

LISP Working Group
Internet-Draft
Intended status: Experimental
Expires: September 12, 2013

S. Barkai
ConteXtream Inc.
D. Farinacci
. .
D. Meyer
Brocade
F. Maino
V. Ermagan
Cisco Systems
March 11, 2013

LISP Based FlowMapping for Scaling NFV
draft-barkai-lisp-nfv-00

Abstract

This draft describes distributed flow-mapping applied according to RFC 6830 Locator ID Separation Protocol (LISP) for dynamic scaling of virtualized network functions (NFV) . Network functions such as subscriber management-mobility-security-quality, are typically delivered using proprietary appliances topologically embedded into the network as service-nodes or service-blades. Next generation virtualized network functions are pure software instances running on standard servers - unbundled building blocks of processing capacity and modular functionality. LISP based flow-mapping dynamically wires VNF instances into the data-path, and scales virtualized functions by steering the right traffic in the right sequence to the right process.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

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 September 12, 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	2
2. Connectivity Models Used	3
3. XTR FlowMapping Reference Architecture	6
3.1.	7
4. Intra-Provider Mappings	7
5. LCAF Mapping Subscription	7
6. Inter-Provider Mappings	7
7. QOS and Echo Measurements	7
8. Security Considerations	7
9. IANA Considerations	7
10. Acknowledgements	7
11. Normative References	7
Authors' Addresses	7

1. Introduction

This draft describes distributed flow-mapping applied according to RFC 6830 Locator ID Separation Protocol (LISP) for dynamic scaling of virtualized network functions (NFV) .[RFC6830]Network functions such as subscriber management-mobility-security-quality are typically delivered using proprietary appliances topologically embedded into the network as service-nodes or service-blades.

This monolithic service delivery method increases the complexity of roll-out and capacity planning, limits functional choices, and inhibits service innovation. Next generation virtualized network

functions are pure software instances running on standard servers - unbundled building blocks of compute capacity and modular functionality. This component based model opens up service provider networks to the savings of elasticity and the innovation of an open architectures. However this model also presents the network (or the virtual network rather) with the brand new challenges of assembling components into whole solutions by forwarding the right traffic to the right process at the right sequence and the right time.

While it is possible, to some extent, to use traditional virtual networking mechanisms such as virtual-LANs (VLAN) and virtual-private-networks (VPN) in-order to map traffic to functions-processes, these mechanisms are relatively static and require complex and intense configuration of physical network interfaces with vbridge vrf configuration. NGN software-defined flow based models on the other hand are much more programmable and dynamic, and if orchestrated properly offer a better fit to next generation service-provider data-centers. LISP based flow-mapping enables such a dynamic global orchestration of flows. LISP xTRs wire virtual function instances into the data-path, based on dynamic identity-function and identity-location mappings, perform these actions in a dynamically programmable and elastic manner, and operate based on subscriber-profiles and function-demand global information.

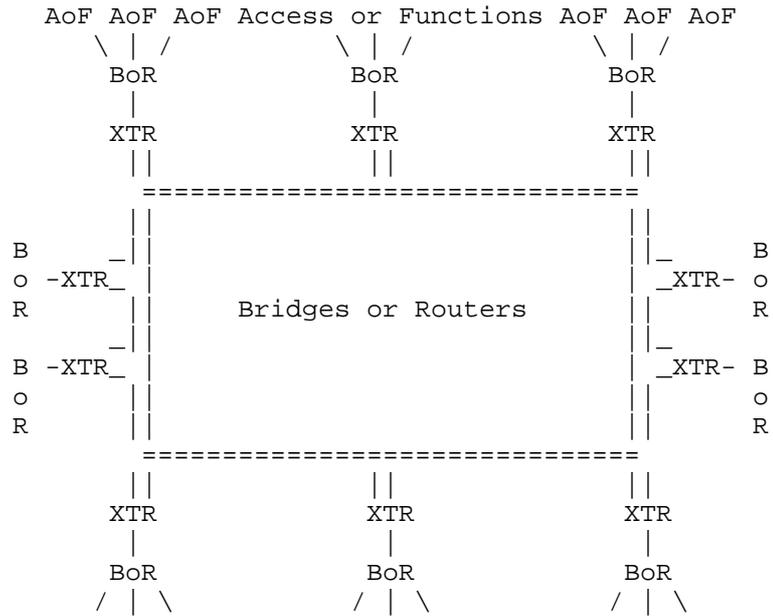
2. Connectivity Models Used

The basic connectivity model used to map flows to function is an identity-matrix forwarding. Unlike topological forwarding models which are based on source-subnet >> routed hop by hop >> to destination-subnet, identity-matrices are based on flow-identity "patched" to function-identity. This model is implemented using the LISP distributed overlay and LISP distributed database mechanisms. These mechanisms are applied over in-place physical networks in a manner described below:

- o The topological network basis or the "location" network is assumed to be implemented using standard bridging and routing. Basic design principles are applied in order to achieve both physical capacity and physical availability of connectivity. Typical examples of these practices include spine-leafs switching for cluster many server racks that are used as the compute and virtual compute foundations to functions. These practices also include core-edge routing for inter-connecting server clusters across points of presence, as well as for inter-connecting these points of presence to the access networks and to the public Internet.
- o The functional network or the "identity" matrix is there to map identified subscriber-flows carrying an application thread to the

	v <<	FunctionA	FunctionB	..	FunctionN	
	v					
Recursion	Instance1..i	Instance1..j	Instance1..k			
	v					
	v					
Subscriber1	o o o o	- - -+ o o o	- - -o o o o			
Subscriber2	o + o o	- - -o o o o	- - -o o o o			
.			
.			
.			
SubscriberM	o o o o	- - -o o o o	- - -+ o o o			

Functional Identity Matrix

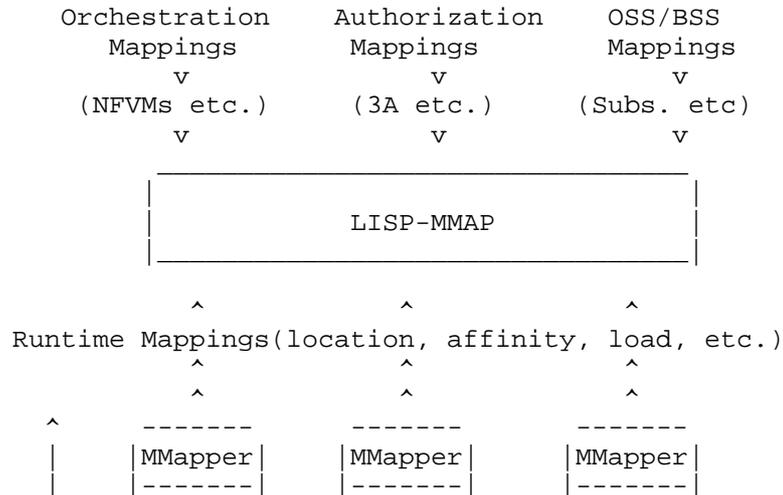


Identity-Location Overlay

3. XTR FlowMapping Reference Architecture

In order to map subscriber flows to virtualized function instances and essentially to overlay identity grid onto topology based bridge-routed network we use the following XTR 3-tier reference architecture:

1. Flow-Switching: is a set of n-tuple LOCALLY defined masks, BEST matched against each packet in-order to separate traffic to LOCALLY significant sequences representing application threads. Flows are either Encapsulated in-to the overlay, Decapsulated out-of the overlay, Forwarded up-to xTR internally registered Flow-Handlers ..
2. Flow-Handlers: are a set of control processes invoked for each flow where a specific identity-location mapping has not been defined and provisioned into the Flow-Switching. The default "catch-all" Flow-Handler maps IP flows to locations and gateways based on RFC 6830. In addition protocol-specific handlers can be loaded into the xTR for dealing with specific mapping and AFFINITY requirements of network functions such as SIP, GTP, S1X etc.
3. Global-Mappers: is how GLOBALLY significant key-value mappings is translated by Flow-Handlers to LOCALLY defines masks and encapsulation headers. Examples of such mappings include functional VIP to actual function processes EIDs, application specific SubscriberIDs to function EIDs, public IP-Port to SubscriberID, and EIDs to RLOCs.



X	F F F F	F F F F	F F F F
T	H H H H	H H H H	H H H H
R	n n n n	n n n n	n n n n
	d d d d	d d d d	d d d d
	-----	-----	-----
v	FlowX	FlowX	FlowX
	-----	-----	-----

Identity-Location Overlay Ring

3.1.

- 4. Intra-Provider Mappings
- 5. LCAF Mapping Subscription
- 6. Inter-Provider Mappings
- 7. QOS and Echo Measurements
- 8. Security Considerations

There are no security considerations related with this memo.

9. IANA Considerations

There are no IANA considerations related with this memo.

10. Acknowledgements

11. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, January 2013.

Authors' Addresses

Sharon Barkai
ConteXtream Inc.
California
USA

Email: Sharon@Contextream.c

Dino Farinacci
.
California
USA

Email: farinacci@gmail.com

David Meyer
Brocade
California
USA

Email: dmm@1-4-5.net

Fabio Maino
Cisco Systems
California
USA

Email: fmaino@cisco.com

Vina Ermagan
Cisco Systems
California
USA

Email: vermagan@cisco.com