Simulating the National-Level Impact of Routing Attacks in Sweden

Michael Liljenstam
mili@omicron.se
Omicron Ceti AB
Borgarfjordsgatan 7, 164 40 Kista

Abstract

This paper describes a simulation study of potential consequences to Swedish Internet users from propagation of false routing information, conducted as part of a larger study of potential vulnerabilities in the interdomain routing system by the Swedish National Post&Telecom Agency. Based on the route filtering information available in the routing registry, preliminary results for route deaggregations indicate that: i) the point where the attack is initiated determines the impact to a greater extent than which organization is targeted; and ii) due to defensive filtering and market concentration, attacks from most points in the network have limited impact. However, in some cases the results indicate that at least 60-70% of the users could be affected. As one might expect, the largest players have a key security role for the system as a whole. More interestingly, the model also suggests a few specific cases where attacks initiated from players with small market share could potentially affect larger sections of users within other ISPs.

1. Introduction

The interdomain routing system is a critical function in the Internet. Unfortunately, it was designed without fundamental security features and thus has known vulnerabilities [1, 2]. Consequently, there is currently a lot of activity devoted to developing improvements to routing security (including [5, 12]). The study reported here is part of a larger study by the Swedish National Post&Telecom Agency (PTS), where the aim is to perform a risk analysis of certain types of attacks by attempting to quantify their potential impact for Swedish Internet users. In this paper we consider injection of false routing information into the system, as might occur by accident or through a deliberate attack. A detailed simulation model of the Swedish part of the interdomain routing system at the AS-level has been constructed using data from public and non-public sources. We describe validation of the model, and preliminary results from systematic experiments where attacks are initiated from different parts of the network targeting a select set of organizations.

There is increasing interest in the use of simulation to study security aspects of the routing system [7]. However, we are not aware of any other studies that have attempted to build a model of the real system with the detail that is done here; including policy information, user market shares, and other data, to quantify the consequences of security events.

The starting point for building the simulation model was the SSFNet packet-level network simulator [8], which has previously been used in other studies concerning BGP [4, 6]. Section 2 describes the construction of the simulation model; followed by validation of the model in Section 3. Section 4 describes preliminary results from simulated attacks, and finally, conclusions and future directions are discussed in Section 5.

2. Model Construction

A model of the Swedish part of the interdomain routing system, primarily at the AS-level, was constructed using data from several sources:

- **Public BGP data.** BGP table dumps (RIBs) collected by Route Views [10] and RIPE NCC [9] were used to derive AS adjacencies.
- **A list of registered Swedish ISPs** kept by the PTS was used to select ASes for the model to create a focus on Swedish organizations and users.
- **ISP market share data** collected by the PTS was used to assign fractions of Swedish users to different ASes in the model.
- **Routing policy data from the Internet Routing Registry** is the only publicly avail-

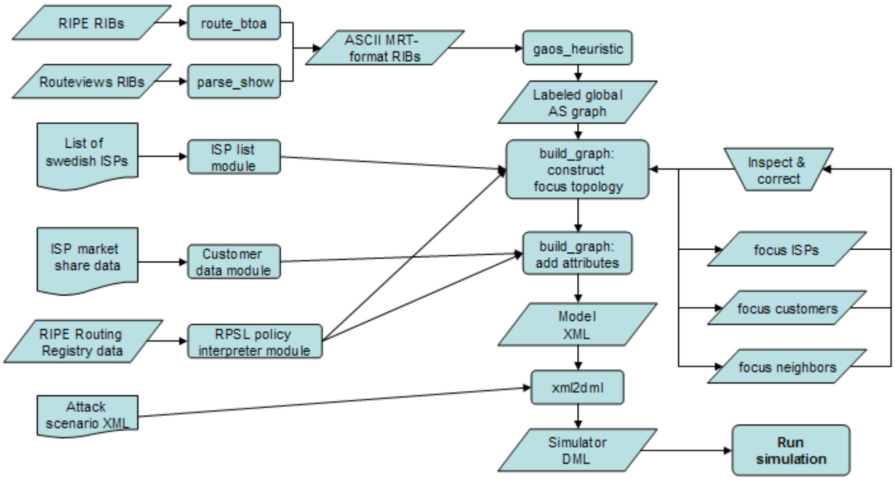

Figure 1. Model construction process.

able information on routing policy, and was used to supply routing policies for the model.

