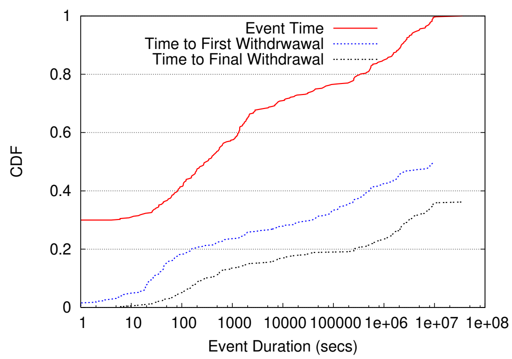

(b) Origin AS-based events.

Figure 5: Persistence of bogon prefix-based events.

4.2 *Do multiple ASes advertise the same bogon?*

We suspected that certain bogon prefixes and prefix space might be advertised by multiple origin ASes, particularly if the announcement was accidental. Since private address space is commonly used to address infrastructure inside a single AS, we expected to see different origin ASes leaking routes from the same bogon IP address space or even leaking the same prefix.

Table 3 also summarizes the bogon prefix *space* from which multiple ASes advertised routes. We initially thought that many leaked bogon prefixes would come from private address space (e.g., 172.16.0.0/12, 10.0.0.0/8, etc.), but other portions of unallocated space are also commonly advertised and leaked to portions of the Internet. In particular, 71 different ASes leaked routes from the 0.0.0.0/7 space.

Looking in more detail at the actual bogon *prefixes* that multiple ASes advertised, Table 4 also shows that some of the prefixes commonly leaked by multiple ASes from 0.0.0.0/7 include 0.1.0.0/16, 0.0.0.0/16, 1.1.1.0/24, and 1.0.0.0/8, which might be test routes or otherwise internal routes that were mistakenly leaked to the global Internet. Interestingly, many of these routes appear to be default routes that some ASes filter while others do not, because announcements for these prefixes are not visible at many of our monitors. On the other hand, 1.0.0.0/8 was visible at *six* of our eight monitors, considerably more than any prefixes in 0/8. This observation indicates that many ASes are likely filtering 0.0.0.0/8, but are *not* filtering 0.0.0.0/7. We asked the operator of the network where one of our iBGP monitors saw this prefix, who confirmed that this was the case for his network.

Bogon Prefix	*Prefix-based Ev.*	*Origin AS-based Ev.*	*Monitors*	*Origin ASes*
192.0.2.0/24	17	40	2	3
0.1.0.0/16	32	32	2	32
72.1.64.0/19	82	28	8	2
1.1.1.0/24	150	19	8	9
0.0.0.0/16	9	9	3	8
169.254.0.0/16	6	5	3	4
1.0.0.0/8	29	5	6	4
0.16.0.0/23	3	3	2	3
2.0.0.0/8	13	3	5	2
0.0.0.0/13	2	2	1	2
50.0.0.0/8	6	2	4	2
99.0.0.0/8	6	2	4	2
0.16.0.0/13	2	2	1	2
0.16.0.0/17	2	2	1	2
100.0.0.0/8	6	2	4	2

Table 4: Bogon IP *prefixes* that were originated by at least 2 distinct origin ASes, sorted by the number of distinct origin ASes that leaked them. Announcements of 1.1.1.0/24 were prevalent, both in the number of origin AS-based events and the number of ASes that announced it. Many other routes appear to be routes that were mistakenly leaked.

