

Understanding Resiliency of Internet Topology Against Prefix Hijack Attacks *

Mohit Lad [†]

Ricardo Oliveira *

Beichuan Zhang [‡]

Lixia Zhang *

Abstract

A prefix hijack attack involves an attacker announcing victim networks' IP prefixes into the global routing system. As a result, data traffic from portions of the Internet can be diverted to attacker networks. Prefix hijack attacks are a serious security threat in the Internet and it is important to understand the factors that affect the resiliency of victim networks against these attacks. In this paper, we conducted a systematic study to gauge the effectiveness of prefix hijacks launched at different locations in the Internet topology. Our study shows that direct customers of multiple tier-1 networks are the most resilient, even more than the tier-1 networks themselves. Conversely, if these customer networks are used to launch prefix hijacks, they would also be the most effective launching pads for attacks. We verified our results through case studies using real prefix hijack incidents that had occurred in the Internet.

1 Introduction

On January 22, 2006, a network (AS-27506) wrongly announced the IP prefix 65.173.134.0/24 representing an address block of 2^{24} IP addresses, into the global routing system. This prefix belonged to another network (AS-19758) and because routers do not have a means to accurately verify the legitimate origin of each prefix, they accepted announcements from both the true origin (AS-19758) and the false one (AS-27506), and selected one of them based on the local routing policies and other criteria. As a result, some networks sent for data traffic destined to 65.173.134.0/24, to AS-27506 instead of the true owner. This is a typical incident of a prefix hijack, where a network announces an

address space it does not own and *hijacks* traffic destined to the true owner.

Prefix hijacking is a serious security threat in the Internet. Prefix hijacks can potentially be launched from any part of the Internet and can target any prefix belonging to any network. A hijack attack has a large impact if the majority of routers choose the path leading to the false origin. Conversely, if the majority of routers choose the path leading to the true origin, the network of the prefix owner is considered to be resilient against prefix hijack attacks. Although there have been several results on preventing prefix hijacks (e.g., [6][11]) and monitoring potential prefix hijack attempts (e.g., [8, 10]), there is a lack of a general understanding on the impact of a successful prefix hijack and networks' resiliency against such attacks. This lack of understanding makes it difficult to assess the overall damage once an attack occurs, and to provide guidance to network operators on how to improve their networks' resilience.

In this paper, we conduct a systematic study to gauge the impact of prefix hijacks launched at different locations in the Internet topology, and identify topological characteristics of those networks that are most resilient against hijacks of their prefixes. Specifically, we deal with a type of prefix hijack referred to as false origin hijacks where a network announces the exact prefix announced by another network. Using simulations on an Internet scale topology and measurements from real data, we estimate how many nodes in the Internet may believe the true origin and how many believe the false origin during a hijack. Our results show that the Internet topology hierarchy and routing policies play an essential role in determining the impact of a prefix hijack. Our study shows that the high degree networks (e.g., tier-1 ISPs) are not necessarily most resilient against prefix hijacks. Instead, small networks that are direct customers to multiple tier-1 ISPs are seen to be most resilient. Conversely, attacks launched from these multi-homed customer networks would also have the biggest impact. Implications of our results are twofold. First, networks that desire high resilience against prefix hijacks should connect to multiple providers, and be as close as possible to multiple tier-1 ISPs and networks that cannot achieve such topological connectivity, should use reactive means to learn about their prefix being hijacked. Second, securing only the big ISP networks is not adequate nor effective, since high impact attacks come

*This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No N66001-04-1-8926 and by National Science Foundation(NSF) under Contract No ANI-0221453. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the DARPA or NSF.

[†]University of California, Los Angeles. 4732 Boelter Hall, Los Angeles, CA 90095. Email: {mohit,rveloso,lixia}@cs.ucla.edu

[‡]University of Arizona. 1040 E. 4th Street Tucson, AZ 85721. Email: bzhang@cs.arizona.edu

from well connected small networks.

The rest of the paper is organized as follows. Section 2 reviews Internet routing and prefix hijacking. Section 3 defines evaluation metrics and Section 4 uses simulations on an Internet scale topology to evaluate the resiliency of different networks. Section 5 presents evidence of our findings in real hijack incidents. Section 6 discusses the insights and implications of our findings. Section 7 presents related work and Section 8 concludes the paper.

2 Background

In this section we present the relevant background on Internet routing and describe prefix hijacking with an example.

2.1 Internet Routing

The Internet consists of more than twenty thousand networks called “Autonomous Systems” (AS). Each AS is represented by a unique numeric ID known as its AS number, and may advertise one or more IP address prefixes. For example, the prefix 131.179.0.0/16 represents a range of 2^{16} IP addresses belonging to AS-52 (UCLA). Internet Registries such as ARIN and RIPE assign prefixes to organizations, who then become the owner of the prefixes. Autonomous Systems run the Border Gateway Protocol (BGP) [15] to propagate prefix reachability information among themselves. In the rest of the paper, we abstract an autonomous system into a single entity called *AS node* or *node*, and the BGP connection between two autonomous systems as *AS link* or simply *link*.

BGP uses routing update messages to propagate routing changes. As a path-vector routing protocol, BGP lists the entire AS path to reach a destination prefix in its routing updates. Route selection and announcement in BGP are determined by networks’ routing policies, in which the business relationship between two connected ASes plays a major role. AS relationship can be generally classified as customer-provider or peer-peer¹. In a customer-provider relationship, the customer AS pays the provider AS for access service to the rest of the Internet. The peer-peer relationship does not usually involve monetary flow; The two peer ASes exchange traffic between their respective customers only. Usually a customer AS does not forward traffic between its providers, nor does a peer AS forward traffic between two other peers. For example in Figure 1, AS-1 is a customer of AS-2 and AS-3, and hence would not want to be a transit between AS-2 and AS-3, since it would be pay both AS-2 and AS-3 for traffic exchange between themselves. This results in the so-called *valley-free* BGP paths [3] generally observed in the Internet. When ASes choose their best

¹Sometimes the relationship between two AS nodes can be “siblings,” usually because they belong to the same organization.