The complete process to create the model is fairly complex, as shown in Figure 1. AS path information from BGP RIBs was used to derive AS adjacencies and to infer peering relationships using Gao's heuristic [3]. In order to get good coverage of the Swedish part of the routing system, the Route Views RIBs containing 8.6 million routes collected globally were supplemented with 1.86 million routes collected by the RIPE route collector at the Stockholm Internet exchange point (run by Netnod) and the London Internet exchange. From these data sets, a global AS-topology is created, annotated with peering relationships (provider/customer, peer/peer, or sibling/sibling).

The resulting global AS-topology is unwieldy to simulate and we are interested primarily in parts related to Swedish interests. Hence, a subgraph focused on Sweden, was extracted using a simple iterative heuristic method that iteratively grows the focus topology starting from the ASes NETNOD and TELIANET-SE subject to manual inspection steps. The resulting topology consists of 167 ASes, and is shown in Figure 2. Since private peering connections, that are not globally visible, and unused backup connections are generally missing from collected routing tables, we also add adjacencies indicated by policy information from the Internet Routing Registry[1] The network inside each AS is abstracted away, and a single BGP speaking router represents the AS.[2] The internal topology of the ASes is abstracted away, and represented by a single router.

Once the topology has been created, additional attributes can be added to the ASes. The collected ISP market share data consists of reported number of subscribers, both individuals and organizations, broken down into counts of how many connections there are of certain different categories (dial up or a few different categories of broad band connections). For the model we count customers, not distinguishing between individual customers, organizations, or different types of connections.[3] From the data we identify the 30 ISPs with the largest market shares and assign fractions of the total market sum of subscriptions to ASes. This way, 94% of the Swedish users, i.e. total sum of subscription counts, is covered by the model.

A critical part of the model is the configuration of routing policies. Unfortunately, the detailed policies are generally regarded as sensitive business information, and thus ISPs are reluctant to reveal them. The only publicly available information on routing policy is what has been registered into the Internet Rout-

[1] Only adjacencies where policies occur at both end point ASes were included, to try to avoid including stale information.

[2] This is presumably not very far from reality for some of the customer ASes, that may only have a small number of BGP speaking routers. For larger ISPs, however, this is clearly a radical simplification. To study attacks based on DoS traffic against single routers, this kind of simplification would be unacceptable. However, to study the propagation of false routing information, the simplification is not expected to be a problem.

[3] Hence, a large organization of a thousand people is equated with a single home user. Moreover, because it is not uncommon that home users keep a dial up subscription as a backup when switching to a broadband connection, and the data does not allow us to identify these cases, there may be some degree of double counting of subscriptions.

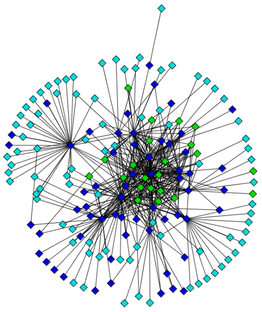

Figure 2. AS-topology. (Swedish ISPs = dark/blue, Customers = light/light blue, Foreign = middle/green.

