Failure Isolation based Defense against Internet CXPST-like Attack

Hongjun Liu[1], Xiaofeng Hu[1], Dan Zhao[1] and Xicheng Lu[1]

[1]*National University of Defense Technology, Changsha, China*
Seeker_lhj@163.com, xfhu@nudt.edu.cn, danzhao.nudt@gmail.com, xclluu@163.com

Abstract

Attacking on inter-domain routing system degrades the availability and performance of Internet severely. It is challenge to defend against the extreme attacks which exhaust the resources of routers by generating a great number of update messages. In this paper, we propose two mechanisms to protect Internet from such attacks: to isolate attacks in local region, unnecessary updates are suppressed without affecting the correctness of routing; to break down the route flapping which repeatedly generates updates, the paths selected are validated to detour the attacked links, which diffuses the deliberately attacks to random attacks . Simulation shows our methods greatly decrease the number of updates under such attacks, and isolate the attacks in local region of network.

Keywords: *Failure isolation, CXPST attack, BGP, Path exploration*

1. Introduction

Internet is now the home to an increasing number of critical business applications. Some attacks aiming to the underlying infrastructure (such as the Border Gateway Protocol, short for BGP) of internet may result in significant harm to individuals or institutions. The well known coremelt [1] and CXPST [2] attacks can severely disrupt the routing system, which leads to network instability and losing of connectivity and data. We refer to these attacks as CXPST-like attacks, for the attacks share the common idea of exhausting the resource of routers by generating enormous update messages which results from terminating of BGP sessions.

To defend against the CXPST-like attacks, it is vital to reduce the number of updates. One ideal method is to prevent BGP session from being terminated, such as S-BGP [3] validates the data passed between ASes using public key certificates. IRV [4] validates each data item by directly querying the AS from whence it came. However, to preserve message integrity, these methods are costly both in terms of computation and storage. As it is impossible to ensure absolute security of BGP sessions, HLP [5] and SCION [6] propose new protocol to divide the Internet into a set of isolated regions, which prevent the change of topology and routing information from transmitting out of the region where the change occurs. LSRP [7] isolates the instability locally in shortest path routing through different propagation speed. Unluckily, these works provide new protocols which change greatly from BGP, so it is not easy to apply them into practice. Especially, they do not try to isolate attacks to the smallest scale.

In this paper, to defend against CXPST-like attacks, we propose two mechanisms to eliminate unnecessary update messages. Firstly, to isolate attacks in local region, some unnecessary updates are suppressed without affect the correctness of routing. Secondly, to break down the route flapping caused by repeatedly attacking the recovered links, the

selected paths are validated to detour the attacked links, which results in the attacks are diffused to avoid deliberately attacking one target.

2. Our Approach

The key characteristic feature of the CXPST-like attacks is that route flapping generates enormous update messages to exhaust the resource of routers. In this section, to eliminate the globally visibility, we suppress the unnecessary updates without affect the correctness of routing. To breakdown the route flapping caused by repeatedly session failure and reestablishing, the selected path are validated to detour the attacked links, which results in the attacks losing the deliberately selected targets. In this paper, we suppose a routing instability only contains the updates generated by single attack, which is feasible by applying the method in [8] to distinguish updates triggered by one event from others.

2.1. Suppressing Unnecessary Updates

As illustrated in [5], BGP has very poor fault isolation properties, small-scale local perturbations can be propagated globally across network. To avoid the global visibility, it is rational to insulate the effect of attacks in local region by only propagating necessary updates. As BGP advertises updates when topology and policy changes, it is vital to determine which ASes needs to know the changes, i.e. affirming whether the updates are necessary to be advertised. In this paper, we propose the judging criteria by considering the path length variation of different types of updates.

Theorem 1: Given the best path from AS v to prefix $d \in u$ is $best_d(v)$, and AS m belongs to the path, i.e. $m \in best_d(v), m \neq u \neq v$. If the best path of m to reach d changes from p_1 to p_2, and the length variation is $0 < |p_2| \leq |p_1|$, advertising the updates that contain p_2 does not change the partial path between m and v in $best_d(v)$, i.e. $best_d(v)^{m-v}$.

