

Understanding the Challenges in Securing Internet Routing

Ricardo Oliveira Mohit Lad Lixia Zhang
UCLA Nokia UCLA
rveloso@cs.ucla.edu mohit.lad@nokia.com lixia@cs.ucla.edu

Abstract

The Internet routing system plays an essential role of glueing together tens of thousands of individual networks to create a global data delivery substrate. Over the years many efforts have been devoted to securing the routing system and a plethora of solutions have been proposed. Yet none of the solutions has seen wide adoption in the operational Internet and the routing system security remains a serious concern. In this paper we articulate the fundamental challenges in rolling out new security solutions to the global routing system by categorizing the various proposed solutions into a few classes and identifying the difficulties and remaining issues in deploying each class of solutions. Our examination of the solution space shows that monitoring is an essential component in securing the routing system, and that the “detect and react” class of solutions have the lowest hurdle in deployment and thus are most readily acceptable by the network operational community.

I. Introduction

As the Internet penetrates into every corner of the human society ranging from daily life information search to transactions in financial sector to management of critical infrastructure such as power supply systems, securing the global routing system also becomes of paramount importance. All applications of the Internet depend on the reliable functioning of the routing system to deliver their data to the right destinations. A routing system failure can lead to the failure of all applications, and a routing fault can result in denial of services to applications, or even compromises of applications security. As Huston commented in a recent talk [6], the most effective attack against Internet could be to target its routing system.

The importance of securing the global routing system has long been recognized. Over the years many research efforts have been devoted into this problem and have produced a number of specific solutions, for instance [19], [2], [5], [10], [21], just to name a few. Unfortunately as of today, none of the proposed solutions has been widely deployed, or even considered for deployment, in the operational Internet. The continued growth of the Internet both in size and in its complexity also makes it increasingly challenging to deploy new cryptographic protections. Naturally the following questions

come to one’s mind: Where are we today in securing the global routing system? What are the fundamental challenges in deploying secure routing solutions?

This paper aims to assess the current stage of affairs in securing the Internet routing system, with a focus on BGP security. Our objective is to examine different approaches to securing the global routing system, and to categorize them into a few classes. Rather than analyzing each and every security proposal, we aim to extract the basic approaches behind each solution class and identify its pros and cons. In particular, we articulate the deployment challenges faced by each class of solutions. We conclude the paper by identifying remaining challenges as well as promising directions for future efforts.

II. Current Best Practices in Securing Routing

In the operational Internet, routing system security practices can be broadly sorted into two categories: protecting BGP operations (router access and BGP session establishment), and protecting BGP routing information content.

A. Protecting Critical Elements

Today’s Internet routing operations are built on a mutual trust model between neighboring ASes. If AS X tells AS Z that it has a path X-Y to reach a destination network, AS Z does not attempt to verify the existence of the link X-Y. There is an implicit assumption that whatever one learns from a BGP neighbor is legitimate information. The philosophy of the current practice can be summarized as “Secure your own routers and hope that everybody else does the same”. The following key steps constitute the basic parts of the current routing security practice.

- 1) **Protecting Router Configurations:** Network operators (often manually) configure into each BGP router R a list of neighbor BGP routers that R is allowed to establish a BGP session with. Therefore protecting the configuration control of routers is an essential part of network security. The routing configuration access is usually protected through the use of SSH access to the routers with preconfigured cryptographic keys.
- 2) **Protecting BGP Sessions:** Even under the strict control over BGP session establishment, another potential security threat is false BGP update injection by a third party into established BGP sessions. The current

practice in protecting BGP sessions from intrusions is by using session passwords and MD5 checksums, and the access list to protect TCP port 179 (the port number used by BGP). One can also protect the router R from resource exhaustion by setting the maximum number of prefixes (max-prefix-limit) R may accept from each given neighbor BGP router. Another simple yet effective means of protecting the BGP session from DDoS attacks is the use of TTL hack [4].

- 3) **Keeping Configuration Updated:** Routers configurations change over time, and it is important to ensure that the configuration files reflects the current connectivity of the router. An outdated configuration may allow older BGP neighbors to connect when they should not.

Unfortunately no software system is bullet proof against break-ins, and human involvement in the above steps means that mistakes are bound to happen from time to time. Therefore active monitoring has been recognized as a necessary component in the process. For example in addition to use SSH for router access, one must also closely monitor any configuration changes.

All the above prevention measures are only for BGP connectivity and resource protection of the control plane. Even with a properly functional control plane, the content that flows through it can still be wrong, due to operational errors or malicious attacks. Thus one must also check the validity of routing data content to detect potential faults.

