LISP
Internet-Draft
Intended status: Informational
Expires: August 24, 2014

S. Freitas
P. Bellagamba
Y. Hertoghs
Cisco Systems

February 20, 2014

Using LISP for Secure Hybrid Cloud Extension
draft-freitasbellagamba-lisp-hybrid-cloud-use-case-00

Abstract

The purpose of this draft is to document how the Locator/Identifier Separation Protocol (LISP) can be used to enable a secure layer 3-based Hybrid Cloud, which is a composition of two or more distinct cloud infrastructures that remain unique entities, but are bound together to enable data and application portability.

It describes how LISP, in combination with IPsec, can be implemented on a virtualized router that is deployed on a public cloud and on the enterprise Data Center to enable cloud bursting, workload migration, rapid provision of new applications in the cloud and disaster recovery use cases. This draft covers how LISP allows virtual machines (VMs) to be moved to the cloud without requiring changes to the VMs IP and/or MAC-address.

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 RFC 2119 [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 August 24, 2014.

Table of Contents

1. Introduction

Many enterprises are pursuing an hybrid cloud computing strategy within the next three years. Some of the use cases for Hybrid Cloud are automated on-demand compute capacity (cloud bursting), split application architectures, workload migration, rapid provision of new applications in the cloud, and disaster recovery.

One common requirement from enterprises that wish to move to a Hybrid Cloud is the ability to migrate the servers to the Cloud without making any changes to them. In particular, server administrators would like to avoid changing the server's IP address, subnet mask and default gateway configurations. Also, enterprises would like to adopt their own IP addressing scheme in the cloud and not been limited by the addressing scheme of the cloud provider infrastructure.

Locator/Identifier Separation Protocol (LISP) [RFC6830] can be used to address those requirements. LISP separates location and identity, thus allowing the enterprise to migrate or create new servers on the cloud with the same IP address (server's identity), subnet mask, and default gateway configurations then the one used inside their own private data centers. LISP will update the EID-to-RLOC mapping of the server to reflect the new location that, in this this example, is moved to the cloud. No changes are required to the end systems, users or servers, as the mapping between identity (server's IP address) and location (enterprise DC or public cloud) is handled by LISP, transparently to the users trying to reach the server.

LISP operates as an overlay, encapsulating the original packet from the server into an UDP packet along with an additional outer IPvN header, which holds the source and destination RLOCs. This allows the server administrators to address the server in the cloud according to their own IP addressing scheme, independent of the cloud provider's addressing structure.

Another important property of LISP that is relevant to enable a Secure Hybrid Cloud extension is that it enables IP portability to the cloud by routing (layer 3) to the right location where the server

is. This provides total isolation of broadcast (Layer 2) domains between the Enterprise and the Public Cloud.

Non-LISP enabled sites communicates to the servers moved to the cloud via the enterprise data center (DC), where LISP is also deployed. The solution proposed in this draft does not require LISP to be enabled globally, but can be deployed by enabling LISP on just the enterprise DC and the public cloud, with minimal impact on the operations of both the DC and the cloud.

The optional deployment of LISP at individual user's site, provides data path optimization, as the LISP-encapsulated traffic is routed directly to the public cloud or the DC, depending on the server location.

The LISP service is provided in the network, to any virtual machine, independently of the hypervisor type used in the enterprise or the public cloud provider. In fact, the hypervisor type used on the enterprise and on the cloud infrastructure can be different. LISP works with any VM in the public cloud without the need to modify the VM configuration before migration.

LISP is deployed in virtualized routers in the cloud, and can be deployed in either physical or virtual router in the enterprise DC. The LISP-enabled router deployed within the enterprise DC does not need to be the default gateway for the local servers (physical and VMs).

The communication between the LISP-enabled router deployed within the

enterprise DC and within the public cloud SHOULD be secured by an IPsec (Internet Protocol Security (IPsec)) tunnel. LISP encapsulated traffic is protected with an IPsec tunnel, that can provide data origin authentication, integrity protection, anti-reply protection and confidentiality between the cloud and the enterprise.

