# Architectural Considerations for Mapping Distribution Protocols

Kotikalapudi Sriram, Young-Tak Kim, and Doug Montgomery

**Abstract**

In this contribution, we present a discussion of some architectural ideas pertaining to the mapping distribution protocol. The efficiency of this protocol in terms of response time and the volume of traffic load it generates are important considerations. We consider how Egress Tunnel Routers (ETRs) can perform aggregation of end point ID (EID) address space belonging to their downstream delivery networks. This aggregation may be useful for reducing the processing load and memory consumption associated with mapping messages, especially in some resource-constrained components of the mapping distribution system. Some interesting architectural issues, their potential solutions and trade-offs are discussed. The overarching goal is to expose and discuss some subtleties in design considerations for mapping distribution and management.

## I. Introduction

We present a discussion of some architectural principles and issues pertaining to the mapping distribution protocol. The efficiency of this protocol in terms of response time and the volume of traffic load it generates are important considerations. We consider how Egress Tunnel Routers (ETRs) can perform aggregation of end point ID (EID) address space belonging to their downstream delivery networks. We discuss incorporation of an exception message as part of the map announcement message to indicate that portions of the ID space (some small number of more specific prefixes or subprefixes) under a less specific prefix have moved to or reside at different ETRs. This aggregation may be useful for reducing the processing load and memory consumption associated with map messages, especially at some resource-constrained components of the mapping distribution system. Some interesting architectural issues, their potential solutions and trade-offs are discussed. The general goal is to expose and discuss some subtleties

The authors are with the National Institute of Standards and Technology, Gaithersburg, MD, USA (Email: {ksriram, young-tak.kim, dougm}@nist.gov).

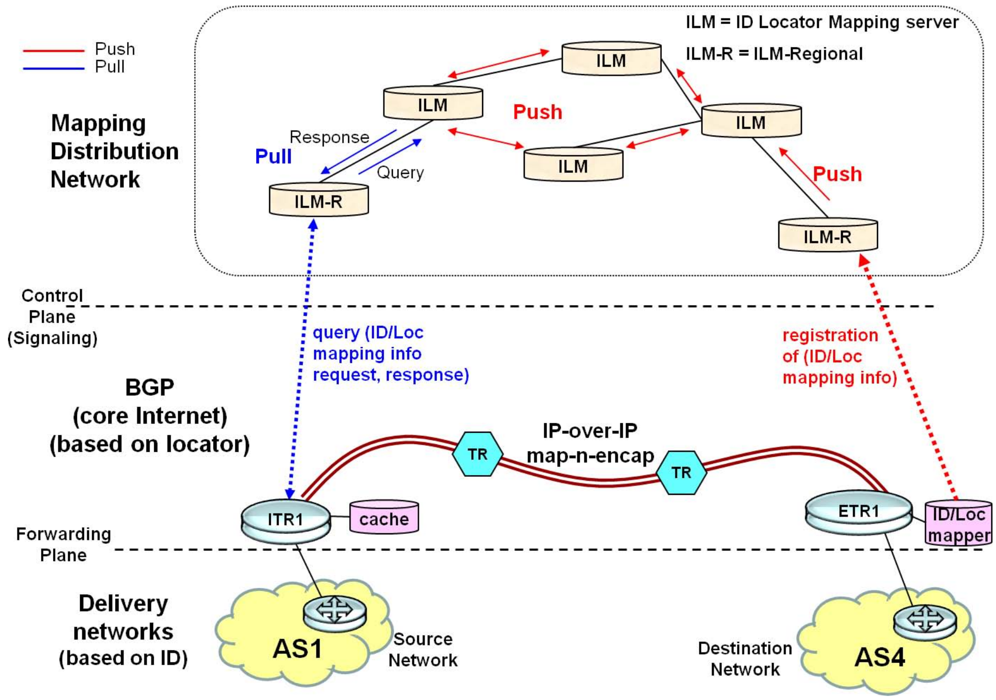


Fig. 1. An example of map-and-encap and mapping distribution architecture.

in the mapping distribution protocol early so that some architectural roadblocks are anticipated and possibly averted, and better efficiency and performance are achieved.