B. Protecting Critical Routing Content

The goal here is to make sure that all the routes one router learns from its BGP peers are authentic and correct. Since there has been no explicit binding between autonomous systems numbers (ASN) and the address prefixes each AS announces in BGP, AS y may announced a prefix owned by another AS x , resulting in traffic that was supposed to go to x being routed to y . This is called a *prefix hijack* attack. A recent prefix hijack attack happened to Youtube's prefixes, which were announced from an ISP in Pakistan [1]. If some user typed `www.youtube.com` into his browser while the attack was in force, that Web request would end up being directed to Pakistan instead of going to the Youtube servers.

To prevent prefix hijacks, one solution is for an ISP to validate the routing announcements from its customers by comparing them against a preconfigured customer prefix list. Some ISPs verify the route using a centralized routing registry called IRR[14] before accepting it; some others use a locally managed routing registry. ISPs may also adjust their neighbor eBGP route filters to accept route advertisements for each prefix. Other security mechanisms include using a bogon list to filter out false announcements of non-existing prefixes. However bogon prefixes may change rapidly over time, but automated and secure bogon list update mechanisms are yet to be developed.

C. The Need for More Protection

In today's Internet, securing one's own routers is necessary but not sufficient to secure the global routing system.

One badly configured router anywhere may cause false information to propagate globally, and such tragedy has happened time and again in the past. Every network is different and managed autonomously, even the best practices may not be rigorously followed. As a result, bad behavior emanating from one network (like spam and bots) may not be suppressed.

Adding operational security measures is not about being able to create and maintain absolute security. It is about a pragmatic approach to *effective* risk mitigation, and it represents a trade-off between cost, complexity, flexibility and outcomes. We argue that the first requirement to achieve effective fault mitigations is a solution that does not require significant changes to the internals of Internet service providers, so that the solution will have a better chance getting deployed. In many ways, the Internet is like human society consisting of diverse, independently operating elements, and some bad elements are bound to exist. Instead of striving to make every individual parts perfect, we need to monitor the overall activity to detect any irregular behavior and react promptly. Monitoring seems a common denominator approach that is needed in all security solutions.

III. Solution Space

The solution space for securing the global routing system can be broadly divided into three potentially overlapping areas: *prevention*, *mitigation*, and *detect-and-react*.

A. Prevention

Generally speaking, the proposed secure BGP solutions in this category are based on cryptographic authentications [15], [10], [13], [18], [7]. Unfortunately these solutions share three deployment obstacles: the absence of a global PKI infrastructure in today's Internet, the high computational overhead of verifying BGP update signatures, and the requirement of changing implementations of *all* routers to achieve effective protection. For example the design of Secure BGP (S-BGP) [10] requires two PKIs, one for address ownership attestation and the other for route announcement attestation. In S-BGP, the prefix owner has an asymmetric private key for each prefix, generated by a global trust entity. Each AS along a path uses the corresponding public key to verify the prefix origin authenticity. To ensure that the route cannot be tampered by malicious ASes, each AS also signs each received update with its own private key before further forwarding. Other cryptographic-based secure BGP work, e.g., SPV [13], KC-BGP [26], and path stability based improvement [8], focused on the performance improvement of the second phase generating/verifying route attestation. SoBGP reduces the route verification overhead by providing only the origin AS ownership authentication using *Authorization Certificate* assigned by some global PKI.

An effective fault prevention mechanism seems the ideal solution to securing the global routing system, as they could block attacks from ever getting into the system. S-BGP even once attracted serious attention from a few network operators. However the deployment obstacles, as mentioned

above, effectively blocked the road towards the actual deployment.

B. Mitigation

Instead of relying on cryptographic authentications to verify received routes, in [20] Wang *et al.* proposed to protect the routes to top level DNS servers by letting each router observe the origin ASes and AS paths of these routes over time, and deferring the adaption to any sudden route changes until the changes can be verified through other means. Because of the high degree of redundancy of DNS servers, their evaluation using BGP logs shows that the delay before adapting to legitimate routing changes has little impact on the overall system performance. This solution is also incrementally deployable as it introduces no routing protocol changes.

PGBGP [9] further generalized the above approach to all prefixes. In PGBGP, each router monitors the origin AS nodes in BGP announcements for each prefix over time; any newly occurred $\langle \text{origin AS, prefix} \rangle$ pair is quarantined for a certain time until the origin of the route is verified; the route is also assigned a low local-preference value which prevents it from being used as long as any alternative route is available (which is normally the case, due to the ever increasing topological connectivity density of the Internet).

Zhang *et al.* [23] proposes yet another scheme where a victim AS relies on a set of other ASes called lifesavers to mitigate ongoing hijack attacks by shortening the real routes through condensing the entire AS path into an AS_SET. Other ways to mitigate an attack is to announce more specific prefixes, as was done during the Youtube prefix hijack [1]. Of course this solution works only if the attacker is not already announcing the longest prefixes.