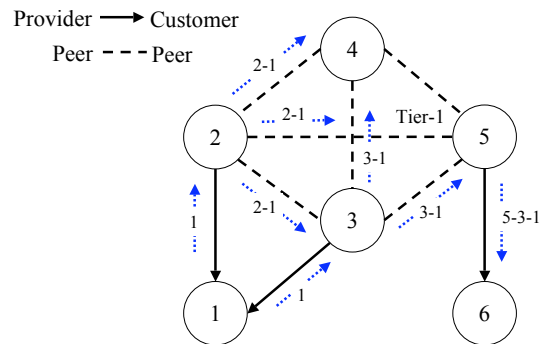


Figure 1. Route propagation.

path, they usually follow the order of customer routes, peer routes, and provider routes. This policy of *no valley prefer customer* is generally followed by most networks in the Internet. As we will see later, the *no valley prefer customer* policy plays an important role in determining the impact of prefix hijacks and hence we present a simple example to illustrate how this policy works.

Figure 1 provides a simple example illustrating route selection and propagation. AS-1 announces a prefix (e.g. 131.179.0.0/16) to its upstream service providers AS-2 and AS-3. The AS announcing a prefix to the rest of the Internet is called the *origin AS* of that prefix. Each of these providers then prepends its own AS number to the path and propagates the path to their neighbors. Note that AS-3 receives paths from its customer, AS-1, as well as its peer, AS-2, and it selects the customer path over the peer path thus advertising the path {3 1} to its neighbors AS-4 and AS-5. AS-5 receives routes from AS-2 and AS-3 and we assume AS-5 selects the route announced by AS-3 and announces the path {5 3 1} to its customer AS-6. In general, an AS chooses which routes to import from its neighbors and which routes to export to its neighbors based on import and export routing policies. An AS receiving multiple routes picks the best route based on policy preference. Metrics such as path length and other BGP parameters are used in route selection if the policy is the same for different routes. The BGP decision process also contains many more parameters that can be configured to mark the preference of routes. A good explanation of these parameters can be found in [4].

2.2 Prefix Hijacking

A prefix hijack occurs when an AS announces prefixes that it does not own. Now, suppose AS-6 wrongly announces the prefix that belongs to AS-1, as shown in Fig-

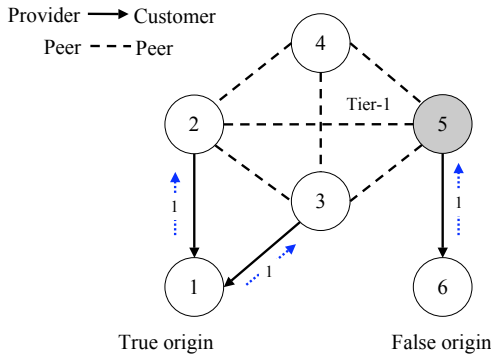


Figure 2. Hijack scenario.

Figure 2. Note that AS-5 previously routed through AS-3 to reach AS-1. On receiving a customer route through AS-6, it prefers the customer route over the peer route and hence believes the false route. This is an example of a prefix hijack, in which a false origin AS-6 announces a prefix it does not own, and deceives AS-5. In current routing practice, it is difficult for an AS to differentiate between a true origin and a false origin. Even though Internet Routing Registries (IRR) provide databases of prefix ownership, the contents are not maintained up-to-date, and not all BGP routers are known to check these databases. Hence, when presented multiple paths to reach the same prefix, a BGP router will often choose the best path regardless of who originates this prefix, thus allowing hijacked routes to propagate through the Internet. Prefix hijacks can be due to malicious attacks or router mis-configurations. When legitimate data traffic is diverted to the false origin, the data may be discarded, resulting in a traffic blackhole, or be exploited for malicious purposes. A recent study [14] reported that some spammers hijack prefixes in order to send spams without revealing their network identities.

The hijack depicted in Figure 2 is called a false origin prefix hijack, where an AS announces the *exact* prefix owned by another AS. Another type of hijack, called sub-prefix hijack, involves an AS announcing a more specific prefix (e.g. hijacker announces a /24, when the true origin announces a /16). In this case, BGP routers will usually treat them as different prefixes and maintain two separate entries in routing tables. However, due to longest prefix matching in routing table lookups, data destined to IP addresses in the /24 range will be forwarded to the false origin, instead of the true origin. Prefix hijack can also involve a false AS link advertised in the AS path without a change of origin. Our aim in this paper is to understand that, how the topological characteristics of two AS nodes announcing the same prefix influence the impact of the hijack. Studying the impact of

sub-prefix hijacks and false link hijacks involves different considerations and is beyond the scope of this paper.

Terminology

In the rest of this paper, we use the term *prefix hijacks* to refer to false origin prefix hijacks. We call the AS announcing a prefix it does not own as the *false origin*, and the AS whose prefix is being attacked as the *true origin*. Upon receiving the routes from both the false origin as well as the true origin, an AS that believes the false origin is said to be *deceived*, while an AS that still routes to the true origin is said to be *unaffected*.

3 Hijack Evaluation Metrics

For our simulations, we model the Internet topology as a graph, in which each node represents an AS, and each link represents a logical relationship between two neighboring AS nodes. Note, two neighboring nodes may have multiple physical links between themselves. However, BGP paths are represented in the form of AS-AS links, and hence we abstract connections between two AS nodes as a single logical link. For simplicity, each node owns exactly one unique prefix, i.e. no two nodes announce the same prefix except during hijack. A prefix hijack at any given time involves only one hijacker, and the hijacker can target only one node.

