

FIRMS – Demonstration of a Mapping System for Loc/ID Split Internet Routing in G-Lab

Matthias Hartmann, David Hock, Michael Höfling, Tim Neubert, and Michael Menth
University of Wuerzburg, Germany

I. INTRODUCTION

The interdomain routing system faces scalability and flexibility problems caused by the coupling of endpoint identifiers (EIDs) and routing locators (RLOCs) in today’s IP addresses. More scalable architectures (e.g. LISP [1]) based on the locator/identifier (Loc/ID) split paradigm are currently developed [2]. They separate addressing into two separate parts: the EID, which identifies an end system within the edge network, and the RLOC, which indicates the location of the edge network in the Internet. In LISP domains, data packets are addressed with EIDs and forwarded to an ingress tunnel router (ITR) that connects the edge network to the core Internet (see Figure 1). The ITR uses a special mapping service to obtain an RLOC for the destination EID and tunnels the packets to the egress tunnel router (ETR) with that RLOC. The ETR decapsulates the packets and forwards them to the destination.

We developed the distributed EID-to-RLOC mapping service FIRMS (“Future InterNet Mapping System”) [3], [4] that can be used with any LISP-like Loc/ID split architecture. We implemented a prototype of FIRMS that is currently running in the German-Lab (G-Lab) experimental facility [5].

In the following, we explain the FIRMS architecture and show its design features. Then we describe the implementation and its deployment in G-Lab, and finally provide technical details for the demonstration setup.

II. THE FIRMS ARCHITECTURE

Figure 2 illustrates the basic structure and operation of FIRMS. We assume that EIDs are assigned to their owners in prefix blocks. Each prefix owner provides a map-base (MB) holding the EID-to-RLOC mappings for all its EIDs. The operation of the MB may be delegated to a specialized company. The mappings from EID-prefix to MB are stored in a map-base pointer (MBP). FIRMS has a global MBP distribution network, where all prefix owners register their MBPs. Each ITR is configured with a map-resolver (MR). ITRs and MRs are different entities, but the MR functionality may be integrated in an ITR to save communication overhead. The MR registers at the MBP distribution network and receives a copy of the global MBP table. When the ITR requires an EID-to-RLOC mapping for an EID, it sends a map-request to its MR. The MR looks up the address of the responsible MB in its local copy of the MBP table and forwards the map-request to that MB. The MB returns a map-reply containing the desired EID-to-RLOC mapping to the MR which forwards it to the ITR. If a non-existing mapping is queried, a negative map-reply is returned.

This work was funded by the Federal Ministry of Education and Research of the Federal Republic of Germany (support code 01 BK 0800, G-Lab). <http://www.german-lab.de/>

ITRs and MRs have local caches with appropriate time-to-live values to accelerate the lookup process and to minimize the number of map-requests. MRs and MBs must have globally reachable RLOC addresses.

We assume that EIDs are assigned hierarchically, similar to IP addresses today. EID address blocks are delegated to the five regional Internet registries (RIRs). They delegate subsets thereof to local Internet registries (LIRs). Both RIRs and LIRs partition the address space in prefix blocks and assign prefixes to organizations (prefix owners). Every RIR or LIR runs a map-base pointer exchange node (MBPX). Figure 3 shows that the MBPX of a LIR (LIR-MBPX) is connected to the MBPX of its RIR, and the MBPXs of the RIRs (RIR-MBPX) are fully meshed. This constitutes the MBP distribution network. The prefix owner adds, changes, or removes MBPs for its EID prefixes at the MBPX of its LIR or RIR. An LIR-MBPX forwards this data to its superordinate RIR-MBPX. The RIR-MBPX collects the MBPs for all EID prefixes under its control and compiles a regional MBP table. The MBP tables are exchanged among all RIR-MBPXs so that each of them has a copy of the global MBP table. They push this information to their subordinate LIR-MBPXs which forward it to all MRs that have registered for that service.

When an ITR receives a packet addressed to an outbound EID, it tries to retrieve the EID-to-RLOC mapping from its local cache and, if successful, tunnels the packet to the ETR whose RLOC was given in the mapping. In case of a cache miss, the ITR retrieves the mapping over the network which is a




Fig. 1. Packet flow and destination addresses in a Loc/ID split architecture.

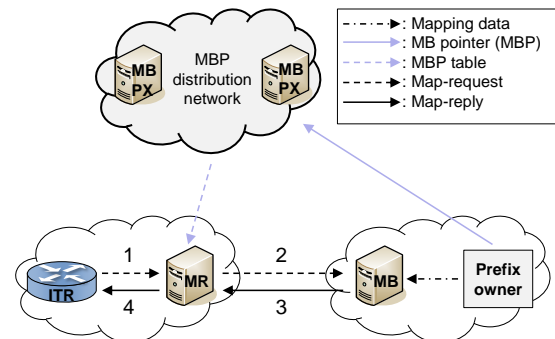


Fig. 2. Basic operation of FIRMS.

time-consuming action. In the meantime, packets must either be cached or dropped, which both has several drawbacks. FIRMS uses packet relaying to handle the outbound packet: In case of a cache miss, the ITR tunnels the packet to the MR. If the MR finds the required mapping in its cache, it tunnels the packet to the ETR. Otherwise, the MR tunnels the packet to the appropriate MB. The MB has the mapping in its database and tunnels the packet to the ETR. Every tunneled packet arriving at an MR or MB is also interpreted as a mapping request and the mapping is sent to the ITR. This design has several nice properties that are listed in [3].

