Mapping Systems for the Future Internet

[Extended Abstract]

Michael Hoefling
University of Tübingen, Chair of Communication Networks, Germany
hoefling@informatik.uni-tuebingen.de

ABSTRACT

The Internet routing research community has performed extensive research to solve impending problems in the current Internet. One of the most promising approaches for a new routing architecture is the separation of core routing and host identification. The address space of the Internet is divided into core and edge network addresses. To provide fast and reliable routing, a mapping system is needed to resolve the mapping between both address spaces. This extended abstract presents the scalable, resilient, and secure "Future InteRnet Mapping System" (FIRMS).

1. INTRODUCTION

The Internet routing research community has performed extensive research to solve impending problems in the current Internet. One of the most promising approaches for a new routing architecture is the separation of core routing and host identification. The idea behind the locator/identifier (Loc/ID) split is to divide the global address space into core network addresses, the routing locators (RLOCs), and edge network addresses, the endpoint identifiers (EIDs). EIDs are reachable in the local scope and are used for intra-domain routing. RLOCs are globally reachable addresses, specify the location of a given EID in the global Internet, and are used for inter-domain routing. Through Loc/ID, the global state of reachability information in the Internet core can be minimized to only necessary information. Communication in a Loc/ID Internet involves middleboxes which have to resolve the destination EID to the destination's RLOC. The mapping system used by the middleboxes has to be fast-responding, scalable, reliable, secure, and be able to relay initial packets.

The contribution of this extended abstract is the presentation of the "Future InteRnet Mapping System" (FIRMS), a mapping system which can be used with and by future Internet routing architectures. The remainder of the extended abstract is structured as follows. Section 2 briefly introduces LISP, a Loc/ID routing architecture. Section 3 describes the design of FIRMS in detail. Section 4 shortly summarizes related work in the area of mapping systems and Section 5 concludes the extended abstract.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 2011 ACM X-XXXXX-XX-X/XX/XX ...$10.00.

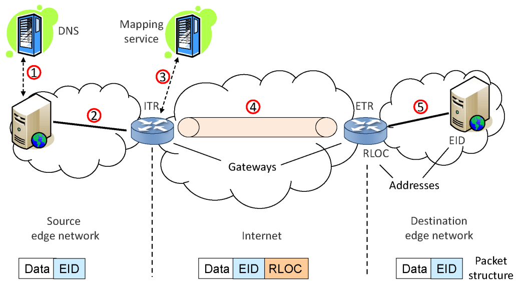

Figure 1: Locator/Identifier Separation Protocol.

2. BASICS OF LISP

The IETF currently develops a prototype implementation of a Loc/ID split architecture called the Locator/Identifier Separation Protocol (LISP)[1]. The basic architecture is shown in Figure 1. LISP-enabled networks are called LISP domains and are connected to the Internet via LISP gateways. Each host is identified by an EID while each LISP gateway is identified by a set of RLOCs. Communication between two hosts in different LISP domains requires tunneling. First, the source host queries the DNS for the EID of the destination host. Then, it forwards the packet to the LISP gateway which acts as an ingress tunnel router (ITR). The ITR queries the mapping system for the RLOC of the gateway that belongs to the LISP domain hosting the destination EID given in the packet. The ITR encapsulates the packet and sends it towards the destination LISP domain. The receiving LISP gateway acts as an egress tunnel router (ETR), i.e., it decapsulates the received packet and forwards it to the destination EID.

We make use of the LISP nomenclature (EID, RLOC, ITR, ETR) in the following, but FIRMS is also applicable to other routing approaches that are based on the Loc/ID split.

3. FIRMS

In this section, we present FIRMS, describe its architecture, specify its operation, and discuss its resilience and security features. We assume that EIDs are assigned in prefix-blocks from authorities to so-called prefix owners. Figure 2 illustrates the basic structure and operation of FIRMS.

3.1 Components

Mapping data is stored in so-called Map-Bases (MBs) which are the only authoritative source for mappings in FIRMS. MBs can be replicated to provide resiliency against MB failures. Primary MBs should be located near the destination site to enable forwarding of initial data packets.

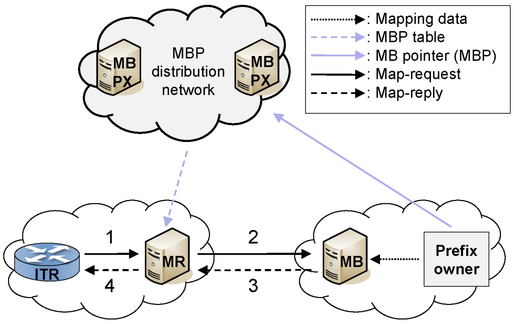

Figure 2: Basic structure and operation of FIRMS.

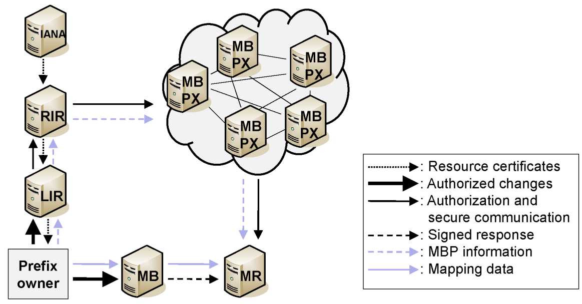

Figure 3: Security concept for FIRMS.