4.3 *Which ASes leak bogon prefixes, and who sees them?*

By examining the origin AS in the AS path corresponding for each bogon prefix announcement, we observed 110 different ASes originated bogon prefixes. Table 5 lists all 19 ASes that caused at least three origin AS-based events; the remaining 91 ASes only caused one or two origin AS-based events each, although some of these events were quite large: on March 25, 2004, AS 19962 leaked 87 bogon prefixes to Cogent, who passed these bogon routes to MIT. Table 5 shows that a few ASes are responsible for a large number of prefix-based and origin AS-based events. Most of these events were typically visible at only one or two monitoring points, implying that, despite their frequency, these bogon route leaks are not visible at many places on the Internet. However, the routes leaked by large tier-1 ISPs (shown in bold) were visible at nearly all of our monitoring points. This observation has important implications: routes that are advertised by smaller ASes are likely to be filtered by some upstream ISP, although there is no guarantee of this. However, routes with bogon prefixes that are originated by large ISPs are almost always widely observed. Our observation is consistent with the commonly held belief that many tier-1 ISPs filter routes based on the AS path attribute only, not based on specific prefixes. A malicious party that wanted to inject widely visible bogon prefixes would appear to have reasonable success injecting these prefixes from a tier-1 ISP.

4.4 *How many prefixes leak per origin AS-based event?*

We wanted to know how many bogon prefixes were leaked in any given origin AS-based event (i.e., some measure of the magnitude of the event). Origin AS-based events that leak large numbers of bogon prefixes may suggest the temporary misconfiguration of a filter. Figure 6 shows the distribution of the number of distinct prefixes that an origin AS leaked in each bogon prefix event. Most events resulting in the leaks of bogon prefixes involved only a handful of bogon prefixes, and the majority of events involved only one or two prefixes, but 31 events leaked more than 10 prefixes, and 14

Origin AS	*Prefix-based Ev.*	*Origin AS-based Ev.*	*Prefixes*	*Monitors*
10753	3191	65	214	1
577	23	38	2	6
1276	121	21	10	1
32880	82	18	3	8
3356	115	17	5	2
23504	8	15	1	7
7563	33	13	1	2
1239	383	11	6	8
13536	8	10	1	6
16482	16	10	1	5
4471	28	9	1	6
26230	6	7	1	6
701	50	6	1	7
30528	8	6	1	6
1299	43	6	5	7
1	38	6	3	1
3845	96	4	2	8
1784	4	3	2	3
9064	2	3	1	2

Table 5: Distribution of bogon prefix announcements and the ASes that originate them, for all ASes that caused at least three origin AS-based events. Tier-1 ISPs are shown in bold; routes originated from these networks tend to be more widely visible (i.e., we observe them at more of our monitors.)

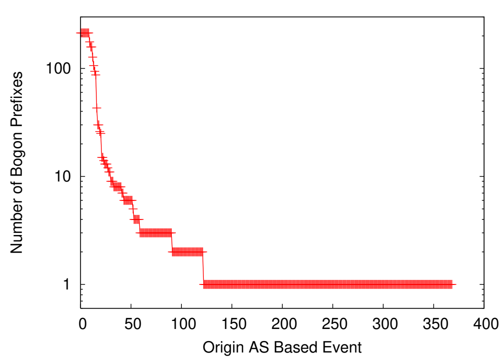

Figure 6: Distribution of number of bogon prefixes seen for each origin AS-based event. All events larger than 10 prefixes were originated by AS 10753.

events involved more than 100 prefixes. Upon further examination, we discovered that a single AS, AS 10753, was responsible for all origin AS-based events involving more than 100 prefixes, and 29 of the 31 events involving at least 10 prefixes. The exceptions were two distinct events where AS 19962 leaked 87 and 93 bogon prefixes, respectively.

4.5 Do ASes update filters when bogons are allocated?

Maintaining up-to-date prefix filters is important: it not only prevents malicious misuse of the reserved IP space, but also allows global reachability by ensuring that routes from previously bogon address space is widely advertised. During the 15 months of our dataset, several regions of previously bogon IP address space were allocated to regional Internet registries by IANA; Table 6 summarizes when various "former bogons" were allocated. Because operators regularly complain on the NANOG mailing list [20] that other operators are slow to update their route filters when new prefix space is allocated, we hypothesized that routes advertised from this space would remain filtered by certain ASes for a considerable amount of time after the allocation. (RIPE has a project specifically aimed at allowing ASes to test reachability to active prefixes within recently allocated IP address space [3].)

[4]58.0.0.0/8 and 59.0.0.0/8 were allocated on April 28, 2004; 71.0.0.0/8 was allocated on August 2, 2004. We have not yet seen any announcements from this prefix space.