To capture the interaction between the entities involved in a hijack, we introduce a variable $\beta(a, t, v)$, function of false origin a , true origin t and node v as follows:

$$\beta(a, t, v) = \begin{cases} 1 & : \text{if node } v \text{ is deceived by false} \\ & \text{origin } a \text{ for true origin } t\text{'s prefix} \\ 0 & : \text{otherwise} \end{cases} \quad (1)$$

Due to the rich connectivity in Internet topology, a node often has multiple equally good paths to reach the same prefix. Figure 2 shows a case where AS-4 has three equally good paths to reach the same prefix, two to the true origin AS-1 (through AS-2 and AS-3), and one to the false origin AS-6. In our model, we assume a node will break the tie randomly. Therefore, we define the expected value of β as follows. Let $p(v, n)$ be the number of equally preferred paths (e.g. same policy, same path length) from the node v to node n . E.g., in Figure 2, $p(4, 1) = 2$ since AS-4 has two paths via AS-2 and AS-3 to reach AS-1, and $p(4, 6) = 1$ since AS-4 has only one route via AS-5 to reach AS-6. If nodes use random tie-break to decide between multiple equally good preferred paths, then the expected value for β is defined as:

$$\bar{\beta}(a, t, v) = \frac{p(v, a)}{p(v, a) + p(v, t)} \quad (2)$$

yielding $\bar{\beta}(6, 1, 4) = \frac{1}{3}$ for the example in the figure. $\bar{\beta}$ is the probability of a node v being deceived by a given false origin a announcing a route belonging to true origin t .

Impact

We use the term *impact* to measure the attacking power of a node launching prefix hijacks. We define impact of a node a as the fraction of the nodes that believe the false origin a during an attack on true origin t . More formally, the impact of a node a is given by:

$$I(a) = \sum_{t \in \mathcal{N}} \sum_{v \in \mathcal{N}} \frac{\bar{\beta}(a, t, v)}{(N-1)(N-2)} \quad (3)$$

Note that the outer sum is over $N-1$ true origins (we exclude the false origin) and the inner sum is over $N-2$ nodes (excluding both the false origin and true origin).

Resilience

We use the term *resilience* to measure the defensive power of a node against hijacks launched against its prefix. We define the resilience of a node t as the fraction of nodes that believe the true origin t given an arbitrary hijack against t . More formally, the node resilience $R(t)$ of a node t is given by:

$$R(t) = \sum_{a \in \mathcal{N}} \sum_{v \in \mathcal{N}} \frac{\bar{\beta}(t, a, v)}{(N-1)(N-2)} \quad (4)$$

Note, higher $R(t)$ values indicate better resilience against hijacks, and higher $I(a)$ values indicate higher impact as an attacker.

Relation between Impact and Resilience

The true origin t and false origin a compete with each other to make nodes in the Internet route to itself. For example in Figure 2, false origin AS-6 is hijacking a prefix belonging to true origin AS-1. In this case, only AS-5 believes the false origin and AS-4 has a $1/3$ chance of being deceived. Therefore, the chances that a node believes the false origin AS-6 when it hijacks AS-1 is given by $\frac{1+1/3}{4} = \frac{1}{3}$.

Now if AS-1 was to hijack a prefix belonging to AS-6, then AS-5 would still believe AS-6 and AS-4 will believe it with a probability of $1/3$. Thus, in this case, the chances that a node believes the true origin AS-6 when it is hijacked by AS-1 is $\frac{1+1/3}{4} = \frac{1}{3}$.

We see that the resilience of the node as a true origin is equal to its impact as a false origin. We note that in our model, when the roles of attacker and target are switched, the impact of a node becomes its resilience. In the rest of the paper, we focus on resilience, while keeping in mind that a highly resilient node can also cause high impact as a false origin.

4 Evaluating Hijacks

In this section, we aim to understand the topological resilience of nodes against prefix hijacks by performing sim-

ulations on an Internet derived topology. We first explain the simulation setup, followed by the main results of our simulation and the insight behind the results.

4.1 Simulation Setup

For our simulations, we use an AS topology collected from BGP routing tables and updates, representing a snapshot of the Internet as of Feb 15 2006 (available from [19]). The details of how this topology was constructed are described in [20]. Our topology consists of 22,467 AS nodes and 63,883 links. We assume each AS node owns and announces a single prefix to its neighbors. We classify AS nodes into three tiers: Tier-1 nodes, transit nodes, and stub nodes. To choose the set of Tier-1 nodes, we started with a well known list, and added a few high degree nodes that form a clique with the existing set. Nodes other than Tier-1s but provide transit service to other AS nodes, are classified as *transit* nodes, and the remainder of nodes are classified as *stub* nodes. This classification results in 8 Tier-1 nodes, 5,793 transit nodes, and 16,666 stub nodes. We classify each link as either customer-provider or peer-peer using the PTE algorithm[3] and use the *no valley prefer customer* routing policy to infer routing paths (also used in previous works such as [18]). We abstracted the router decision process into the following priorities (1)local policy based on relationship, (2)AS path length, and (3)random tie-breaker.

Of the 22,467 AS nodes in our topology, we randomly picked 1,000 AS nodes to represent false origins that would launch attacks on other AS nodes. We checked the degree distribution of this set of 1,000 AS nodes, and found it to be similar to the degree distribution of all the AS nodes. For each of the 22,467 AS nodes as a true origin, we simulated a hijack with the 1,000 false origins. Thus we simulated $22,467 \times 1,000 \simeq 22.5$ million hijack scenarios in total.

4.2 Characterizing Topological Resilience

Figure 3 shows the distribution of the resilience (average curve) for all the nodes in our topology from our simulated hijacks. Since the resilience of each node results from the average over 1,000 attackers, we also show the standard deviation range. Note, higher values of resilience imply more resilience against hijacks.

This distribution shows that node resilience varies fairly linearly except at the two extremes. Figure 3 also shows that the deviations at the two extremes are quite small compared to the middle, indicating that some nodes(top left) are very resilient against hijacks, while some others (bottom right) are easily attacked, regardless of the location of the false origin.