Map-Base Pointers (MBPs) contain aggregated information about EID-prefixes. This separation of mapping data and MBPs enables FIRMS to scale well even if the number of managed EIDs is going to be extremely large. MBPs are stored in a distributed table, the MBP distribution network or MBP table. This distribution network is composed of MBP eXchange nodes (MBPXs). The organization of MBPXs represents the organization of the authorities which are responsible for the assignment of EIDs and EID-prefixes.

Map Resolvers (MRs) combine the information of the MBP table and the MBs. Each MR has a valid, complete copy of the MBP table and uses caching to speed up consecutive mapping look-ups. A mapping look-up contains an MBP table look-up and an MB query. ITRs use MRs to query the mapping system for information.

3.2 Basic Operation

An ITR which requires a mapping for an EID sends a map request to an MR. The MR knows which MB stores the requested mapping and forwards the map-request accordingly. The MB returns a map-reply containing the desired EID-to-RLOC mapping which is returned to the ITR. FIRMS assumes that EIDs are assigned to their owners in a block that can be described by an EID prefix. Each prefix owner provides an MB holding the EID-to-RLOC mappings for all its EIDs. The operation of the MB can also be delegated to a specialized company. A MBP is a data structure containing information about the MB. The prefix owner registers this information in the global MBP distribution network which collects all MBPs and provides all MRs with a full MBP table. When a MR receives a map-request for a certain EID, it can retrieve the address of the appropriate MB from that table and can forward the map-request. MRs and MBs must be globally reachable by RLOCs.

3.3 Security

The security concept is visualized in Figure 3. MRs rely on the authenticity of MBPs and EID-to-RLOC mappings provided by MBs. MBPs and mapping data can only be changed by the prefix owner. This information reaches the MR unchanged. We use RPKI to transfer the right-to-use for IP prefixes from IANA through the RIRs and LIRs to prefix owners. MRs receive the MBP information through a secured connection from authenticated MBPXs. They can trust the received MBPs and do not need to verify them before map-requests are issued to MBs. When an MB returns a digitally signed map-reply to the MR, the MR verifies the data with the MB's public key, which is provided in the MBP. This is necessary to prevent the data from being maliciously altered on the way to the MR.

4. RELATED WORK

In LISP, ETRs provide the same functionality as MBs. A mapping system in LISP defines the way how a map-request finds the ETR. To enable the exchange of the mapping system, LISP-Map-Server[2] defines a map-resolver interface for ITRs and a map-server interface for ETRs. LISP+ALT[3] is the currently preferred mapping system for LISP. It defines a semi-hierarchical overlay structure over which map-requests from the ITR are forwarded to the appropriate ETR in an efficient way. LISP-DHT[4] uses a distributed hash table for that purpose. The structure of LISP-TREE [5] is similar to the one of the DNS system. In [6], we propose a classification of mapping systems and compare FIRMS with many others in detail.

5. CONCLUSION

In this extended abstract, my research results gained during the last 2 years were presented. Part of the research was done during my time at the University of Würzburg. FIRMS has been published in [7], a simulation demo was presented at Euroview 2009[8], and a prototype running in the G-Lab testbed has been presented at Euroview 2010[9]. A comprehensive summary of existing mapping systems is currently under submission, a preliminary report can be found in [6].

References

[1] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, "Locator/ID Separation Protocol (LISP)," http://tools.ietf.org/html/draft-ietf-lisp, Mar. 2011.

[2] D. Farinacci and V. Fuller, "LISP Map Server," http://tools.ietf.org/html/draft-ietf-lisp-ms, Apr. 2011.

[3] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, "LISP Alternative Topology (LISP+ALT)," http://tools.ietf.org/html/draft-ietf-lisp-alt, Mar. 2011.

[4] L. Mathy and L. Iannone, "LISP-DHT: Towards a DHT to Map Identifiers onto Locators," in *Re-Architecting the Internet (ReArch)*, Madrid, Spain, Dec. 2008.

[5] L. Jakab, A. Cabellos-Aparicio, F. Coras, D. Saucez, and O. Bonaventure, "LISP-TREE: A DNS Hierarchy to Support the LISP Mapping System," *IEEE Journal on Selected Areas in Communications, Special Issue on Internet Routing Scalability*, vol. 28, no. 8, Oct. 2010.

[6] M. Menth, M. Hartmann, and M. Hoefling, "Mapping Systems for Loc/ID Split Internet Routing," University of Würzburg, Institute of Computer Science, Technical Report, No. 472, May 2010.

[7] ——, "FIRMS: A Mapping System for Future Internet Routing," *IEEE Journal on Selected Areas in Communications, Special Issue on Internet Routing Scalability*, vol. 28, no. 8, Oct. 2010.

[8] ——, "Demo: a Future InteRnet Mapping System (FIRMS)," *Wuerzburg Workshop on IP: Visions of Future Generation Networks (EuroView)*, Jul. 2009.

[9] M. Hartmann, D. Hock, M. Hoefling, T. Neubert, and M. Menth, "FIRMS – Demonstration of a Mapping System for Loc/ID Split Internet Routing in G-Lab," *Wuerzburg Workshop on IP: Visions of Future Generation Networks (EuroView)*, Aug. 2010.