ing Registry for coordination of policies between different ASes. “Mid-sized” ISPs (i.e., larger ISPs, but not the largest tier-1 ISPs) tend to use this information to generate BGP filters automatically from the policies expressed in Routing Policy Specification Language (RPSL) in the registry. We use parser code from the `RtConfig` tool (maintained by RIPE) to parse the RPSL objects in the registry, and then construct prefix filters for the simulation model similarly to what `RtConfig` can generate for real routers. However, because the information in the routing registry is entered manually, on a voluntary basis, there are known problems with missing and/or stale information. Hence, basing the model on the routing registry entails certain limitations. On the other hand, other studies have indicated that the routing registry maintained by RIPE NCC (Europe) is the best maintained part [11]. In cases where policy information is missing, or for parts of the global topology that have been abstracted away, a conservative assumption was made such that filters were assumed to block illegitimate routes from propagating. The intention is to find a *lower bound* for the consequences rather than overestimating them. We focus on traffic from Swedish users contacting Swedish organizations and in the model thus announce only prefixes announced by ASes belonging to Swedish ISPs and customer ASes. The last step in the process is to configure an attack scenario into the network model and convert it into the Domain Modeling Language used by SSFNet configuration files.

3. Model Validation

The approach chosen to attempt to validate the model against reality was to compare routes collected at certain points in the system. Routes collected at RIPE’s collection point in Stockholm (rrc07) and routes observed at a Looking Glass server in AS1299 (called lg1299 here) are used for this comparison. Routes are collected in the same manner in the model to compare against the real routes.

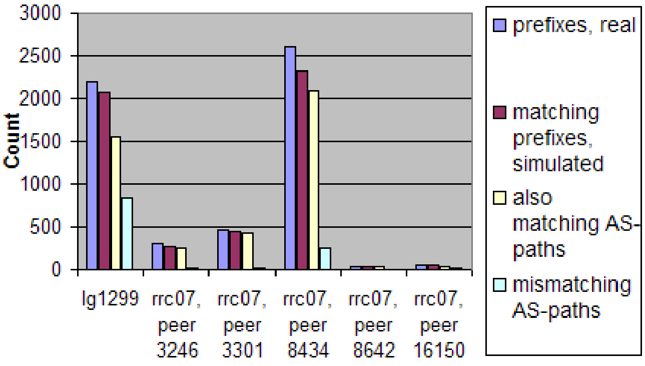

Figure 3. Validation of routes in model.

Initial validation result indicated problems (verified manually) with the policy data from the routing registry. At lg1299 and at rrc07 peer 8434 the real routing tables contain most of the prefixes announced by Swedish ASes, but in the model only about half of the prefixes reached the measurement points (data not shown). The only avenue open to improve on this is to attempt to infer some policy information from the routing data, where we can deduce paths from origins where the prefixes are permitted. This information is added to the model by making small corrections to the model policies to permit these routes, and the results are shown in Figure 3. The graph also shows, for all matching prefixes, the number of received routes where also the AS path matched or did not match, respectively. All in all, the results indicate reasonable agreement, keeping in mind that some issues with the policy data are evident.

4. Simulated Attacks

Attacks based on prefix deaggregations were simulated against several different organizations (one at a time) representing a few different sectors. Results for attacks against a bank is shown in Figure 4. A deaggregation attack against the bank could lead to the bank’s customers being unable to reach their online accounts or, more seriously, being redirected to a forged site (cf., phishing or pharming). Simulations were carried out for a large number of scenarios where the point (AS) where the attack was initiated from was varied.

