

Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 29, 2013

L. Iannone
Telecom ParisTech
D. Lewis
Cisco Systems, Inc.
D. Meyer
Brocade
V. Fuller
February 25, 2013

LISP EID Block
draft-ietf-lisp-eid-block-04.txt

Abstract

This is a direction to IANA to allocate a /16 IPv6 prefix for use with the Locator/ID Separation Protocol (LISP). The prefix will be used for local intra-domain routing and global endpoint identification, by sites deploying LISP as EID (Endpoint Identifier) addressing space.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 29, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Definition of Terms	3
3.	Rationale and Intent	5
4.	Expected use	6
5.	Block Dimension	7
6.	Action Plan	7
7.	Routing Considerations	8
8.	Security Considerations	8
9.	Acknowledgments	9
10.	IANA Considerations	9
11.	References	11
11.1.	Normative References	11
11.2.	Informative References	11
Appendix A.	Document Change Log	12
	Authors' Addresses	13

1. Introduction

This document directs the IANA to allocate a /16 IPv6 prefix for use with the Locator/ID Separation Protocol (LISP - [\[RFC6830\]](#)), LISP Map Server ([\[RFC6833\]](#)), LISP Alternative Topology (LISP+ALT - [\[RFC6836\]](#)) (or other) mapping system, and LISP Interworking ([\[RFC6832\]](#)).

This block will be used as global Endpoint IDentifier (EID) space ([Section 2](#)).

2. Definition of Terms

LISP operates on two name spaces and introduces several new network elements. This section provides high-level definitions of the LISP name spaces and network elements and as such, it must not be considered as an authoritative source. The reference to the authoritative document for each term is included in every term description.

Legacy Internet: The portion of the Internet that does not run LISP and does not participate in LISP+ALT or any other mapping system.

LISP site: A LISP site is a set of routers in an edge network that are under a single technical administration. LISP routers that reside in the edge network are the demarcation points to separate the edge network from the core network. See [\[RFC6830\]](#) for more details.

Endpoint ID (EID): An EID is a 32-bit (for IPv4) or 128-bit (for IPv6) value used in the source and destination address fields of the first (most inner) LISP header of a packet. A packet that is emitted by a system contains EIDs in its headers and LISP headers are prepended only when the packet reaches an Ingress Tunnel Router (ITR) on the data path to the destination EID. The source EID is obtained via existing mechanisms used to set a host's "local" IP address. An EID is allocated to a host from an EID-prefix block associated with the site where the host is located. See [\[RFC6830\]](#) for more details.

EID-prefix: A power-of-two block of EIDs that are allocated to a site by an address allocation authority. See [\[RFC6830\]](#) for more details.

EID-Prefix Aggregate: A set of EID-prefixes said to be aggregatable in the [\[RFC4632\]](#) sense. That is, an EID-Prefix aggregate is defined to be a single contiguous power-of-two EID-prefix block. A prefix and a length characterize such a block. See [\[RFC6830\]](#)

for more details.

Routing LOcator (RLOC): A RLOC is an IPv4 or IPv6 address of an egress tunnel router (ETR). A RLOC is the output of an EID-to-RLOC mapping lookup. An EID maps to one or more RLOCs. Typically, RLOCs are numbered from topologically aggregatable blocks that are assigned to a site at each point to which it attaches to the global Internet; where the topology is defined by the connectivity of provider networks, RLOCs can be thought of as Provider Aggregatable (PA) addresses. See [[RFC6830](#)] for more details.

EID-to-RLOC Mapping: A binding between an EID-Prefix and the RLOC-set that can be used to reach the EID-Prefix. The general term "mapping" always refers to an EID-to-RLOC mapping. See [[RFC6830](#)] for more details.

Ingress Tunnel Router (ITR): An Ingress Tunnel Router (ITR) is a router that accepts receives IP packets from site end-systems on one side and sends LISP-encapsulated IP packets toward the Internet on the other side. The router treats the "inner" IP destination address as an EID and performs an EID-to-RLOC mapping lookup. The router then prepends an "outer" IP header with one of its globally routable RLOCs in the source address field and the result of the mapping lookup in the destination address field. See [[RFC6830](#)] for more details.

Egress Tunnel Router (ETR): An Egress Tunnel Router (ETR) receives LISP-encapsulated IP packets from the Internet on one side and sends decapsulated IP packets to site end-systems on the other side. An ETR router accepts an IP packet where the destination address in the "outer" IP header is one of its own RLOCs. The router strips the "outer" header and forwards the packet based on the next IP header found. See [[RFC6830](#)] for more details.

Proxy ITR (PITR): A Proxy-ITR (PITR) acts like an ITR but does so on behalf of non-LISP sites which send packets to destinations at LISP sites. See [[RFC6832](#)] for more details.