Proof: Before the update, the best path p_1 and $best_d(v)$ are selected among the available paths P_d^m and P_d^v in m and v respectively. If the update is a withdrawal, the available paths in m and v change to $P_d^m - p_1$ and $P_d^v - p$, where $p \in P_d^v, p_1 \subset p$. Thus the withdrawal do not change the local preference and policies in m and v. If the update is an advertisement, when the advertisement is induced by policy change, the available paths in m and v are still P_d^m and P_d^v respectively; when it is induced by link restoration, the available paths in m and v are $P_d^m \cup p_2$ and $P_d^v \cup p$ respectively, where p is the path to reach d from v that meets $p_2 \subset p$. In the former, the local preference and other policies in m and v are both not changed by the update. In the latter, all the added paths of $P_d^v \cup p - P_d^v$ are formed by concatenating partial path $\{p_i^{m-v} \mid p_i \in P_d^v, m \in p_i\}$ and p_2, for all the updates belong to the same instability under one event. So only the local preference in m are changed but not v.

In summary, the local preference and policies in v are not changed. If v selects $best_d(v)$ based on the highest local preference before the updates, it still selects the path $best_d(v)^{m-v}$ to reach m, for the unchanged highest local preference. Otherwise, v selects $best_d(v)$ based on shortest path length only if the local preferences to all neighbors

are the same. So $best_d(v)$ is the shortest path from v to d among P_d^v before the update. As p_2 is export-rule guided, and $0 < |p_2| \le |p_1|$, $best_d(v)^{m-v} + p_2$ is the shortest path after the update. As a result, the partial path of the best path from v to d between m and v is unchanged during the update, i.e. $best_d(v)^{m-v}$.□

Theorem 1 implies the following conclusion:

Corollary 1: When a new selected path is of no longer length than the replaced best path, suppressing the new selected path does not affect the topology and routing viewed from the ASes outside of the attacked region to the AS who suppressed the updates when reaching the same prefixes.

Taking figure 1 for example, the best path from AS v to d before path change is $p_3 + p_1$. When an attack event triggers the best path changing from p_1 to p_2 at AS m, and p_2 is updated to AS v, if $|p_2| \le |p_1|$, v will select $p_3 + p_2$ as the new best path, even if $p_5 + p_1$, $p_5 + p_2$ and p_4 are of the same path length with $p_3 + p_1$. Thus v will still pass through path p_3 to reach prefix d, which is of the same effect if the path p_2 is suppressed at AS m.

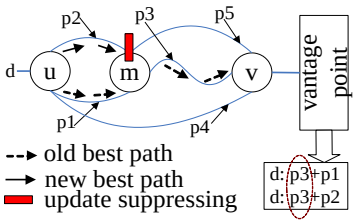

Figure 1. Example of Update Suppressing

According to theorem 1 and corollary 1, we advertise the updates containing paths whose length becomes longer, which leads to the selected path becoming longer as time goes on. It will result in the path can not resume to the shorter length. To avoid this situation, we advertise some path with shorter length only if the path is the preferred path [9], which is the path to one prefix that remains in the routing table for the longest time during a long time period. The underlying idea is that the internet routing will revert to the normal paths that are usually used for the longest time, for the policy rules manually configured are pretty much steady in the internet. Consequently, if there is a transient recovery with shorter paths which is not the preferred path, it is unnecessary to advertise this path, for the transient shorter paths will ultimately be replaced by the preferred path. In conclusion, the criteria of hiding unnecessary updates for isolating attacks in local region include:

C1: when the path length stays unchanged,
C2: when the path length become shorter, and the path is not the preferred path.

As stated in [9], policy change produces a large number of equal length path (around 70% of the total), so criterion C1 will largely reduce a great number of updates. Work in [10] analyzes empirical results of up events (such as link restoration, session reestablishing) in real BGP raw data dumps, and shows that 58.6% of the up events advertise the paths with same length, 7.9% advertise shorter paths 15.5% advertise longer paths, and 18% advertise paths with uncertain length change. With the help of

criteria of C1 and C2, the majority updates are suppressed. All of these illustrate the effectiveness of our method.