The LISP-enabled Hybrid Cloud solution allows Intra-Subnet communication regardless of the location of the server. This means that communication between two servers located on the same subnet can happen even when one server is located at the enterprise DC and another server is located at the cloud. Inter-subnet communication is also supported.

2. LISP-Enabled Secure Hybrid Cloud Diagram

This diagram illustrates the use case described in this draft.

Users' end-devices, typically located at non-LISP sites, are connected to the enterprise WAN to access servers located either at the enterprise DC or at the public cloud.

The enterprise DC hosts some of the servers, and has a PxTR deployed attached to the subnets which host mobile servers.

The PxTR (PxTR-1) deployed within the enterprise DC is configured with an IPsec tunnel towards the cloud's xTR (xTR-2).

Some servers have been moved to the cloud and an xTR is deployed within the cloud to allow IP mobility between the enterprise DC and the public cloud.

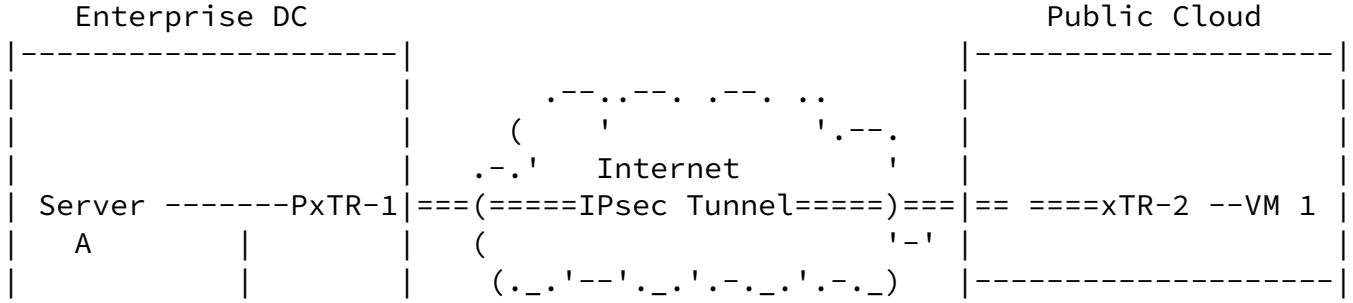

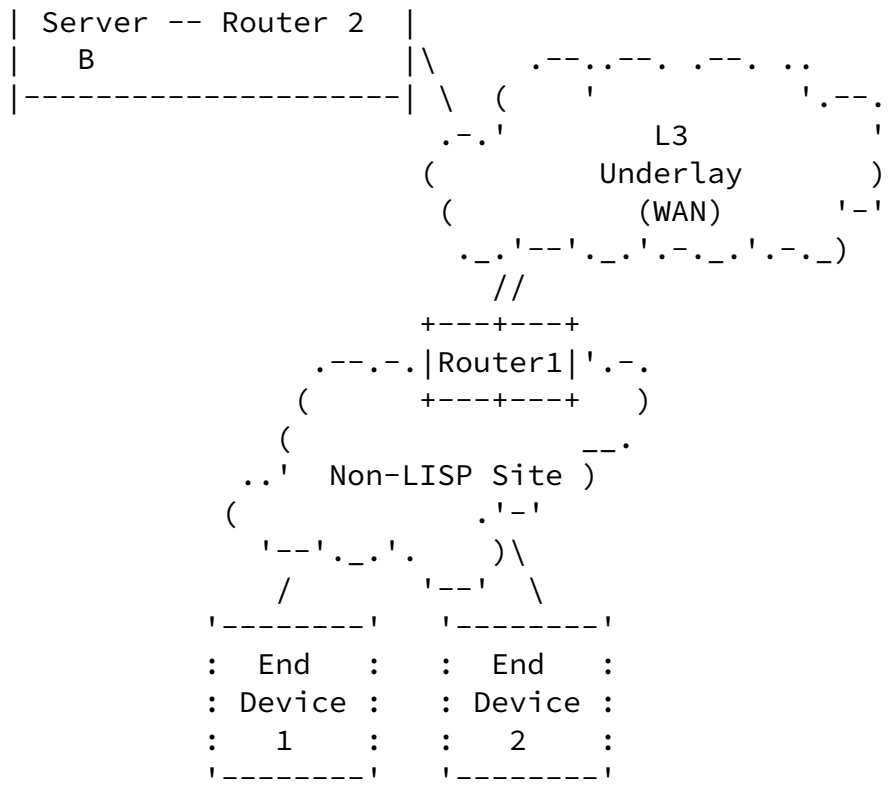

Figure 1

3. Terminology

Hybrid Cloud = A composition of two or more clouds (private, community or public) that remain unique entities but are bound together, offering the benefits of multiple deployment models.