Proxy ETR (PETR): A Proxy-ETR (PETR) acts like an ETR but does so on behalf of LISP sites which send packets to destinations at non-LISP sites. See [[RFC6832](#)] for more details.

Map Server (MS): A network infrastructure component that learns EID-to-RLOC mapping entries from an authoritative source (typically an ETR). A Map Server publishes these mappings in the distributed mapping system. See [[RFC6833](#)] for more details.

Map Resolver (MR): A network infrastructure component that accepts LISP Encapsulated Map-Requests, typically from an ITR, quickly determines whether or not the destination IP address is part of the EID namespace; if it is not, a Negative Map-Reply is immediately returned. Otherwise, the Map Resolver finds the appropriate EID-to-RLOC mapping by consulting the distributed mapping database system. See [[RFC6833](#)] for more details.

The LISP Alternative Logical Topology (ALT): The virtual overlay network made up of tunnels between LISP+ALT Routers. The Border Gateway Protocol (BGP) runs between ALT Routers and is used to carry reachability information for EID-prefixes. The ALT provides a way to forward Map-Requests toward the ETR that "owns" an EID-prefix. See [[RFC6836](#)] for more details.

ALT Router: The device on which runs the ALT. The ALT is a static network built using tunnels between ALT Routers. These routers are deployed in a roughly-hierarchical mesh in which routers at each level in the topology are responsible for aggregating EID-Prefixes learned from those logically "below" them and advertising summary prefixes to those logically "above" them. Prefix learning and propagation between ALT Routers is done using BGP. When an ALT Router receives an ALT Datagram, it looks up the destination EID in its forwarding table (composed of EID-Prefix routes it learned from neighboring ALT Routers) and forwards it to the logical next-hop on the overlay network. The primary function of LISP+ALT routers is to provide a lightweight forwarding infrastructure for LISP control-plane messages (Map-Request and Map-Reply), and to transport data packets when the packet has the same destination address in both the inner (encapsulating) destination and outer destination addresses ((i.e., a Data Probe packet). See [[RFC6836](#)] for more details.

3. Rationale and Intent

With the current specifications, if an ITR is sending to all types of destinations (i.e., non-LISP destinations, LISP destinations not in the IPv6 EID Block, and LISP destinations in the IPv6 EID Block) the only way to understand whether or not to encapsulate the traffic is to perform a cache lookup and, in case of cache-miss, send a Map-Request to the mapping system. In the meanwhile, packets may be dropped.

There are several use cases for this address block, for instance:

- o In certain circumstances it is possible to configure the router so to natively forward all packets that have not a destination

address in the block, without performing any lookup whatsoever.

- o In some scenarios, in case of cache-miss packets, are routed toward a PETR until a mapping is obtained, if the destination is in a specific EID space packets may be dropped in order to avoid forwarding paths like ITR->PETR->PITR->ETR, avoiding the related overhead.
- o Improved traffic engineering capabilities with respect to LISP vs. non-LISP traffic.

Is worth to mention that new use cases can arise in the future, due to new and unforeseen scenarios. furthermore, this will give a tighter control over the traffic in the initial experimental phase, while facilitating its large-scale deployment.

The EID Block will be used only at configuration level, it is recommended not to hard-code in any way the IPv6 EID Block in the router hardware. This allows avoiding locking out sites that may want to switch to LISP while keeping their own IPv6 prefix, which is not in the IPv6 EID Block.

4. Expected use

Sites planning to deploy LISP may request a prefix in the IPv6 EID Block. Such prefix will be used for routing and endpoint identification inside the site requesting it. Mappings related to such prefix, or part of it, will be made available through the mapping system in use or registered to one or more Map Server(s).

To guarantee reachability from the Legacy Internet the prefix could be announced in the BGP routing infrastructure by one or more PITR(s). The use of PxTRs allow to aggregate several prefixes; the deployment model for this element is described in [[RFC6832](#)] and [[I-D.ietf-lisp-deployment](#)].

As the LISP adoption progress, the EID prefix space will potentially help in reducing the impact on the BGP routing infrastructure with respect to the case of the same number of adopters using global unicast space allocated by RIRs ([[MobiArch2007](#)]). From a short-term perspective, the EID space offers potentially large aggregation capabilities since it is announced by PxTRs possibly concentrating several contiguous prefixes. Such trend should continue with even lower impact from a long-term perspective, since more aggressive aggregation can be used, potentially leading at using few PxTRs announcing the whole EID space ([[FIABook2010](#)]).

The prefix is not supposed to be used as normal prefix announced in the BGP routing infrastructure without the use of LISP.