2.2. Broken Route Flapping

Since the essence of CXPST-like attacks is the engendering of route flapping which oscillates to generate a great number of updates, it is important to break down the route flapping in defending against these attacks. As illustrated in [2], the attacked sessions are deliberately selected. The repeated BGP connection tear-down and re-establishment belongs to the type of basic oscillation. The intuition to eliminate such type of oscillation resulting from CXPST-like attacks is to select path detouring the attacked links that are deliberately chosen, even if the attacked links are available again.

To eliminate the rout flapping resulting from CXPST-like attacks, it is chiefly to identify the occurrence of oscillation. As the deliberately selected links are repeatedly failed and restored, we only consider the stable paths before and after the attacks. As an event may change the paths of many prefixes, we check the stable paths of the nearest affected prefix d. Thus we save the stable paths of the nearest prefix before and after the instability which is induced by link failure and restoration, and the stable paths are denoted as p_1 and p_2. If the path set of p_1 and p_2 appears several times continuously, we claim that there is oscillation between path p_1 and p_2. To detour the attacked links, it is vital to infer the location of the links. By intersecting the two paths, we affirm the link belongs to p_1 which is the stable path before the instability triggered by the attacks, and the link is directly connected to the nearest prefix d, for the CXPST-like attack is well known to terminate BGP session, and d is the nearest prefix.

With the knowledge of attacked links, if one new selected path contains the attacked link, the path is suppressed. As a result, the new selected path will not pass through the attacked link again until the routing is stable. Consequently, our approach can greatly improve the stability of routing system, for Internet has high tolerance to random failures but deliberate attacks. Under random failures, Internet will not break until more than 95% of nodes have failed.

3. Evaluation

We choose SSFNET to generate events terminating and reestablishing BGP sessions, and the simulations are performed using internet like topology, which is generated as follows: An AS-level topology of the Internet mapped by CAIDA from April 29 2009 has 31212 ASes and 60052 links. For the great scale of the topology, we adopt the method of Dimitropoulosto to generate small scale internet-like policy annotated AS graph, which has 200 ASes and 359 links.

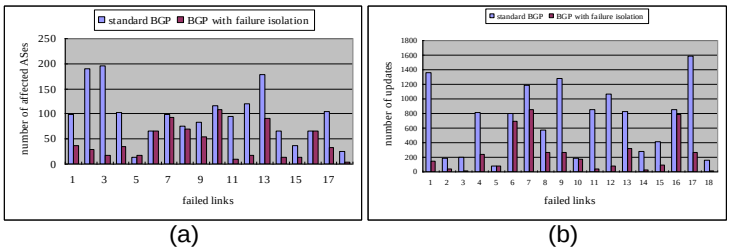

(a) (b)

Figure 2. Effect of Isolation on Link Tear-down: (a) Affect Scale (b) Number of Updates

To demonstrate the effectiveness, we compare our method with the standard BGP on the metrics of the number of update messages and the affected scale which is measured by the number of ASes observing the attacks. To simulate the CXPST-like attacks, we deliberately select 18 links which have the highest betweenness. When one link is failed or recovered, all the ASes in the topology are monitored. When link failed, our method can isolate the effect of link failures to a smaller scale, and largely decrease the number of updates, as shown in figure 2. When these links are recovered, the likewise effects still hold, which is shown in figure 3.

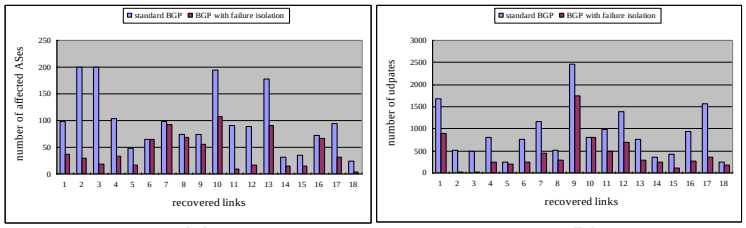

(a) (b)

Figure 3. Effect of Isolation on Link Recovery: (a) Affect Scale (b) Number of Updates

4. Conclusion

CXPST-like attacks are of great threat to internet, for they generate enormous update messages which exhaust the resources of routers. In the state of art, there is still no applicable method to effectively defend against this kind of attacks. In this paper, we propose two mechanisms to protect internet. To isolate the attacks in the local region, the unnecessary updates are suppressed. To break down the route flapping, the selected paths detour the attacked links. With these techniques, our method can insulate attacks locally, which is validated through simulation. In future work, we will suppress the unnecessary obsolete and stale paths in propagating process.