Prefix	*Allocated*	*Monitors that saw it initially (and now)*	*Days until first announcement*	*Days until Full Visibility*
83.0.0.0/8	Nov 16, 2003	6 (8)	15	150
84.0.0.0/8	Nov 16, 2003	6 (8)	100	150
70.0.0.0/8	Jan 15, 2004	7 (8)	41	90
86.0.0.0/7	Apr 1, 2004	1 (7)	21	—
88.0.0.0/8	Apr 1, 2004	1 (7)	21	—
85.0.0.0/8	Apr 1, 2004	6 (8)	5	176
72.0.0.0/8	Aug 2, 2004	8 (8)	39	39

Table 6: Visibility of routes from previously bogon IP address space. [4]

We can estimate such a phenomenon from our observations when announcements from newly allocated IP address space were not visible at all of our monitoring points. With the exception of a single announcement for 85.0.0.0/8 observed at MIT on March 25, 2004 and 72.0.0.0/8, which we observed regularly at all monitoring points, we did not observe route announcements for any of these prefixes prior to when they were allocated; as such, we cannot determine whether operators removed these prefixes from their filters or were never filtering them in the first place. However, our previous analysis in Table 3 suggests that most leaks for unallocated space are not visible at all 8 of our monitors, so it is reasonable to assume that announcements for these newly allocated prefixes would not have been visible at all of our monitors prior to being allocated. Table 6 also shows that the initial announcement of a prefix from newly allocated space was never visible at all 8 of our monitors.

As of October 9, 2004, announcements from AS 12654 for 8 prefixes from the two prefix spaces allocated on April 1, 2004, were still not visible at the Aros iBGP monitor. Because these announcements are visible at all but one of the monitors, it is likely that at least one regional AS is filtering this prefix space.

5. Summary and Future Work

The integrity and robustness of the interdomain routing system requires ASes to correctly filter invalid route advertisements, but our preliminary findings indicate that many ASes appear to leak invalid routes and fail to filter invalid routes. Many invalid routes, particularly those that originated from certain tier-1 ISPs, were visible at nearly all of our monitoring points. An attacker with access to a router in one of these networks could inject invalid routes that would likely reach a significant fraction of ASes. ASes that *do* filter invalid routes often do not update their route filters until many months after new prefix space is allocated, thus mistakenly filtering *valid* routes. Clearly, more diligent filtering or better route authentication is needed, but a reasonable first step would be for ASes to simply deploy filters for *private* address space; such filters are static (i.e., they don't need to be updated as prefix space is allocated) and would block the majority of invalid route leaks. Routers could even filter private address space by default.

We plan to extend our analysis in several ways. First, we plan to study the properties of BGP updates, such as the AS path, to infer information such as which ASes are (or are not) filtering certain invalid prefixes. Second, we would like to study the forwarding path properties of invalid routes by running traceroutes to invalid routes when they are received at our monitoring points. Where possible, we would also like to monitor traffic from bogon prefixes to determine whether networks are using these invalid routes for nefarious activities (e.g., spam, DDoS, etc.). Finally, using information about possibly hijacked prefixes [11] as a starting point, we plan to extend our analysis to hijacked prefixes.

6. References

[1] American Registry for Internet Numbers. `http://www.arin.net/`.

[2] Asia Pacific Network Information Center. `http://www.apnic.net/`.

[3] De-bogonising new address blocks. `http://www.ris.ripe.net/debogon/debogon.html`.

[4] Internet Assigned Numbers Authority. `http://www.iana.org/`.

[5] Internet Protocol V4 Address Space. `http://www.iana.org/assignments/ipv4-address-space`.

[6] Regional Latin-America and Caribbean IP Address Registry. `http://lacnic.net/en/index.html`.

[7] Réseaux IP Européens. `http://www.ripe.net`.

[8] Bogon Report. `http://www.cidr-report.org/bogons/`, Nov. 2004.

[9] Possible Bogus Routes and AS Announcements. `http://www.cidr-report.org/#Bogons`, Nov. 2004.