5. Block Dimension

The working group reached consensus on an initial allocation of a /16 prefix out of a /12 block which is asked to remain reserved for future use as EID space. The reason of such consensus is manifold:

- o The working group agreed that /16 prefix is sufficiently large to cover initial allocation and requests for prefixes in the EID space in the next few years for very large-scale experimentation and deployment.
- o As a comparison, it is worth mentioning that the current LISP Beta Network ([\[BETA\]](#)) is using a /32 prefix, with more than 250 sites using a /48 sub prefix. Hence, a /16 prefix looks as sufficiently large to allow the current deployment to scale up and be open for interoperation with independent deployments using EIDs space in the new /16 prefix.
- o A /16 prefix is sufficiently large to only allow deployment of independent (commercial) LISP enabled networks by third parties, but may as well boost LISP experimentation and deployment.
- o The /16 size and alignment allows the use to current policies to allocate and distribute prefixes out of this space, without the need to introduce any new specific address management policy.
- o The proposed alignment provides as well a natural support for DNS. In particular, reverse DNS for IPv6 in the special ip6.arpa domain is represented as sequence of nibbles. A different alignment would force to a binary representation.
- o The use of a /16 prefix is in line with previous similar prefix allocation for tunnelling protocols ([\[RFC3056\]](#)) and is considered a useful practice ([\[RFC3692\]](#)).

6. Action Plan

This document requests IANA to initially assign a /16 prefix out of the IPv6 addressing space for use as EID in LISP (Locator/ID Separation protocol).

It is suggested to IANA to temporarily avoid allocating any other address block the same /12 prefix the EID /16 prefix belongs to.

This is to accommodate future requests of EID space without fragmenting the EID addressing space. This will also help from an operational point of view, since it will be sufficient to change the subnet mask length in existing deployments. If in the future there will be need for a larger EID Block the address space adjacent the EID Block could be allocate by IANA according to the current policies.

IANA should assign the requested address space by September 2013 for a duration of 10 (ten) years (through September 2023). By the end of this period, the IETF will provide a decision on whether to transform the prefix in a permanent assignment or to put it back in the free pool.

The allocation and management of the Global EID Space will be detailed in a separate document.

7. Routing Considerations

In order to provide connectivity between the Legacy Internet and LISP sites, PITRs announcing large aggregates of the IPv6 EID Block could be deployed. By doing so, PITRs will attract traffic destined to LISP sites in order to encapsulate and forward it toward the specific destination LISP site. Routers in the Legacy Internet must treat announcements of prefixes from the IPv6 EID Block as normal announcements, applying best current practice for traffic engineering and security.

Even in a LISP site, not all routers need to run LISP elements. In particular, routers that are not at the border of the local domain, used only for intra-domain routing, do not need to provide any specific LISP functionality but must be able to route traffic using addresses in the IPv6 EID Block.

For the above-mentioned reasons, routers that do not run any LISP element, must not include any special handling code or hardware for addresses in the IPv6 EID Block. In particular, it is recommended that the default router configuration does not handle such addresses in any special way. Doing differently could prevent communication between the Legacy Internet and LISP sites or even break local intra-domain connectivity.

8. Security Considerations

This document does not introduce new security threats in the LISP architecture nor in the Legacy Internet architecture.

9. Acknowledgments

Special thanks to Roque Gagliano for his suggestions and pointers. Thanks to Brian Carpenter, Roger Jorgensen, Terry Manderson, Brian Haberman, Adrian Farrel, Job Snijders, Marla Azinger, Chris Morrow, and Peter Schoenmaker, for their insightful comments. Thanks as well John Curran, Paul Wilson, Geoff Huston, Wes George, Arturo Servin, Sander Steffann, and to all participants to the fruitful discussion on the IETF mailing list.

10. IANA Considerations

This document instructs the IANA to assign a /16 IPv6 prefix for use as the global LISP EID space using a hierarchical allocation as outlined in [[RFC5226](#)] and summarized in Table 1.

Attribute	Value
Address Block	XXX0::/16 [1]
Name	EID Space for LISP
RFC	[This Document]
Allocation Date	September 2013
Termination Date	September 2023
Source	True [2]
Destination	True
Forwardable	True
Global	True
Reserved-by-protocol	True [3]

[1] XXX value to be provided by IANA before published as RFC. [2] Can be used as a multicast source as well. [3] To be used as EID space by LISP [[RFC6830](#)] enabled routers.

Table 1: Global EID Space

During the discussion related to this document, the LISP Working Group agreed in suggesting to IANA to reserve adjacent addressing space, more specifically the /12 covering the assigned /16 prefix, for future use as EID space if needs come. Table 2 summarizes the request. Following the policies outlined in [[RFC5226](#)], such space will be assigned only upon IETF Review.