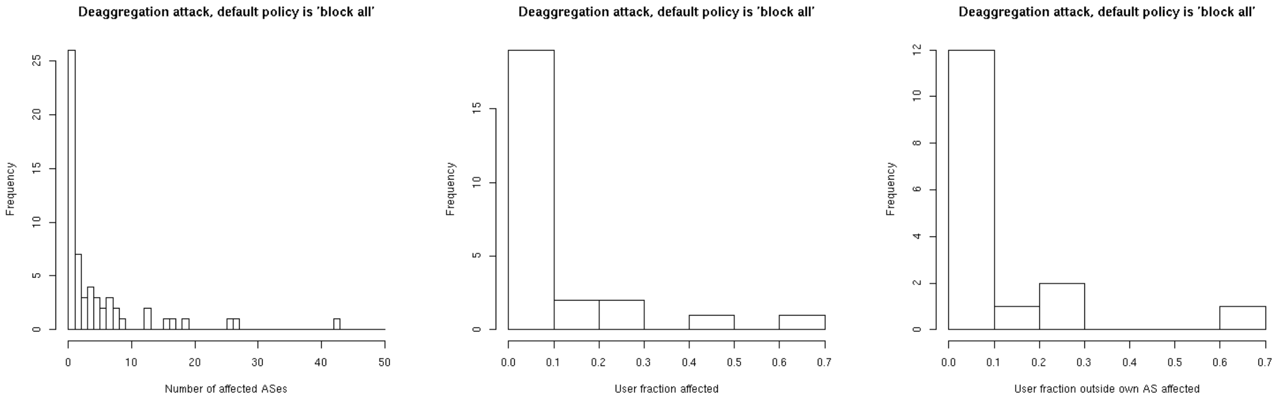

Figure 4. Histograms of consequences from deaggregation attack against a Swedish bank initiated from different ASes. Each graph plots the number of attack points (ASes belonging to Swedish ISPs) in the network that result in a certain impact on the system. Left: Number of ASes affected other than the AS where the attack was initiated; Middle: Fraction of users affected; Right: Fraction of users excluding initiating AS.

The preliminary results presented here are for attacks initiated from ASes belonging to the Swedish ISPs. The results indicate that in about half the cases, no other ASes are affected. In most other cases, the influence is limited to less than 10 ASes. Similarly, the fraction of users being affect is generally also limited to less than 10%. However, in some cases significantly larger influence is observed, up to 60-70% of the users, or more than 40 ASes; and given our assumptions the estimates should be considered lower bounds rather than absolutes. Varying the targets it became apparent that the results are very similar for different targets, but the point the attack is launched from is what determines the extent of impact on users.

5. Conclusions and Future Directions

The preliminary results indicate that filtering policies prevent deaggregation attacks in about half of the simulated scenarios, and the impact in other cases varies significantly. Most scenarios see limited impact, as market concentration requires some specific ASes to be affected for the attack to penetrate to up to about 60-70% of the users. However, the model also indicated a few specific cases of interest in that players with a small market share on their own could potentially affect ISPs with a larger market share. Thus, targeting the routing system might still extend the reach of certain network attacks. The simulation model might also be used to experiment with different proposed attack detection mechanisms that have been proposed, to evaluate their effectiveness in reducing the impact on users.

References

[1] K. Butler, T. Farley, P. McDaniel, and J. Rexford. A survey of bgp security. Draft report, www.patrickmcdaniel.org/pubs/td-5ugj33.pdf, 2005.

[2] Convery and Franz. Bgp vulnerability testing: Separating fact from fud v. 1.1. NANOG presentation, www.nanog.org, 2003.

[3] L. Gao. On Inferring Automonous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, Dec 2001.

[4] T. Griffin and B. J. Premore. An Experimental Analysis of BGP Convergence Time. In *Proceedings of the 9th International Conference on Network Protocols (ICNP)*, 2001.

[5] Kent, Lynn, and Seo. Secure border gateway protocol (secure-bgp). *IEEE Journal on Selected Areas in Communications*, 18(4), April 2000.

[6] Mao, Govindan, Varghese, and Katz. Route flap dampening exacerbates internet routing convergence. *ACM SIGCOMM Computer Communication Review, Vol.32, Issue 4, Proceedings SIGCOMM'02*, 2002.

[7] P. McDaniel. Iseb: Trace Driven Modeling of Internet-Scale BGP Attacks and Countermeasures. DETER/EMIST Workshop Presentation, Sept. 2005.

[8] SSFNet Project. http://www.ssfnet.org/, 2004.

[9] RIPE. RIPE NCC. http://www.ripe.net/info/ncc/index.html, 2006.

[10] RouteViews. University of Oregon Route Views Project. http://www.routeviews.org/, 2005.

[11] G. Siganos and M. Faloutsos. Analyzing BGP Policies: Methodology and Tool. In *INFOCOM*, 2004.

[12] White. Securing bgp through secure origin bgp. *The Internet Protocol Journal*, 6(4), Sept. 2003.