LISP-enabled virtualized router = A virtual machine / appliance that supports Routing and LISP functions, including Host Mobility.

The terms Ingress Tunnel Router (ITR) and Egress Tunnel Router (ETR) are defined in [RFC6830]

The terms Proxy-ITR (PITR) and Proxy-ETR (PETR) are defined in [RFC6832]

4. Deploying LISP for Secure Hybrid Cloud Extension

As represented in Figure 1, LISP routers are deployed at both the enterprise DC (PxTR-1) and the public cloud (xTR-2).

At the enterprise DC, PxTR-1 does not need to be the default gateway for the local servers (physical and VMs), but it MUST be directly connected to the subnet where IP mobility will be provided. PxTR-1 MUST have an interface (physical or virtual sub-interface) connected to the same subnet of the servers that are eligible for moving. This allows the LISP router to detect the servers and to provide IP mobility for this subnet. PxTR-1 can detect server's EIDs by various way, including listening to ARPs that may be sent by the servers, for example during boot up time, or by initiating traffic (ICMP requests) to the servers. PxTR-1 MUST perform both proxy-itr and proxy-etr functions, so that non- LISP enabled sites can reach the servers moved to the cloud via the enterprise data center. Since PxTR-1 is not the default gateway and is not on the regular data path (i.e. the data path before there is any migration to the cloud), we refer to the deployment of LISP in the enterprise DC as non-intrusive. To redirect traffic from the enterprise data center to the cloud, the PxTR utilizes Proxy-ARP for both Intra-Subnet and Inter-Subnet communication.

The map-resolver (MR) and map-server (MS) functions can be enabled on either PxTR-1 in the enterprise DC, or xTR-2 running within the cloud. By enabling the MS/MR functions on one of the LISP routers used to provide the hybrid cloud extension, the solution can be deployed without adding other infrastructure at the cloud provider or at the enterprise.

Within the cloud, the LISP-enabled virtualized router (xTR-2) MUST be the default gateway for the Virtual Machines on those subnets that

require IP mobility. xTR-2 MUST be configured as a LISP ITR and ETR node so that it can perform LISP encapsulation and de-encapsulation of the packets coming from or going to the VMs located within the cloud. Whenever a route to the destination is not found on xTR-2 routing table, xTR-2 must route that traffic via PxTR-1 (at the enterprise DC). This function is useful to ensure that the traffic flow is symmetric between non-LISP enabled sites and the cloud, and MUST be used when there are firewalls or other stateful devices

located at the enterprise data center. xTR-2, being the default gateway on the cloud, can detect servers' EIDs (Endpoint IDentifiers) by listening to ARPs that may be sent by the server themselves. For example during boot up time, or whenever the host needs to communicate outside its subnet, because the host will ARP or send the packet towards its default gateway xTR-2. To support Intra-subnet communication between the cloud and the enterprise, xTR-2 attracts traffic local to the cloud using proxy-arp. Whenever a VM on the cloud ARPs for another IP located on the same subnet, the xTR will respond to this ARP request (proxy-arp) unless it has detected that the EID is local to the cloud .

5. Deployment Considerations

This section will cover what needs to be considered for a successful deployment of a LISP-Enabled Secure Hybrid Cloud.