As a first step in understanding how different nodes differ in their resilience, we classify nodes into the three classes

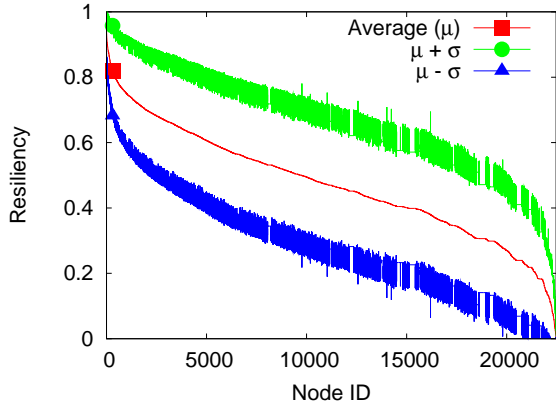


Figure 3. Distribution of node resilience.

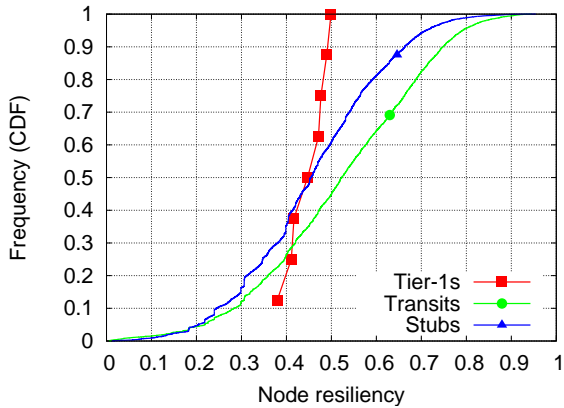


Figure 4. Resilience of nodes in different tiers.

already described: tier-1, transit and stub and plot the average resilience distribution (CDF) of each class of nodes in Figure 4. We observe that the resilience distribution is very similar for transits and stubs, with transit nodes being a little more resilient than stubs.

In contrast, tier-1 nodes show a very different distribution from the stubs and transits. From Figure 4 we observe that all the tier-1 nodes have an average resilience value between 0.4 and 0.5. In addition, we note that about 40% of stubs and 55% of transit nodes are more resilient than all tier-1 nodes. With tier-1 nodes being the ones with the highest degree, it is surprising to see that close to 50% of the nodes in the Internet are more resilient than tier-1s. Next, we explain why tier-1 nodes are more vulnerable to hijacks than a lot of other nodes and generalize this explanation to understand the characteristics impacting resilience.

4.3 Factors Affecting Resilience

We first understand the resilience of tier-1 nodes with a simple hijack scenario in Figure 5. AS-2, AS-3, AS-4 and AS-5 represent 4 tier-1 nodes inter-connected through a peer-peer relationship. AS-1 and AS-6 are small ISPs connected to tier-1 AS nodes through a customer-provider relationship. Finally AS-7 is a multi-homed customer of AS-1 and AS-6. In Figure 5, AS-7 represents the false origin that hijacks a prefix belonging to a tier-1 node, AS-4.

Recall in no-valley prefer customer policy, a customer route is preferred over a peer route which in turn is preferred over a provider route. When AS-7 hijack's AS-4's prefix and announces the false route to AS-1 and AS-6, both AS-1 and AS-6 prefer the hijacked route over the genuine route to AS-4 since its a customer route. AS-1 in turn announces the hijacked route to its tier-1 providers AS-2 and AS-3. These tier-1 AS nodes, AS-2 and AS-3 now have to choose between a customer route through AS-1(hijacked route), and a peer route through AS-4 (genuine route). Again due to policy preference, the tier-1 nodes will choose the customer route which happens to be the hijacked route. Similarly, AS-5 will also choose the hijacked route. Once big ISPs like tier-1 nodes are deceived by the hijacker, their huge customer base (many of whom are single homed) are also deceived, thus causing a high impact. One can see from this example, that the main reason for the low resilience in the case of a hijack on a tier-1 node is that tier-1 nodes inter-connect through peer-peer relationship thus rendering a genuine route less preferred to other tier-1 nodes than hijacked routes from customers.

The key to high resilience is to make the tier-1 nodes and other big ISPs always believe the true origin. The way to achieve this is to reach as many tier-1 nodes as possible using a provider route. In addition, when a node has to choose between two routes of the same preference, path length becomes a deciding factor, and thus the shorter the number of hops to reach the tier-1 nodes, the better the resilience. From our observations from simulation results, we found that the most resilient nodes are direct customers of many tier-1 nodes and other big ISPs. As an example, in our simulations, the node with highest resilience is a stub (AS-6432 DoubleClick) directly connected to 6 tier-1 nodes, having a resilience value of 0.95. The nodes with lowest resilience were single-home customers, connected to poorly connected providers.

To better understand the influence of tier-1 nodes, we classified the nodes in the Internet based on the number of direct tier-1 providers. Figure 6 shows the distribution of resilience for nodes with different connectivity to Tier-1. Note, the closer the curve to the right hand side of the figure ($x=1$), the better the resilience of that set of nodes. There are about 21,888 nodes with less than 3 connections to Tier-1, and we observe in Figure 6 that these nodes are the least