Acknowledgements

The work described in this paper is partially supported by the grants of the project of National Science Foundation of China under Grant No.61103189 and No. 61070199; and the Program for Science and Technology Innovative Research Team in Higher Educational Institutions of Hunan Province: "network technology",Changjiang Scholars and Innovative Research Team in University(No.IRT1012) and Hunan Province Natural Science Foundation of China (11JJ7003).

Preference

[1] A. Studer and A. Perrig, “The coremelt attack”, In Proceedings of the 14th European Symposium on Research in Computer Security (ESORICS 2009), **(2009)** Sept.

[2] M. Schuchard, E. Y. Vasserman, A. Mohaisen, D. F. Kune, N. Hopper and Y. Kim, “Losing Control of the Internet: Using the Data Plane to Attack the Control Plane”, In: Proc. of the NDSS, **(2010)**.

[3] S. Kent, C. Lynn and K. Seo, “Secure Border Gateway Protocol (S-BGP)”, IEEE Journal on Selected Areas in Communication Special Issue on Network Security, 18(4):582-592, **(2000)** April.

[4] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel and A. Rubin, “Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing”, In: Proc. of NDSS, **(2003)**.

[5] L. Subramanian, M. Caesar, C. T. Ee and M. Handley, "HLP: A next generation interdomain routing protocol", In: Proc. of SIGCOMM2005.Philadelphia, PA, USA: ACM Press, pp. 13-24. **(2005)**.

[6] X. Zhang, H. Hsiao, H. Geoffrey, H. Chan, P. Adrian and D. G. Andersen, "SCION: Scalability, Control, and Isolation on Next-Generation Networks", In Proc. of IEEE Symposium on Security and Privacy, pp. 212-227, **(2011)**.

[7] A. Arora and H. Zhang, "LSRP: local stabilization in shortest path routing", *IEEE/ACM Trans. Netw.*, 14(3):520–531, **(2006)** June.

[8] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger and B. Maggs, "Locating Internet routing instabilities", in Proc. ACM SIGCOMM, pages 205–218, Portland, OR, **(2004)** Aug.

[9] R. Oliveira, B. Zhang, D. Pei and L. Zhang, "Quantifying Path Exploration in the Internet", IEEE/ACM Transactions on Networking, vol. 17, no. 2, pp. 445-458, **(2009)** Apr.

[10] A. Bremler-Barr, N. Chen, J. Kangasharju, O. Mokryn and Y. Shavitt, "Bringing order to BGP: decreasing time and message complexity", Computer Networks, 53(12), pp. 2241-2256, **(2009)** August 13.

Authors

Hongjun Liu received the B.S. and M.S. degree in computer science from National University of Defense Technology, China, in 2005 and 2007 respectively. He is currently pursuing the Ph.D. degree in the School of Computer at the same university. His research interest lies in Internet measurement, network survivability and inter-domain routing.

Xiaofeng Hu received his Ph.D. degree in computer science from National University of Defense Technology, China, in 2004. He is currently an associated professor in the School of Computer at the same university. His research interest includes Internet architecture, routing protocol, and high performance router design.

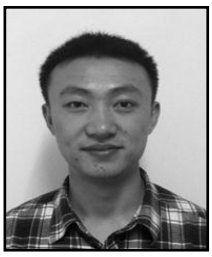

Dan Zhao is currently a Ph.D. student in the School of Computer of National University of Defense Technology in Changsha, China. He received the B.S. and M.S. degree in computer science in the same university, in 2006 and 2008 respectively. His research interest includes network architectures, Internet routing and protocols.

Xicheng Lu received his B.S. degree in computer science from Harbin Engineering Institute, Harbin, China, in 1970. He was a visiting scholar at the University of Massachusetts from 1982 to 1984. He is currently a professor with School of Computer Science of National University of Defense Technology (NUDT), Changsha, China. His research interests include distributed computing, computer networks, and parallel computing. He has served as a member of editorial boards of several journals and has co-chaired many professional conferences. He is an academician of the Chinese Academy of Engineering and a member of the IEEE.