5.1. Requirements on the underlying network

The network at the enterprise DC and the network within the cloud MUST allow the flooding of ARP packets within the subnets (i.e. VLAN, or similar Layer 2 technologies) where IP mobility is required. The reason for this requirement is that the xTR deployed within the cloud or the PxTR deployed within the enterprise DC uses Proxy-ARP to attract traffic to themselves to enable intra-subnet communication. For example, when a server within the enterprise DC wants to communicate with a server that is located within the cloud within the same subnet, the server within the enterprise DC will ARP for the IP address of the server that has moved to the cloud, in this case this ARP request MUST be flooded on the broadcast domain (i.e. VLAN) such that the LISP enabled router within the enterprise DC can respond to this request with its own MAC (proxy-ARP) so that traffic to a server moved to the cloud is then sent to the router located within the enterprise DC which then performs LISP encapsulation and send the traffic to the RLOC of the xTR located on the cloud. The same process happens on the reverse direction when traffic is returned from the cloud to the enterprise DC.

5.2. Routing Considerations

Deploying LISP for Secure Hybrid Cloud extension requires routing considerations related with: (1) RLOC advertisement, (2) advertisement of LISP enabled subnets to enable communication from non-LISP sites, and (3) traffic symmetry.

5.2.1. RLOC Advertisement

The RLOCs used on the LISP enabled routers at the enterprise DC and at the cloud MUST be reciprocally reachable via the IPsec tunnel that connects the enterprise DC to the cloud. Those RLOCs SHOULD belong to the private IP address space controlled by the enterprise. This keeps the LISP deployment independent from both the cloud and the enterprise infrastructures. Reachability information for those RLOCs SHOULD be announced via the dynamic routing protocol of choice (IGP or EGP) used on top of the IPsec tunnel connecting the enterprise to the cloud. In alternative, if preferred, static routing COULD be used as well.

5.2.2. LISP Stretched Subnets Advertisement

The subnets that are going to be stretched from the enterprise to the cloud already exist within the enterprise data center. Those subnets SHOULD already be advertised towards the enterprise WAN by the existing routing protocol. This ensures that non-LISP remote sites have a route to the LISP-enabled subnets via the enterprise data center, where PxTR-1 attracts all the traffic directed to the subnets that have been "stretched" to the cloud.

5.2.3. Traffic symmetry

For the cases where there are stateful devices (i.e. Firewalls, Load balancers) located within the enterprise data center, traffic symmetry is mandatory. To achieve that, as described in the LISP Stretched Subnets Advertisement section, traffic is first attracted to the enterprise and then tunneled to the cloud. On the return from the cloud, the iTR MUST ensures that the traffic towards non-LISP sites is first returned to the PeTR at the enterprise DC.

5.3. No changes to Virtual or Physical Servers

The network-based solution described in this draft, deployed with LISP routers at the enterprise DC and at the cloud, does not require changes to the servers (physical or virtual). Neither to those that are eligible for mobility, nor to those that are not eligible for mobility.

5.4. Integration with other network services

Although not the focus of this draft, it's important to consider that enterprises may have some network services, like firewalls or WAN acceleration devices that SHOULD be integrated as part of a holistic Hybrid Cloud solution. LISP SHOULD not prevent the integration of such services.

5.4.1. WAN acceleration

It may be desirable by an enterprise to accelerate traffic from the enterprise DC to the cloud. WAN acceleration SHOULD happen before the traffic is LISP and IPsec encapsulated. Defining all the options for how the WAN acceleration device is inserted within the traffic flow is outside the scope of this draft. However, one approach that may be supported by the routers used on the enterprise and on the cloud, is to have the LISP-enabled router redirect the traffic to the WAN acceleration device(s) before the traffic is LISP and IPsec encapsulated.

5.4.2. Firewalls

Within enterprise data centers, firewalls can be implemented at several points of the network. The section "Traffic Symmetry" covers how this solution works when there are firewalls located within the enterprise data center.

At the cloud, the LISP-enabled virtualized router (xTR-2 in Figure 1) is the default gateway for the servers VMs. In order to simplify the deployment xTR-2 SHOULD also function as a firewall for the servers located at the cloud. xTR-2 SHOULD secure communication between subnets located at the cloud, and communication from the cloud to the enterprise or the Internet.