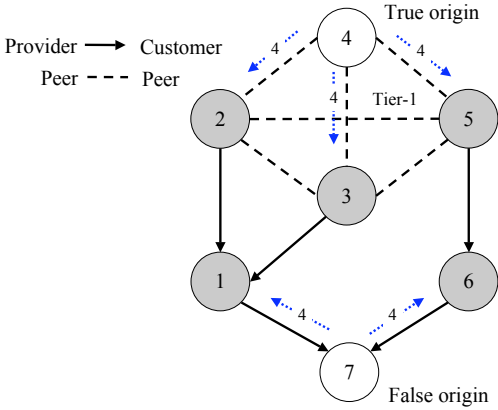


Figure 5. Understanding resilience of tier-1 nodes

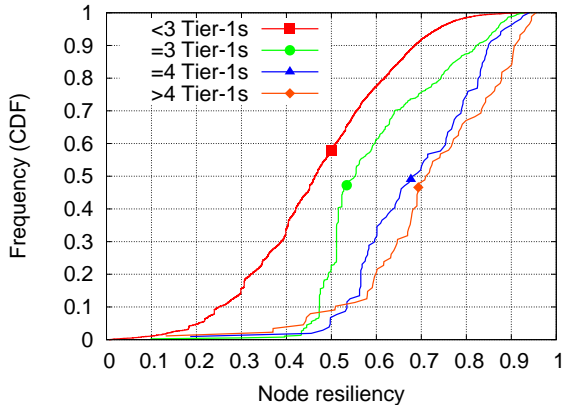


Figure 6. Resilience of nodes with different number of Tier-1 providers.

resilient. A total of 379 nodes are directly connected to 3 Tier-1s and 104 nodes are connected to 4 Tier-1s. Only 88 nodes are connected to more than 4 Tier-1s, and these nodes prove to be the most resilient, highlighting the role of connecting to multiple tier-1 nodes.

Summary: In this section, we used an Internet scale topology with no-valley prefer customer policy routing to evaluate the resilience of nodes against random hijackers. The key to achieve high resilience is to protect tier-1 nodes and other big ISPs from being deceived by the hijacker. Our main result shows that the nodes that are direct customers of multiple tier-1 nodes are the most resilient to hijacks. On the other hand, the tier-1 nodes themselves in spite of being so well connected, are much less resilient to hijack. The next question we seek to answer in Section 5 is whether there is evidence of such behavior in reality, where the routing deci-

sion process is much more complex.

5 Prefix Hijack Incidents in the Internet

In this section we examine two hijack events, one from January 2006 which affected a few tens of prefixes, and the other from December 2004 when over 100,000 prefixes were hijacked. To gauge the impact of the prefix hijacks, we analyzed the BGP routing data collected by the Oregon collector of the RouteViews project. The Oregon collector receives BGP updates from over 40 routers. These 40 routers belong to 35 different AS nodes (a few AS nodes have more than one BGP monitor) and we consider an AS as deceived by a hijack if at least one BGP monitor from that AS believes the hijacker. We call these 35 AS nodes as *monitors*, as they provide BGP monitoring information to the Oregon collector. The impact of a hijack is then gauged by the ratio of monitors in the Internet that were deceived.

5.1 Case I: Prefix Hijacks by AS-27506

On January 22, 2006, AS-27506 announced a number of prefixes that did not belong to it. This hijack incident was believed to be due to operational errors, and most of the hijacked prefixes were former customers of AS-27506. We observed a total of 40 prefixes being hijacked by AS-27506. These 40 prefixes belonged to 22 unique ASes. We present two representative prefixes; for the first prefix the false origin could only deceive a small number of monitors, while for the second prefix the false origin deceived the majority of the monitors. We examine the topological connectivity of the true origins as compared to that of the false origin and the relation to the true origin's resiliency.

5.1.1 High Resiliency against Hijack

We examine a hijacked prefix that belongs to the true origin AS-20282. The impact of hijacking this prefix is just over 10%, that is 4 out of the 35 monitored ASes were deceived by the hijack. Figure 7(a) depicts the connectivity of some of the entities involved in this hijack incident. The nodes colored in gray are the nodes deceived by the false origin AS-27506, and the white nodes persisted with the true origin. The true origin AS-20282 is a direct customer of two tier-1 nodes, AS-701 and AS-3356. Before the hijack incident, all the 35 monitors used routes containing one of these two tier-1 ASes as the last hop in the AS path to reach the prefix. The hijacker AS-27506 is a customer of AS-2914, another tier-1 node. When AS-27506 hijacked the prefix, AS-2914 chose the false customer route from AS-27506 over an existing peer route through AS-701. The false route was further announced by AS-2914 to other tier-1 peers including AS-701 and AS-3356, however neither of them adopted the new route because they

chose the customer route announced by the true origin AS-20282. Other tier-1 ASes, such as AS-1239 (not shown in the figure), did not adopt to the false route from AS-2914 either, most likely because the newly announced false route was 2 hops in length, the same as that of their existing route through AS-701 or AS-3356, and the recommended practice suggests to avoid unnecessary best path transitions between equal external paths [2]. However we note that AS-3130, who is a customer of both a deceived and an unaffected tier-1 providers, also got deceived, possibly because the new path {2914, 27506} is shorter than the original path which contained 3 AS hops.

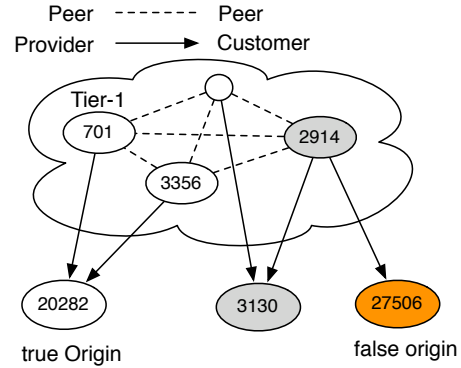
5.1.2 Low Resiliency against Hijack

Next, we examine another hijacked prefix which belonged to AS-23011. The average impact of this hijacked prefix is 0.6, i.e. 21 out of the 35 monitors were deceived by the hijack. Figure 7(b) shows the most relevant entities involved in this prefix hijack. The true origin of this prefix was an indirect customer of 5 tier-1 ASes (not all of them are shown in the figure) through its direct providers AS-12006 and AS-10910. The connectivity of the hijacker is the same as before, and AS-2914 was deceived by the hijack. The 5 tier-1 ASes on the provider path of the true origin stayed with the route from the true origin AS-23011, however the rest of the tier-1 ASes were deceived this time, possibly because the peer route to false origin through AS-2914 was shorter than any other peer route to the true origin. AS-286 is a customer of the providers of both the true and false origins, and it picked the false route through AS-2914 because it was shorter. We note that, in this case, the true origin being indirect customers of multiple tier-1 ASes ensured that those tier-1 ASes themselves did not get deceived, however due to its longer distance to reach these tier-1 providers (compared to the true origin in Figure 7(a)), other tier-1 ASes and their customers chose the shorter route to the false origin.