[10] CAESAR, M., SUBRAMANIAN, L., AND KATZ, R. Towards localizing root causes of BGP dynamics. Tech. Rep. UCB/CSD-04-1302, U.C. Berkeley, Nov. 2003.

[11] CompleteWhois. http://www.completewhois.com/, 2004.

[12] Team Cymru bogon route server project. `http://www.cymru.com/BGP/bogon-rs.html`.

[13] FEAMSTER, N., AND BALAKRISHNAN, H. Detecting BGP Configuration Faults with Static Analysis. In *Proc. 2nd Symposium on Networked Systems Design and Implementation* (Oct. 2004).

[14] HUSTON, G. Hunting the bogon. `http://www.potaroo.net/papers/isoc/2004-04/bogons.html`, Apr. 2004.

[15] IANA. *Special Use IPv4 Addresses*. Internet Engineering Task Force, Sept. 2002. RFC 3330.

[16] KAZUYOSHI, S. 1.1.1.0/24. `http://www.cctec.com/maillists/nanog/historical/9805/msg00258.html`.

[17] KRUEGEL, C., MUTZ, D., ROBERTSON, W., AND VALEUR, F. Topology-based detection of anomalous BGP messages. In *RAID* (2003), vol. 2820 of *Lecture Notes in Computer Science*, Springer.

[18] LABOVITZ, C., AHUJA, A., BOSE, A., AND JAHANIAN, F. Delayed Internet Routing Convergence. *IEEE/ACM Transactions on Networking 9*, 3 (June 2001), 293–306.

[19] MAHAJAN, R., WETHERALL, D., AND ANDERSON, T. Understanding BGP misconfiguration. In *Proc. ACM SIGCOMM* (Pittsburgh, PA, Aug. 2002), pp. 3–17.

[20] The North American Network Operators' Group mailing list archive. `http://www.cctec.com/maillists/nanog/`.

[21] REKHTER, Y., AND LI, T. *A Border Gateway Protocol 4 (BGP-4)*. Internet Engineering Task Force, Mar. 1995. RFC 1771.

[22] TODD, J. AS number inconsistencies, July 2002. `http://www.merit.edu/mail.archives/nanog/2002-07/msg00259.html`.

[23] GNU Zebra. `http://www.zebra.org/`.

[24] ZHANG, K., YEN, A., ZHAO, X., MASSEY, D., AND WU, S. F. On detection of anomalous routing dynamics in BGP. In *IFIP-TC6 Networking 2004* (2004), vol. 3042 of *Lecture Notes in Computer Science*, Springer.

Appendix

A. Bogon Prefixes

As of December 8, 2004, there are 94 /8 IPv4 address blocks that are reserved or not allocated to any registries [5]. Table 7 lists those unallocated IP address ranges in dotted notation.

Address range
0.0.0.0 - 2.255.255.255
5.0.0.0 - 5.255.255.255
7.0.0.0 - 7.255.255.255
23.0.0.0 - 23.255.255.255
27.0.0.0 - 27.255.255.255
31.0.0.0 - 31.255.255.255
36.0.0.0 - 37.255.255.255
39.0.0.0 - 39.255.255.255
41.0.0.0 - 42.255.255.255
49.0.0.0 - 50.255.255.255
73.0.0.0 - 79.255.255.255
89.0.0.0 - 126.255.255.255
173.0.0.0 - 187.255.255.255
189.0.0.0 - 190.255.255.255
223.0.0.0 - 223.255.255.255
240.0.0.0 - 255.255.255.255

Table 7: Unallocated IP address space

In addition, RFC 3330 [15] defines several address blocks that are reserved for specialized purposes and should not appear on the public Internet (Table 8).

Address block	Allocated for
10.0.0.0/8	Private address space
127.0.0.0/8	Loopback
169.254.0.0/16	Communication between hosts on a single link
172.16.0.0/12	Private address space
192.0.2.0/24	Use in documentation and example code
192.168.0.0/16	Private address space
198.18.0.0/15	Benchmark tests of network interconnect devices
224.0.0.0/4	IPv4 multicast

Table 8: Specialized address blocks defined in RFC 3330