## II. Distribution and Management of Mapping Messages: Aggregation and Multi-Homing Considerations

### *A. Management of Mapping Distribution of Subprefixes Spread Across Multiple ETRs*

We feel that the architectural principles examined here are generally applicable to any of the proposals being discussed in the RRG [1]-[8] and their associated mapping distribution protocols. However, just to assist in this discussion, we start with the high level architecture of a map-and-encap approach as illustrated in Fig. 1. This helps anchor the discussion of the principles of the mapping distribution protocol to an architectural framework; the specific architecture can, however, vary and the ideas presented here are generally relevant to any of the proposals that are currently being reviewed in RRG. As shown in Fig. 1, the Egress Tunnel Routers (ETRs) generate map messages to inform the ID-Locator Mapping (ILM) servers of end-point ID prefixes or delivery-network address ranges that can be reached through them. The ILMs are repositories for complete mapping information, while the ILM-Regional (ILM-R)

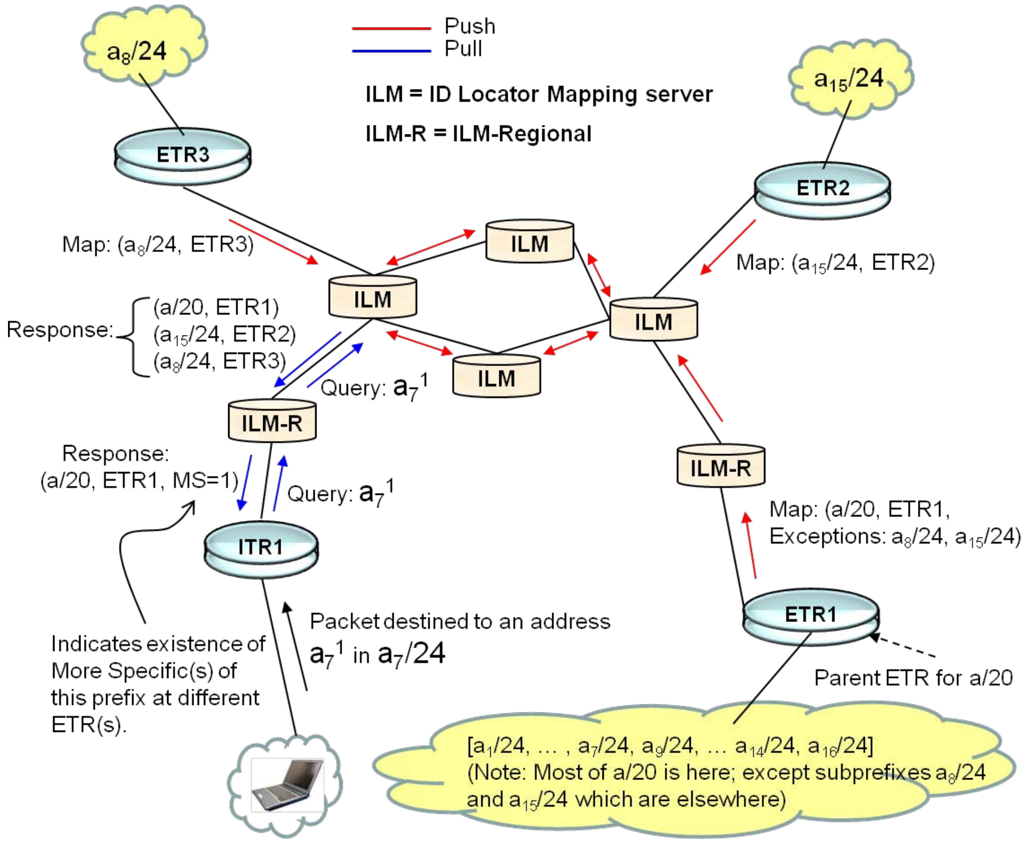


Fig. 2. Illustration of endpoint ID aggregation and exception concepts.

servers can contain partial and/or regionally relevant mapping information. The ETRs can push ID-to-locator mapping information pertaining to the delivery networks under their purview to an ILM-R or an ILM they are associated with. The ILM-Rs push the mapping information received from ETRs to the ILMs, and the ILMs update each other by pushing the mapping updates as they are received. When an Ingress Tunnel Router (ITR) does not have mapping information for an EID, it initiates pull requests (query/response) to the ILM-R that it has access to. The ILM-R can then pull the relevant mapping information from an ILM (if necessary) and respond back to the ITR.