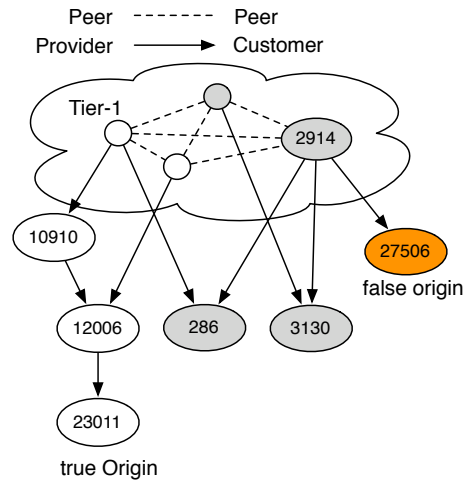
One of the tier-1 providers that propagated the false route is known to verify the origin of received routes with the Internet Routing Registries (IRR). However, it did not block the hijack because the registry entries were outdated and still listed AS-27506 as an origin for the hijacked prefixes, and hence the hijack announcements passed the registry check.

5.2 Case II: Prefix Hijacks by AS-9121

In this hijack incident, operational errors led AS-9121 to falsely announce routes to over 100,000 prefixes on December 24, 2004. We use this case to evaluate the resiliency of tier-1 ASes as compared to that of direct customers of multiple tier-1 ASes. Due to the large number of prefixes being falsely announced, some BGP protection mechanisms such as prefix filters and maximum prefix limit, where an AS sets an upper limit on the number of routes a given neighbor may



(a) High resiliency: Tier-1 provider 2914 preferred the customer route to false origin 27506 instead of the peer route. Similarly tier-1 providers 701 and 3356 stayed with their customer routes to the true origin 20282. Other tier-1 providers like X received a peer route to false origin that is no better than existing route and did not change route. 3130 routed to the false origin since the route via one of its providers, 2914, was shorter



(b) Low resiliency: Tier-1 providers like Y with a customer route to true origin 23011 were not deceived by false origin. Other tier-1 providers like X received a shorter peer route through 2914 and hence routed to false origin. 286 preferred the shorter route to 27506 via 2914 and was deceived.

Figure 7. Case study: AS-27506 as false origin

announce, were triggered and made an effect on the overall impact. Given that multiple factors were involved in such a large scale hijack event, it is difficult to accurately model the impact on an AS as a function of its topological connectivity. Our objective in examining this case is to find supporting evidence for our observations made in Section 4, as opposed to a detailed study over all the hijacked prefixes. Similar to case-1, we observed how many monitors were deceived for each hijacked prefix and used this result to gauge the resiliency of the true origin AS.

5.2.1 Hijacked Tier-1 AS Prefix

In order to understand how tier-1 ASes fared against AS-9121 hijack, we studied the impact of those hijacked prefixes that belonged to AS-7018, a tier-1 AS. Note that AS-7018 announced over 1500 prefixes, and the impacts of different prefixes varied noticeably, with around 7 to 8 monitors being deceived for most prefixes. For our case study, we examine one of the hijacked prefixes which deceived the majority of the monitors. Figure 8(a) shows the entities involved in the hijack of this tier-1 prefix.

The hijacker AS-9121 was connected to 3 providers, one of which was AS-1239, a tier-1 AS. The true origin of the prefix in question was AS-7018, another tier-1 AS. The grey nodes in the figure indicate those deceived by the hijack. All the 3 providers of AS-9121, namely AS-1239, AS-6762, and AS-1299 were deceived into believing the false origin. AS-1299 also propagated the false route to its tier-1 AS providers. From our observations, a total of 19 out of 35 monitors were deceived by this hijack.

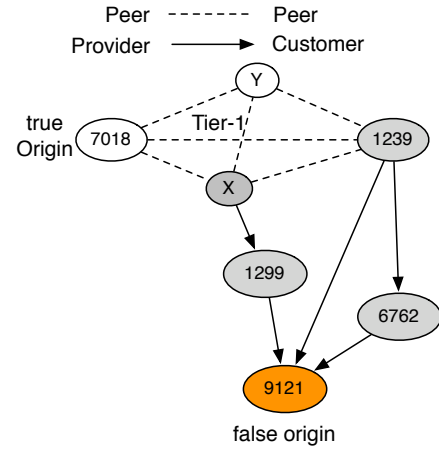
5.2.2 Hijacked Prefix belonging to Customer of Tier-1s

Next, we see how the AS-9121 hijack incident affected the prefixes belonging to an AS that was a direct customer of multiple tier-1 ASes. We picked AS-6461 as an example here because it connected to all the 8 tier-1 ASes. AS-6461 announced over 100 prefixes, 87 of which were hijacked by AS-9121. No more than 2 monitors were deceived by the false origin of all the hijacked prefixes. Figure 8(b) shows the entities involved in the hijack of one of the prefixes belonging to AS-6461. As before, AS-6762 believed the false origin and was one of the monitors deceived of all the hijacked prefixes of AS-6461. However, because all the tier-1 ASes were direct providers of AS-6461, they stayed with the original one-hop customer route to the true origin; in particular, note that AS-1239 was a provider for both the true origin and the hijacker, and it stayed with the original correct route. As a result, the hijack of AS-6461's prefixes made a very low impact.

