How to mitigate the effect of scans on mapping systems

Damien Saucez
Université catholique de Louvain
Louvain-la-Neuve, Belgium
{first.last}@uclouvain.be

Luigi Iannone
Deutsche Telekom Laboratories
Berlin, Germany
luigi@net.t-labs.tu-berlin.de

1. INTRODUCTION

The network research community has recently started to work on the design of an alternate Internet Architecture aiming at solving some scalability issues that the current Internet is facing. The Locator/ID separation paradigm seems to well fit the requirements for this new Internet Architecture. The principle of this paradigm is to separate the identification part from the localization one. In today's Internet, nodes are identified by their IP address and the same IP address is used to localize the node in the Internet. In the Locator/ID separation proposals, locators are used to localize the nodes on the Internet (i.e., packets are routed towards the Locator) while the identification of the node is let independent of the routing infrastructure thanks to the ID. In this paper, among the various solutions, we consider LISP (Locator/ID Separation Protocol), proposed by Cisco [1].

LISP is a map-and-encap solution where the inner header addresses are identifiers and outer header addresses are locators. A set of locators is associate to each identifier via a mapping. Mappings are obtained by querying a mapping system ([2], [3], [4]) like in DNS where the DNS is queried to resolve a name.

Unfortunately, with on-demand mappings, some delay is observed between the reception of a packet by the border router and the time it can effectively be sent when the router sees for the first time the destination identifier. To reduce the impact of such mapping resolution delay, the *gleaning* functionality (discussed in Sec. 2) has been proposed in [1]. However, in it original form, the gleaning, introduces a cache poisoning threat. This paper presents an extended gleaning approach that limits the risk of cache poisoning and mitigates the impact of port scanning on the mapping system load.

2. EXTENDED GLEANING

Fig. 1(a) shows what happens when a LISP router has to send traffic to an identifier that it has never seen. When a packet arrives from the LISP site with a destination that is unknown (the red block), the LISP router drop the packet and sends a map request to get a mapping for this EID. The packets towards this EID are dropped until a map reply is received. Afterwards, a locator from the mapping (the yellow block) is used to transfer the packet over the Internet.

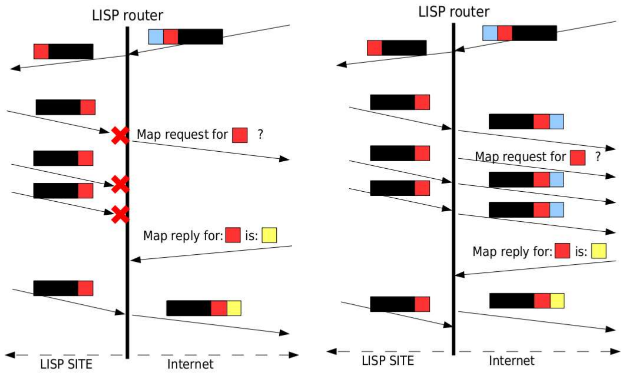

(a) Without gleaning

(b) With gleaning

Figure 1: Exchange of message when a LISP router has to transmit a packet towards an identifier for which it has no mapping.

Fortunately, when the traffic is bidirectional, it is possible to avoid to wait for a mapping. Fig. 1(b) shows how to do that. When a LISP packet arrives from the Internet and is originated from an identifier that is not known by the LISP router, the router installs a temporary mapping in its map cache. This mapping associates the identifier with the RLOC contained in the LISP packet. The LISP router then uses the temporary mapping for this identifier until the mapping has been confirmed by a map reply. Afterwards, the temporary mapping is removed and the confirmed mapping is used. This technique is called *gleaning* [1]. The gleaning avoids to have to wait for a mapping from the mapping system when a flow is initiated from the Internet. However, a gleaned mapping contains very few information and a map request is sent in parallel meaning that a map request is sent even in the case of scans. In addition, the gleaning is related to identifier, meaning that different flow can use the same gleaned entry, opening the doors of DoS and eavesdropping attacks. For example, an attacker can send a packet with the identifier of a target but using its own locator.

We propose an extended gleaning to tackles these problems. To minimize the risk of cache poisoning, the gleaned mapping is associated to a L4-flow instead of an identifier. It thus limit the scope of a gleaned mapping to the flow having generated it. Considering L4-flows independently also offers a way to mitigate the impact of scans on the mapping system load: the map request is sent only once the flow is validated (e.g., the 3-way handshake succeeds for TCP). However, as the inspection of packets is costly, this proposition can be relaxed by only counting the number of outgoing packet: a map request is sent only if at least two packets have been sent to the gleaned identifier.