The key points that we wish to discuss in this contribution are illustrated in Fig. 2 and Fig. 3. With the help of Fig. 2, it is illustrated that while a large endpoint address space contained in a prefix may be mostly associated with the delivery networks served by an ETR, some fragments (subprefixes) of that address space may be located elsewhere at different ETRs. Let $a$/20 denote a prefix that is conceptually viewed as composed of 16 subprefixes of /24 size that are denoted as $a_1$/24, $a_2$/24, ..., $a_{16}$/24. In Fig. 2, for example, $a$/20 is mostly at ETR1, while only two of its subprefixes $a_8$/24 and $a_{15}$/24 are elsewhere at

ETR3 and ETR2, respectively. From the point of view of efficiency of the mapping distribution protocol, it may be beneficial for ETR1 to announce a map for the entire space $a$/20 (rather than fragment it into a multitude of more-specific prefixes), and provide the necessary exceptions in the map information. Thus the map message could be in the form of *Map:(a/20, ETR1; Exceptions:* $a_8$*/24,* $a_{15}$*/24)*. In addition, ETR2 and ETR3 announce the maps for $a_{15}$/24 and $a_8$/24 respectively, and so the ILMs know where the exception endpoint ID addresses are located.

Now the interesting question is in what ways does the mapping distribution system communicate such mapping information to the ITRs? This issue is illustrated in the lower left-hand portion of Fig. 2. The sending host initiates a packet destined for an address $a_7^1$, which is in $a_7$/24 and hence in the normal portion of $a$/20 and not in the exception portion. In this case, we assume that the ITR1 does not have the map information, so it sends a query to its ILM-R. Assuming that the ILM-R does not have the map information either, it will send a query to an ILM it is connected with. The ILM can then send a response to the ILM-R, and this response can contain the full prefix information regarding ($a$/20, ETR1) as well as the maps for the exceptions or the subprefixes that are elsewhere, namely, ($a_{15}$/24, ETR2) and ($a_8$/24, ETR3). Now a question arises as to which of the following approaches would be the best choice:

1) ILM-R provides the complete mapping information for $a$/20 to ITR1 including all the maps for the relevant exception subprefixes.
2) ILM-R provides only the directly relevant map to ITR1 which in this case is ($a$/20, ETR1).
3) The mapping information transaction between ILM-R and ITR1 can dynamically use approach 1 or approach 2 above depending on the context (further explanation of this is provided below).

In the first approach, the advantage is that ITR1 would have the complete mapping for $a$/20 (including exception subprefixes), and it would not have to generate queries for subsequent first packets that are destined to any address in $a$/20, including $a_8$/24 and $a_{15}$/24. This would be true as long as ITR1 holds the received mapping information in its memory. However, the disadvantage is that if there is a significant number of exception subprefixes (in the example in consideration there are only two but in general it can be many more), then the very first packet destined for $a$/20 will experience a long delay, and also the processors at ITR1 and ILM-R can experience overload. In addition, the memory usage at ITR1 can be very inefficient as well.

The advantage of the second approach above is that the ILM-R does not overload resources at the ITR both in terms of processing and memory usage. This will help avoid resource exhaustion at the ITRs and thus provide better response time in handing mapping queries for the packets received from hosts. However, some care must be exercised here. ITR1 might save the mapping information for $a$/20 and reuse it. So it should be at least made aware that possibly there are subprefixes under $a$/20 that are at other ETRs. This can be taken care of by incorporating an More Specific (MS) indicator in the map message sent to ITR1 as shown in Fig. 2. This map response is of the form *Map:(a/20, ETR1, MS=1)*. The MS indicator is set to 1 to indicate to ITR1 that not all addresses in $a$/20 map to ETR1, and accordingly ITR1 will re-enquire next time there is a packet destined for another address in $a$/20 (other than $a_7^1$). One can go into the details of this methodology and improve it some more but those details can be discussed later. The key idea is that aggregation is beneficial and subprefix exceptions must be handled with additional messages or indicators in the map.

The third approach above seeks adaptability to use either the first approach or the second depending on the context. Here a parameter can be potentially defined such as the maximum number of maps (map for the parent prefix plus all the maps for the relevant exceptions subprefixes) that would be involved in a mapping transaction. If this parameter is below a threshold, then the first approach would be used, else, the second approach would be used. This parameter can also be tuned administratively or dynamically (depending on the state in terms of resource usage of the ILM and/or the ITR).