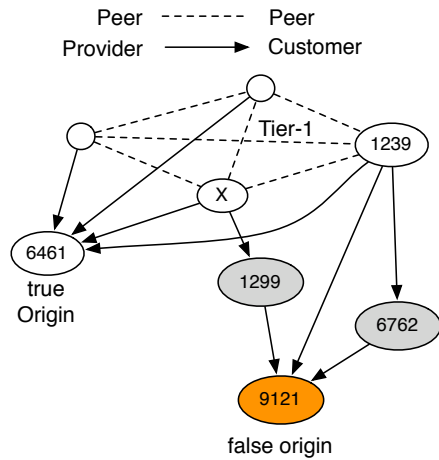
In addition to AS-6461, we also studied the impacts of prefixes belonging to a few other transit ASes that were very well connected to tier-1 ASes, and found the impact pattern for their prefixes to be very similar to the AS-6461 case. To summarize, this real life hijack event showed strong evidence that direct multi-homing to all or most tier-1 ASes can greatly increase an AS's resiliency against prefix hijacks.

6 Discussion

It has been long recognized that prefix hijacking can be a serious security threat to the Internet. Several hijack *prevention* solutions have been proposed, such as SBGP [7], so-BGP [11], and more recently the effort in the IETF Se-



(a) Tier-1 prefix hijacked: Tier 1 providers like 1239 and X, preferred the customer route to false origin 9121, instead of peer route to the true origin 7018, also a tier-1.



(b) Multi-homed customer of tier-1s hijacked: Providers of false origin 9121 got deceived, but all tier-1s including 1239, stayed with the one hop customer route to true origin 6461

Figure 8. Case studies with AS-9121 as false origin

cure Inter-Domain Routing Working Group [1]. These proposed solutions use cryptographic-based origin authentication mechanisms, which require coordinated efforts among a large number of organizations and thus will take time to get deployed. Meanwhile prefix hijack incidents occur from time to time and our work provides an assessment of the potential impacts of these incidents. Several hijack *detection* systems have also been developed, for example MyASN[10] and PHAS[8]. However since these systems are reactive in nature, it is still important for network customers to understand the relations between their networks' topological con-

nectivity and the potential vulnerability in face of prefix hijacks.

Our simulation and analysis show that AS nodes with large node-degrees (e.g., tier-1 networks) are not the most resilient against hijacks of their own prefixes. An AS can gain high resiliency against prefix hijacks by being direct or indirect customers of multiple tier-1 providers with the shortest possible AS paths. Conversely, such customer AS nodes can also make the most impact over the entire Internet, if they inject false routes into the Internet. This finding suggests that securing the routing announcements from the major ISPs alone is not effective in curbing a high impact attack, and that it is even more important to watch the announcements from lower-tier networks with good topological connectivity.

On the other hand, customer networks that are far away from their indirect tier-1 providers can be greatly affected if their prefixes get hijacked. These topologically disadvantaged AS nodes are in the most need for investigating other means to protect themselves. Subscribing to prefix hijack detection systems, such as MyASN and PHAS, would be helpful. To reduce the transient impact during the detection delay, one may also look into another proposed solution called PGBGP [5], which is briefly described in Section 7.

Note that the topological connectivity required for resiliency against prefix hijacks is different from that required for fast routing convergence [12]. Fast convergence benefits from fewer alternative paths when the routes change, thus prefixes announced by tier-1 providers meet the requirement well; while hijack resiliency benefits from being a direct or indirect customer of a large number of tier-1 providers, thus prefixes are better hosted by well connected non-tier-1 AS nodes.

We would like to end this discussion by stressing the importance of understanding prefix hijack impacts, even when the protection mechanisms are put in place. Our evaluations on an Internet scale topology in Section 4 used a *no-valley prefer customer* routing policy and showed that tier-1 AS nodes are not very resilient to hijacks of their own prefixes since other tier-1 AS nodes prefer customer routes to false origin. However, in reality a tier-1 AS may use various mechanisms, such as Internet Routing Registries (IRR), to check the origin of a prefix before forwarding the route. Such mechanisms would probably boost the resiliency of tier-1 AS nodes being hijacked. On the other hand, these protection mechanisms can also fail or backfire, thus exposing the vulnerability of a network. As we saw in case I of Section 5, most of the hijacked prefixes were the former customers of the false origin AS and were recorded in the Internet Routing Registry (IRR), which was not updated. Outdated registries resulted in false routes being propagated to the rest of the Internet.

Another example of a protection mechanism is the maximum prefix filter in BGP that allows an AS to configure

the maximum number of routes received from a neighbor. Thus, by limiting the total number of routes received from a neighbor, an AS can limit the damage in case of the neighbor announcing false routes. In case II from Section 5, AS-9121 announced over 100,000 false routes and one of its neighbors, AS-1299, had a max prefix set to a relatively low value. AS-1299 believed only 1849 routes directly from AS-9121, but since the max prefix limit is per neighbor, AS-1299 received hijacked routes from other neighbors as well. It learned a total of over 100,000 bad routes from all the neighbors combined, thus infecting a major portion of its routing table [?]. These examples show how easily protection mechanisms can fail due to human errors, underlining the need to understand the impact of hijacks in face of protection failures, and the need to protect networks by multiple means such as PGBGP and PHAS.

7 Related Work

Previous efforts on prefix hijacking can be broadly sorted into two categories: *hijack prevention* and *hijack detection*. Generally speaking, prefix hijack prevention solutions are based on cryptographic authentications [17, 11, 7, 9, 16] where BGP routers sign and verify the origin AS and AS path of each prefix. In addition to added router workload, these solutions require changes to *all* router implementations, and some of them also require a public key infrastructure. Due to these obstacles, none of the proposed prevention schemes is expected to see deployment in near future.