The FIRMS design has several important features: The global MBP distribution network stores only stable pointers for EID-prefixes. This information rarely changes, not even when networks change their ISP. Mappings for individual EIDs can be changed at the MB without any global notification. Even mobility can be supported this way. Requests for individual EIDs only involve the sender’s MR and the receiver’s MB. No additional public infrastructure is required to resolve a mapping. This also keeps the lookup delay short. In addition, packets that arrive before a mapping is available can be relayed with the FIRMS system, using only the MB of the destination. This usually avoids long triangle routing and does not burden the remaining global mapping system.

III. DEMONSTRATION OF FIRMS IN THE G-LAB TESTBED

Our prototype implementation of the FIRMS mapping service is running in the German-Lab (G-Lab) experimental facility. We use five nodes as RIR-MBPXs, nine nodes as LIR-MBPXs, and 19 nodes as MRs. Another 60 nodes serve as MBs that store the actual mapping information. In addition, we use three laptops at the demonstration site to illustrate the FIRMS features. All involved nodes in this setup send log-messages to a dedicated log-server which monitors the FIRMS operation. A graphical Java application called FirmsLogViewer (see Figure 4) is running on our first laptop and connects to the log-server. It shows all previously described nodes of the mapping service, visualizes current activity, and can be used to control the mapping service.

For our demonstration, we use a basic Loc/ID-split application that emulates a node as well as the corresponding ITR or ETR. This application sends EID-based packets. A lower layer in the application, which represents the ITR functionality, queries the mapping service for appropriate locators (i.e., IP addresses). Packets are then forwarded to the application on the destination host. Different nodes in G-Lab run the Loc/ID-split




Fig. 3. Propagation of MBP updates.

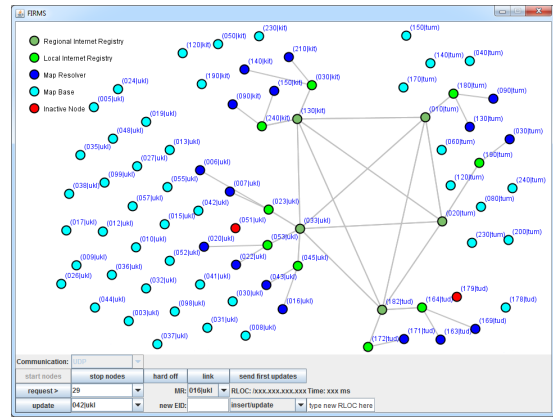


Fig. 4. Visualization of mapping service activity in the FirmsLogViewer.

application and use the FIRMS mapping service to send packets to other nodes running the same application. Our second and third demonstration laptops are also running this application and can send and receive EID-addressed packets.

The demo consists of several different scenarios. We first show the regular FIRMS behavior. For this purpose, the second laptop and several other nodes in G-Lab send packets to an EID that belongs to the application on our third laptop. We monitor lookup delays in the mapping service and path delay between sender and receiver, and show the packet relaying functionality. Then we demonstrate mobility support by changing the mapping of our EID in the MB. The packets are then forwarded to the second laptop instead of the third laptop. We show the switching delay and visualize the resulting packet flow. We also show the same effects when the MB is changed. This involves the whole MBP distribution network and results in longer switching delay. Finally, we illustrate the resilience features of FIRMS by deliberately disabling different components of our architecture.

IV. CONCLUSION

We summarized FIRMS, a fast two-level mapping system for Loc/ID split routing architectures that includes security and resilience features as well as a packet relay service. We demonstrated a proof-of-concept implementation of FIRMS in the G-Lab experimental facility and showed its basic operations as well as its resilience features. In [3] we have shown that FIRMS has structures in common with many other mapping system, but clearly differs in its overall design and stands out in the sum of the achieved benefits.

REFERENCES

- [1] D. Meyer, L. Zhang, and K. Fall, “RFC4984: Report from the IAB Workshop on Routing and Addressing,” Sep. 2007.
- [2] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, “Locator/ID Separation Protocol (LISP),” <http://tools.ietf.org/html/draft-ietf-lisp-07>, Apr. 2010.
- [3] M. Menth, M. Hartmann, and M. Hoefling, “Mapping Systems for Loc/ID Split Internet Routing,” www3.informatik.uni-wuerzburg.de/TR/tr472.pdf, University of Würzburg, Tech. Rep. 472, 2010.
- [4] —, “FIRMS: a Future InterNet Mapping System,” accepted for IEEE Journal on Selected Areas in Communications, 2010.
- [5] <http://www.german-lab.de/>, 2010.