3. EVALUATION

This section evaluates the extended gleaning and its impact on the cache size and the mapping system load.

A one day full NetFlow trace has been collected in our Campus (March 23rd 2009). Our campus is connected to the national research network with a 1 Gbps link. A total of 123, 805 different BGP prefixes are observed in the trace.

The benefits of the extended gleaning are computed thanks to a map cache simulator. The simulator maintains a map cache for outgoing flows and a gleaning cache populated by the incoming flows. The lifetime of an entry in the map cache is reset to 3 minute everytime it is hit. If it is not hit after 3 minutes, it is removed from the cache. Entries in the gleaning cache are stored for at most 3 minutes. The input of the simulator is the NetFlow trace. We assume that the mappings are decomposed following the BGP table of the day.

Fig. 2 shows the evolution of the number of entries in the map cache with and without the extended gleaning. It also shows the evolution of the number of entries to store in the gleaning cache with the time. We first see that the number of entries inserted in the map cache is lower when the extended gleaning is activated. This property shows that some entries do not generate more than two outgoing packet within 3 minutes. It is worth to notice that the cache size is more than 79% smaller than the number of observed BGP prefixes.

The peak observed from 8:30 am to 9:30 am localtime (7:30 to 8:30 GMT) corresponds to the start of the working hours in our campus and is mainly due to web surfing.

The extended gleaning helps to reduce the number of map requests and thus the load on the mapping system (shown by the difference between the two curves). A map request reduction of 15% is observed from 2, 717 requests per minute on average to 2, 300 with the extended gleaning. The gleaning cache has an average of 5, 000 entries and is rather stable.

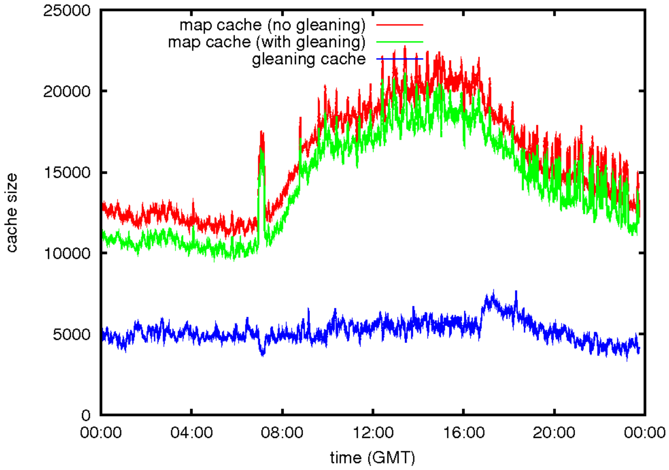

Figure 2: Evolution of the cache size with and without the extended gleaning and size of the gleaning cache.

4. CONCLUSION

The Locator/Identifier Separation Protocol (LISP) is an important proposal for the Future Internet. LISP relies on a separation between the endpoint identifiers and the locators used to send traffic to them. The separation has been proved to be a good property for the Future Internet [5], however, the separation means that identifiers and locators have to be mapped by some technique. We shown that the performance of such mapping system can be influenced by scans and proposed a solution to limit their impact. Securing the gleaning is important but is only a small part of the security problem in LISP. More precisely, extensive studies have still to be performed to secure the mapping systems.

5. REFERENCES

[1] D. Farinacci, V. Fuller, D. Oran, D. Meyer, and S. Brim, “Locator/ID Separation Protocol (LISP),” IETF Network Working Group, Internet Draft *draft-farinacci-lisp-10.txt*, November 2008.

[2] B. Carpenter, “General Identifier-Locator Mapping Considerations,” IETF Network Working Group, Internet Draft draft-carpenter-idloc-map-cons-01.txt, June 2007.

[3] D. Farinacci, V. Fuller, and D. Meyer, “LISP Alternative Topology (LISP-ALT),” IETF Network Working Group, Internet Draft *draft-fuller-lisp-alt-03.txt*, October 2008.

[4] L. Mathy and L. Iannone, “LISP-DHT: Towards a DHT to map identifiers onto locators,” in *Proc. of ReArch'08 - Re-Architecting the Internet.*, Dec. 2008.

[5] B. Quoitin, L. Iannone, C. de Launois, and O. Bonaventure, “Evaluating the Benefits of the Locator/Identifier Separation,” *Proc. 2nd ACM SIGCOMM Workshop on Mobility in the Evolving Internet Architecture (MobiArch)*, August 2007.