Attribute	Value
Address Block	XXX0::/12 [1]
Name) EID Space for LISP
RFC	[This Document]
Allocation Date	September 2013
Termination Date	September 2023
Source	False
Destination	False
Forwardable	False
Global	False
Reserved-by-protocol	True [2]

[1] XXX value to be provided by IANA before published as RFC. [2] To be used as EID space by LISP [[RFC6830](#)] enabled routers.

Table 2: Reserved for Future Use as Global EID Space

This document does not specify any specific value for the requested address block but suggests that should come from the 2000::/3 Global Unicast Space. Furthermore, it is suggested to assign the /16 prefix from the first /16 block out of the reserved /12 prefix. IANA is not requested to issue a AS0 ROA, since the Global EID Space will be used for routing purposes.

The reserved address space is requested for a period of time of ten years starting in September 2013 and ending in September 2023. Following the policies outlined in [[RFC5226](#)], upon IETF Review, by September 2023 decision should be made on whether to keep the assignment making the reserved prefix assignment permanent (this includes final decision on the size of the prefix). If the IETF review outcome will be that is not worth to have a reserved prefix as global EID space, the whole /12 (and all sub-block assigned out of it) will be took out from the IPv6 Special Purpose Address Registry and put back in the free pool managed by IANA.

Allocation and management of the Global EID Space is detailed in a different document. Nevertheless, all prefix allocations out this space must be temporary and no allocation must go beyond September 2023 unless the upon IETF Review the GLoBal EID Space is permanently assigned.

[11.](#) References

11.1. Normative References

- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", [BCP 122](#), [RFC 4632](#), August 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", [RFC 6830](#), January 2013.
- [RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", [RFC 6832](#), January 2013.
- [RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", [RFC 6833](#), January 2013.
- [RFC6836] Fuller, V., Farinacci, D., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)", [RFC 6836](#), January 2013.

11.2. Informative References

- [BETA] LISP Beta Network, "<http://www.lisp4.net>", 2008-2011.
- [FIABook2010] L. Iannone, T. Leva, "Modeling the economics of Loc/ID Separation for the Future Internet.", Towards the Future Internet - Emerging Trends from the European Research, Pages 11-20, ISBN: 9781607505389, IOS Press , May 2010.
- [I-D.ietf-lisp-deployment] Jakab, L., Cabellos-Aparicio, A., Coras, F., Domingo-Pascual, J., and D. Lewis, "LISP Network Element Deployment Considerations", [draft-ietf-lisp-deployment-06](#) (work in progress), January 2013.
- [MobiArch2007] B. Quoitin, L. Iannone, C. de Launois, O. Bonaventure, "Evaluating the Benefits of the Locator/Identifier Separation", The 2nd ACM-SIGCOMM International Workshop on Mobility in the Evolving Internet Architecture (MobiArch'07) , August 2007.

[RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", [RFC 3056](#), February 2001.

[RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", [BCP 82](#), [RFC 3692](#), January 2004.

[Appendix A](#). Document Change Log

Version 04 Posted February 2013.

- o Added Table 1 and Table 2 as requested by IANA.
- o Transformed the prefix request in a temporary request as suggested by various comments during IETF Last Call.
- o Added discussion about short/long term impact on BGP in [Section 4](#) as requested by B. Carpenter.

Version 03 Posted November 2012.

- o General review of [Section 5](#) as requested by T. Manderson and B. Haberman.
- o Dropped [RFC 2119](#) Notation, as requested by A. Farrel and B. Haberman.
- o Changed "IETF Consensus" to "IETF Review" as pointed out by Roque Gagliano.
- o Changed every occurrence of "Map-Server" and "Map-Resolver" with "Map Server" and "Map Resolver" to make the document consistent with [\[RFC6833\]](#). Thanks to Job Snijders for pointing out the issue.

Version 02 Posted April 2012.

- o Fixed typos, nits, references.
- o Deleted reference to IANA allocation policies.

Version 01 Posted October 2011.

- o Added [Section 5](#).

Version 00 Posted July 2011.

- o Updated section "IANA Considerations"
- o Added section "Rationale and Intent" explaining why the EID block allocation is useful.
- o Added section "Expected Use" explaining how sites can request and use a prefix in the IPv6 EID Block.
- o Added section "Action Plan" suggesting IANA to avoid allocating address space adjacent the allocated EID block in order to accommodate future EID space requests.
- o Added section "Routing Consideration" describing how routers not running LISP deal with the requested address block.
- o Added the present section to keep track of changes.
- o Rename of [draft-meyer-lisp-eid-block-02.txt](#).

Authors' Addresses

Luigi Iannone
Telecom ParisTech

Email: luigi.iannone@telecom-paristech.fr

Darrel Lewis
Cisco Systems, Inc.

Email: darlewis@cisco.com

David Meyer
Brocade

Email: dmm@1-4-5.net

Vince Fuller

Email: vaf@vaf.net