Anycast routing has also been used to mitigate, though not *prevent*, route hijacks. By announcing the same prefix from multiple locations, one can reduce the scope of route hijack impact. This approach has been used to protect the DNS root servers, although this approach does not scale as a general solution for all prefixes.

C. Detect-and-react

Recently a number of detection-based security solutions have been developed and even deployed; here we focus on those that at least have had the running code. Techniques that detect routing security attacks usually have two basic components: a *monitoring infrastructure* that collects BGP routing updates, and a *user profile* that provides the ground truth of the networks being observed. The user profile of a network X may contain information about the prefixes announced by X , the origin ASNs, the neighbor ASes of X , and other BGP specific information. This user profile can then be used to compare against observed BGP updates to detect potential faults as well as to filter the alerts that are false positives. Examples of such detection-based systems include PHAS [11], IAR [9], MyASN [17], and more recently Cyclops [3]. All these systems work by sending users alert messages whenever the monitors detect suspicious

BGP announcements, which could be due to prefix hijacks or misconfigurations.

As the successor of PHAS[11], Cyclops[3] is the latest comer in this category. It incorporated feedbacks from network operator community and provides ease of use, tailored alarm generations based on user configuration, extended dimension of monitoring infrastructure and fast response time by using real time BGP feed provided by BGPmon [22]. A major difference between Cyclops and previous approaches is the notion of *network configuration state*, *i.e.* Cyclops uses the network configuration from its users to filter out false positive and false negative alerts. Furthermore, Cyclops allows the users to easily adjust their configuration based on false alerts, as well as collects feedbacks from users to further reduce false positives.

One fundamental advantage of detection-based mechanisms is their ease of deployment: because the detection system is built along the side, instead of within, the existing routing system, they are readily deployable in any existing networks to effectively detect routing faults right away. They do not require any change to the operational system, and do not rely on any individual ISPs participation to work. These advantages of detection-based solutions are in sharp contrast with cryptographic-based prevention solutions, whose deployment requires changes to the entire routing system. In addition, as errors and failures are inevitable in any systems, cryptographic-based prevention solutions, even if they were rolled out, would also require a monitoring system to observe its operation and detect failures.

Of course every coin has two sides, detect-and-react solutions also face their own challenges. First, the monitoring system can become the target of attacks; how best to secure a monitoring system deserved further research. Moreover, detections alone do not prevent damages. The current detection systems only provide input to network operators who must then manually react to the faults or attacks. In the next section we gauge the solution space along the monitoring direction and identify remaining challenges.

IV. Examining Monitoring Based Solutions

We identify three major strengths of monitoring based approaches to routing security. First, they do not require changes to the operational routing system; one simply establishes BGP sessions with operational routers, called *monitors*, to passively connect routing data in real time. As such, one can collect data from a large number of monitors. Second, this class of solutions are not affected by how well individual ISPs may be managing their own networks, as failures or attacks can be detected from BGP updates from multiple routers collectively. Finally, the data collectors can combine inputs from monitors in multiple ASes, together with operations inputs, to form a global picture and identify the scope and magnitude of attack impacts. This class of solutions to routing security is readily deployable; several of them are in operation today. When combined with automated reactions, this class of solutions can also be highly effective in mitigating attack damages.

However we still face a number of challenges to effec-

tively protect the global routing system using monitor-based solutions. To effectively detect route hijacks and other faults, the first open issue is monitor placement. Due to the nature of routing protocols, major routing outages are easily detected with few monitors. However to be able to detect security problems occurring within a limited scope, the topological coverage of monitor locations becomes a critical issue. Lad *et al.* [12] shows some preliminary results; additional studies are needed to guide the monitor placement decisions.

One limitation in monitor placement is our knowledge of Internet topology. The current set of monitors provides a good coverage over top tier ISPs, however due to routing policies, a significant portion of existing AS links are invisible to this set of monitors, especially the peering links between stub ASs. Lack of the complete topology can lead to a bias both in the monitor placement decisions, and in gauging the scope of impact caused by an attack. An important study would be to understand the effect of the missing topological connectivity information on monitor placement, and to identify effective means to compensate the missing information so that we can adequately estimate the scope of impact by detected routing security problems.

We must also consider countermeasures by attackers to avoid being caught. A hijacker may want a low profile hijack to simply inject attack traffic into the network [16], or to attack routes to a specific destination within a specific region. One way to increase the detectability of such low impact attacks is to collect all incoming routes to BGP routers instead of the exported best paths only. The recent BMP proposal [8] will provide us the ability to capture all incoming routes to a router, which would enable us to learn a lot more about topological connectivity as well as detect a wider range of routing attacks. Because certain events are of local scope, there is a need to have geographically distributed monitors, *e.g.* having monitors in top ISPs in each country.