We present the above approaches to provide examples of design methodology and to raise some questions that need to be further studied in order to improve the efficiency and performance of the mapping distribution protocol.

### *B. Management of Mapping Distribution for Scenarios with Hierarchy of ETRs and Multi-Homing*

Now we refer to Fig. 3 to highlight another architectural concept related to mapping management. Here we consider the possibility that ETRs may be organized in a hierarchical manner. For instance ETR7 in Fig. 3 is higher in the loose hierarchy relative to ETR1, ETR2, and ETR3, and like-wise ETR8 is higher relative to ETR4, ETR5, and ETR6. For instance, ETRs 1 through 3 can relegate locator role to ETR7 for their EID address space. In essence, they can allow ETR7 to act as the locator for the delivery networks in their purview. ETR7 keeps a local mapping table for mapping the appropriate EID address space to specific ETRs that are hierarchically associated with it in the level below. In this situation, ETR7 can perform EID address space aggregation across ETRs 1 through 3 and can also include its own immediate EID address space for the purpose of that aggregation. Thus in the example of Fig. 3, ETR7 can aggregate $x_1$/24, $x_2$/24, $x_3$/24, $x_4$/24 into a map message of the form *Map:(x/22, ETR7)* to inform ILM-R. Similarly, ETR8 can aggregate $y_1$/24, $y_2$/24, $y_3$/24, $y_4$/24 into a map message of the form *Map:(y/22, ETR8)*. This architectural principle again lessens the possibility of cluttering the mapping distribution system with

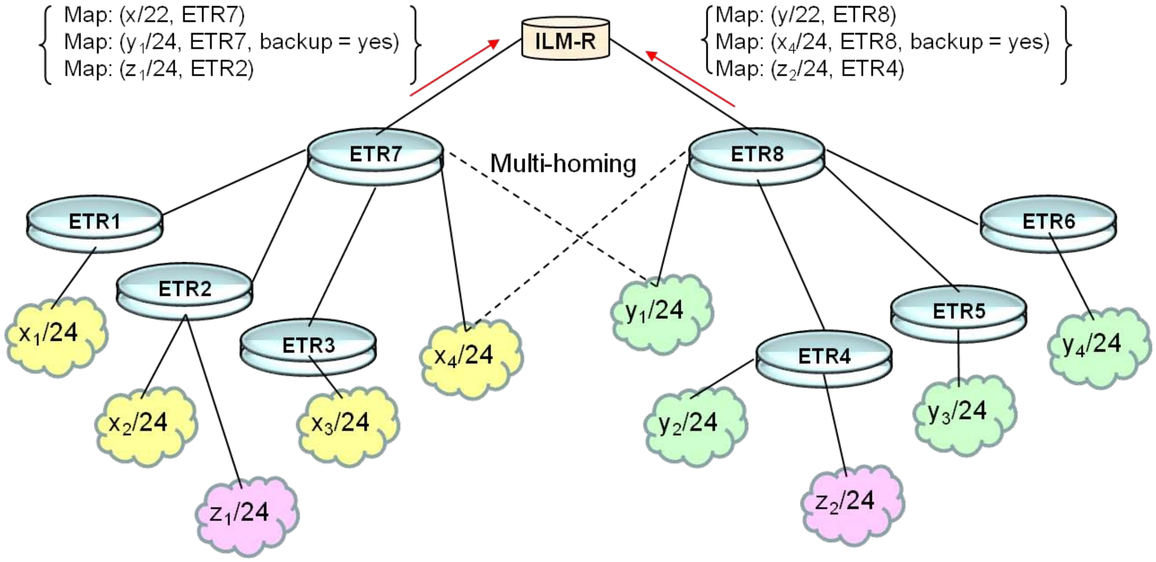


Fig. 3. Illustration of endpoint ID aggregation in the presence of multihoming.

excessive map messages or entries. It may be noted that ETR7 and ETR8 may or may not hide the ETRs below in the hierarchy. This is because there may be prefixes (subnets) such as $z_1$/24 at ETR2 and $z_2$/24 at ETR4, which benefit by simply announcing *Map:(*$z_1$*/24, ETR2)* and *Map:(*$z_2$*/24, ETR4)*, respectively. The idea is that when an ETR at the higher level is not able to aggregate EID prefixes of some ETRs below in the hierarchy, the routing locator may as well be the actual ETR where the delivery network resides.