A number of prefix hijack detection schemes have been developed recently [10, 8, 13, 5]. A commonality among these solutions is that they do not use cryptographic-based mechanisms. In [13], any suspicious route announcements received by an AS trigger verification probes to other AS nodes and the results are reported to the true origin. In PGBGP [5], each router monitors the origin AS nodes in BGP announcements for each prefix over time; any newly occurred origin AS of a prefix is considered anomalous, and the router avoids using anomalous routes if the previously existing route to the same prefix is still available. Different from the above en route detection schemes, MyASN[10] is an offline prefix hijack alert service provided by RIPE. A prefix owner registers the valid origin set for a prefix, and MyASN sends an alarm via regular email when any invalid origin AS is observed in BGP routing updates. PHAS [8] is also an off-path prefix hijack detection system which uses BGP routing data collected by RouteViews and RIPE. Instead of asking prefix owners to register valid origin AS sets as is done by MyASN, PHAS keeps track of the origin AS set for each announced prefix, and sends hijack alerts via multiple path email delivery to the true origin.

Unlike the prevention schemes, a hijack detection mechanism provides only half of the solution: after a prefix hijack is detected, correction steps must follow. A recent proposal

called MIRO [18] gives end users the ability to perform correction after detecting a problem. MIRO is a new inter-domain routing architecture that utilizes multiple path routing. In MIRO, AS nodes can negotiate alternative routes to reach a given destination, potentially bypassing nodes affected by hijack attacks.

The work presented in this paper can be considered orthogonal to all the existing efforts in the area. It examines the relation between an AS node's topological connectivity and its resiliency against false route attacks, or conversely, an AS node's topological connectivity and its impact as a launching pad for prefix hijacks.

8 Conclusion

In this paper we conducted the first investigation into the relation between networks' topological connectivity and the impact of prefix hijacking. Our results show that, AS nodes that are close customers of multiple tier-1 providers are most resilient against hijacks of their own prefixes. Conversely, they can also be the most effective prefix hijackers of other's prefixes.

To gain topological resiliency, our results lead to the following recommendations to customer networks. First, one should try to multi-home directly with as many tier-1 providers as feasible, and choose non-tier-1 providers in a way to maximize the number of tier-1 providers reached through provider-customer AS links. Second, those topologically disadvantaged AS nodes should seek additional means to enhance their resiliency against potential hijack attacks, such as hijack detection services (PHAS) in combination with delayed adoption of suspicious new routes (PG-BGP). Third, operators of those most influential AS nodes should be especially vigilant against faults, and operators of tier-1 providers should pay special attention to routing announcements from those well connected customers.

In departing we note that our results indicate that the AS nodes with highest node-degrees (e.g. tier-1 ISPs) are not the most effective hijack launch pads nor are they the most resilient against prefix hijacks. In hindsight, our results are not surprising as they follow directly from the common routing policy of preferring customer routes over peer and provider routes. Nevertheless, they are counter-intuitive and become obvious only *afterwards*, and we believe it is important to make the results widely disseminated to both network operators as well as the research community.

References

- [1] Secure Inter-Domain Routing (SIDR) Working Group. <http://www1.ietf.org/html.charters/sidr-charter.html>.
- [2] S. S. E. Chen. Avoid BGP Best Path Transitions from One External to Another. Internet Draft, IETF,

- June 2006. <http://www.ietf.org/internet-drafts/draft-ietf-idr-avoid-transition-04.txt>.
- [3] L. Gao. On inferring autonomous system relationships in the Internet. *ACM/IEEE Transactions on Networking*, 9(6):733–745, 2001.
- [4] B. Halabi and D. McPherson. *Internet Routing Architectures*. Cisco Press, 2nd edition, 2000.
- [5] J. Karlin, S. Forrest, and J. Rexford. Pretty good bgp: Protecting bgp by cautiously selecting routes. Technical Report TR-CS-2005-37, University of New Mexico, October 2005.
- [6] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol. *IEEE Journal of Selected Areas in Communications*, 18(4), 2000.
- [7] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *IEEE JSAC Special Issue on Network Security*, 2000.
- [8] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang. PHAS: A prefix hijack alert system. In *15th USENIX Security Symposium*, 2006.
- [9] S. S. M. Zhao and D. Nicol. Aggregated path authentication for efficient bgp security. In *12th ACM Conference on Computer and Communications Security (CCS)*, November 2005.
- [10] myasn system. <http://www.ris.ripe.net/myasn.html>.
- [11] J. Ng. Extensions to BGP to Support Secure Origin BGP. <ftp://ftp-eng.cisco.com/sobgp/drafts/draft-ng-sobgp-bgp-extensions-02.txt>, April 2004.
- [12] R. Oliveira, B. Zhang, D. Pei, R. Izhak-Ratzin, and L. Zhang. Quantifying Path Exploration in the Internet. In *ACM SIGCOMM/USENIX Internet Measurement Conference (IMC)*, October 2006.
- [13] S. Qiu, F. Monrose, A. Terzis, and P. McDaniel. Efficient techniques for detecting false origin advertisements in interdomain routing. In *Second workshop on Secure Network Protocols (NPsec)*, 2006.
- [14] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *Proceedings of ACM SIGCOMM*, 2006.
- [15] Y. Rekhter, T. Li, and S. Hares. Border Gateway Protocol 4. RFC 4271, Internet Engineering Task Force, January 2006.
- [16] B. R. Smith, S. Murphy, and J. J. Garcia-Luna-Aceves. Securing the border gateway routing protocol. In *Global Internet '96*, November 1996.
- [17] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz. Listen and whisper: Security mechanisms for bgp. In *Proceedings of ACM NDSI 2004*, March 2004.
- [18] W. Xu and J. Rexford. Miro: multi-path interdomain routing. In *SIGCOMM*, pages 171–182, 2006.
- [19] B. Zhang, R. Liu, D. Massey, and L. Zhang. Internet Topology Project. <http://irl.cs.ucla.edu/topology/>.
- [20] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the internet's as level topology. In *ACM Sigcomm Computer Communication Review*, 2005.