Furthermore, we need to automate reaction systems to effectively mitigate attacks reported by the monitoring system. Human operators should still be involved in a supervision role to observe how well the automated reactions can work, but need to be removed from the burden of manually invoking the reactions.

Additional challenges include how to scale the monitor data processing. RouteViews and RIPE collect a large amount of routing data everyday and the volume is increasing as more monitors being added. Although the cost of storage resources continues to drop, challenges remain on how to identify routing faults from this vast amount of data. Yet another challenge is how to protect and secure the monitoring system itself. First, as the set of monitors gets big, false information can be injected into the collected data set. Second, the data collecting and distributing servers can easily become targets of attack. We are actively exploring solutions to the above open issues to develop effective solutions based on distributed monitoring systems.

V. Conclusion

The security of the routing system remains an open challenge in today's Internet. In this paper we assess the

state of the art in proposed solutions. Our examination over the advantages and disadvantages of different classes of solutions suggests that *detect-and-react* type of solutions are most promising: they do not require any change in the protocol and they are the only solutions currently used by network operators to battle against attacks and faults in the BGP routing systems. Two major challenges in these type of solutions are to assure the integrity and coverage of the monitoring system, and to develop automated reaction systems to mitigate the damages. These are the focus of our ongoing efforts.

References

- [1] Youtube hijack. <http://www.ripe.net/news/study-youtube-hijacking.html>.
- [2] A. Cavalli and D. Vieira. Working around bgp: An improvement of bgp session maintenance. *Networking and Services, International conference on*, 0:41, 2006.
- [3] Y.-J. Chi, R. Oliveira, and L. Zhang. Cyclops: the as-level connectivity observatory. *SIGCOMM Comput. Commun. Rev.*, 2008. <http://cyclops.cs.ucla.edu>.
- [4] V. Gill, J. Heasley, and D. Meyer. The bgp ttl security hack (btsh). <http://www.ietf.org/internet-drafts/draft-gill-btsh-01.txt>, December 2002.
- [5] Y. Hu, A. Perrig, and M. Sirbu. SPV: Secure path vector routing for securing BGP. In *Proceedings of ACM SIGCOMM.*, August 2004.
- [6] G. Huston. An operational perspective on BGP Security. August 2005.
- [7] G. Huston and G. Michaelson. Validation of Route Origination in BGP using the Resource Certificate PKI. Internet Draft, IETF, 2008. <http://smakd.potaroo.net/ietf/all-ids/draft-ietf-sidr-roat-validation-00.txt>.
- [8] R. F. J. Scudder and S. Stuart. Bgp monitoring protocol. <http://tools.ietf.org/html/draft-ietf-grow-bmp-00>, 2008.
- [9] J. Karlin, S. Forrest, and J. Rexford. Pretty good bgp: Improving bgp by cautiously adopting routes. In *Proceedings of the 14th IEEE International Conference on Network Protocols*, 2006.
- [10] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *IEEE JSAC Special Issue on Network Security*, 2000.
- [11] 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.
- [12] M. Lad, R. Oliveira, B. Zhang, and L. Zhang. Understanding Resiliency of Internet Topology Against Prefix Hijack Attacks. In *IEEE/PIF Dependable Systems and Networks (DSN)*, 2007.
- [13] 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.
- [14] Merit. Internet Routing Registry. <http://www.irr.net/>.
- [15] 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.
- [16] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *Proceedings of ACM SIGCOMM*, 2006.
- [17] RIPE. Routing information service: myASn System. <http://www.ris.ripe.net/myasn.html>.
- [18] B. R. Smith, S. Murphy, and J. J. Garcia-Luna-Aceves. Securing the border gateway routing protocol. In *Global Internet*, 1996.
- [19] T. Wan, E. Kranakis, and P. C. Oorschot. Oorschot. pretty secure bgp (psbgp). In *In The 12th Annual Network and Distributed System Security Symposium (NDSS05)*, 2005.
- [20] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. Wu, and L. Zhang. Protecting BGP Routes to Top Level DNS Servers. In *Proceedings of the ICDCS*, 2003.
- [21] R. White. Architecture and deployment considerations for Secure Origin BGP (soBGP). 2006.
- [22] H. Yan, R. Oliveira, K. Burnett, D. Matthews, L. Zhang, and D. Massey. Bgpmon: A real-time, scalable, extensible monitoring system. In *Cybersecurity Applications and Technologies Conference for Homeland Security(CATCH)*, 2009.
- [23] Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao. Practical defenses against bgp prefix hijacking. In *CoNEXT '07: Proceedings of the 2007 ACM CoNEXT conference*, 2007.