Another architectural consideration for mapping distribution and management arises when some delivery networks (i.e., EID address space prefixes) are multi-homed to different ETRs. Examples of this in Fig. 3 are $x_4$/24 and $y_1$/24 which are each multi-homed to ETR7 and ETR8. In such case, ETR7 distributes a map message for $y_1$/24 which is a more specific prefix of $y$/22 for which ETR8 distributes a map message. The map message for such subprefix multi-homing situations may incorporate a field as shown in Fig. 3 to indicate that the ETR in the map provides a backup or lower priority path. But it is up to the ITRs as to how they make use of the backup/priority information in such cases. They may go by the priority information in the map message or they may prefer more specifics in any case. Recommendations for those actions or choices can be made in the future based on performance and operational considerations. In summary, the take away from Fig. 3 is that additional EID address space aggregation opportunities exist if there is at least a loose hierarchical structure to the ETRs. However, some minor complexities can also arise in some multi-homing cases, which can be handled by suitable recommendations. These recommendations can be established in the future based on network operators' goals and preferences.

The illustration in Fig. 3 also brings to mind a question about whether it would be a good design choice to recursively perform map-and-encap routing via a hierarchy of ETRs. This question has been brought up for discussion earlier in the RRG email list by others. The hierarchical organization of ETRs and delivery networks potentially helps in the future growth and scalability of mapping distribution protocol as well. With recursive map-and-encap, some of the mapping distribution and management functionality will remain local to topologically neighboring delivery networks which are hierarchically underneath ETRs that serve the transit network.

## III. Conclusion

We have discussed various architectural questions related to the mapping distribution and management. We have raised some questions regarding aggregation possibilities for the EID address space associated with delivery networks that are all homed at a the same ETR. We introduced special considerations such as exceptions in the sense that some subprefixes may be located away at different ETRs other than the ETR that is home to bulk of a covering less specific prefix. We also discussed the notion of a loose hierarchy of ETRs with the potential benefit of aggregation of their EID address spaces. The purpose of this contribution is to help generate some discussion of these issues. The overarching goal is to expose and discuss some subtleties in design considerations for mapping distribution and management. Hope is that early awareness of some of these issues may result in anticipating and averting some architectural roadblocks. The end goal is better efficiency and performance for the mapping distribution protocol.

We are planning modeling and simulations to quantify the tradeoffs associated with various architectural concepts pertaining to scalable addressing & routing and mapping distribution protocols, and hope to report such results in the future.

## Acknowledgment

The authors would like to thank Okhee Kim, Patrick Gleichmann, Oliver Borchert, and Rick Kuhn for comments and suggestions.

## References

[1] D. Farinacci et al., "Locator/ID Separation Protocol (LISP)," IETF Internet Draft, draft-farinacci-lisp-05.txt, November 14, 2007.

[2] R. Whittle, "Ivip (Internet Vastly Improved Plumbing) Conceptual Summary and Analysis," February 2008. http://www.firstpr.com.au/ip/ivip/Ivip-summary.pdf

[3] C. Vogt, "Six/One Router  Design and Motivation," July 2008. http://users.piuha.net/chvogt/pub/2008/vogt-2008-six-one-router-design.pdf

[4] C. Vogt, "Six/One Router: A Scalable and Backwards-Compatible Solution for Provider-Independent Addressing," ACM International Workshop on Mobility in the Evolving Internet Architecture (MobiArch), Seattle, WA, USA, August 2008.

[5] D. Jen et al., "APT: A Practical Tunneling Architecture For Routing Scalability," February 2008, http://www.cs.ucla.edu/ meisel/apt-rrg.pdf

[6] S. Brim et al., "LISP-CONS: A Content distribution Overlay Network Service for LISP", IETF Internet Draft, draft-meyer-lisp-cons-03.txt, November 2007.

[7] D. Farinacci et al., "LISP Alternative Topology (LISP-ALT)," IETF Internet Draft, draft-fuller-lisp-alt-01.txt, November 2007.

[8] L. Iannone and O. Bonaventure, "On the Cost of Caching Locator/ID Mappings," 3rd Annual CoNEXT Conference, December 2007. http://inl.info.ucl.ac.be/system/files/Conext-2007-CRV-UCL-v2-Clean.pdf