6. Communication (flow) examples

This section will explain the detailed communication flows referring to the diagram shown in Figure 1

PxTR-1, at the enterprise LISP site, is placed "on a stick", meaning that it does not need to be the default gateway, and its interaction with the local infrastructure is based on Proxy-ARP. PxTR-1 proxy-replies on behalf of all non-local servers, inserting its own MAC address for any EID that is not locally detected. PxTR-1 can be either a physical router, or a virtual appliance. To be able to manage not only locally sourced traffic, but also traffic coming from

the WAN, PxTR -1 is enabled as a PiTR. To be able to receive back traffic from the cloud and deliver it to the WAN, PxTR-1 is also set-

up as a PeTR. In summary this node is a PxTR regarding its role for handling LISP traffic.

At the Cloud LISP site, xTR-2 is a standard xTR for the locally attached subnets. xTR-2 is the default gateway for the extended subnets. xTR-2 is playing the role of eTR for the flow coming from the enterprise site, and acts as an iTR for the flow going back to the enterprise site. For any destination that is not known by the xTR, the iTR encapsulates the traffic to the RLOC of the PeTR specified, in this case the PeTR specified is PxTR-1 deployed at the enterprise site.

6.1. LISP-enabled Intra subnet communication between the Enterprise and the Cloud

Server A at the enterprise site sends an ARP request for the IP address of VM-1 that it wants to communicate with, in order to find the MAC address. PxTR-1, which is on a stick, replies with its own MAC address (proxy-arp) as VM-1 is not detected locally. Traffic is then sent to PxTR-1, after which it will issue a Map-Request to the MS to find the location of VM-1. Finally it will encapsulate the traffic towards xTR-2 at the Cloud site as VM-1 is identified in the Map-server as belonging to that site. Traffic is then delivered towards the cloud-attached subnet.

On the return path, the flow is handled in a symmetrical reverse way compared to the inbound one just described above.

6.2. Communication from non-LISP Enabled Remote Sites to the Enterprise and to the Cloud

Traffic from End Device 1 or 2, which are located at a remote site that is not LISP enabled, is naturally attracted toward the enterprise DC by IP routing. Whenever reaching the Enterprise site, it is crossing the site's security and other services to reach the local subnet that is supposed to host the destination server VM-1. When the local default gateway sends an ARP to find VM-1, PxTR-1 responds to this ARP using the Proxy-ARP function as described above. Traffic is sent LISP-encapsulated to the Cloud site where it is

delivered to VM-1.

xTR-2, which is the default gateway for VM-1, handles the return traffic that is sent by VM-1. As this traffic is not intended to a LISP site, (i.e.it is targeted to End Device 1 or 2), xTR-2 sends the traffic to the PeTR configured on it (PxTR-1).

6.3. Communication from enterprise local LISP enabled subnet to the Cloud LISP enabled subnet

Traffic originated from a LISP enabled subnet (from Server B), intended to another LISP enabled subnet extended to the cloud (to VM-1) will first reach the local gateway (Router 2) and then will be routed locally to the extended subnet where PxTR-1 will respond to the ARP issued by Router 2. On the return path, the traffic will hit xTR-2, which is the default gateway, and then be routed by LISP towards the Enterprise site.

6.4. Communication from enterprise local non-LISP enabled subnet to the cloud LISP enabled subnet

In this case, traffic will first reach the default gateway (Router 2), from which it will follow the same path than the traffic originated from a remote site. The PxTR function will be used in both directions.

6.5. Communication from LISP Enabled Subnets to non-LISP Enabled Subnets

When traffic is sourced from a LISP enabled subnet at the enterprise site (Server A or B) towards a non-LISP enabled subnet at the cloud site, standard routing will take effect.

For the return traffic, xTR-2 will send it back to PxTR-1 as LISP encapsulated traffic.

6.6. Communication between non-LISP enabled subnets

In this case, the traffic will be routed using plain IP routing and

LISP is not involved.

The return traffic also uses plain IP routing, and LISP is not involved.

6.7. Inter-Subnet communication between servers in the Cloud

In the cloud itself, xTR-2 can locally route traffic between local subnets as it is the cloud site default gateway.

6.8. Communication from LISP Enabled Remote Sites to the Enterprise and to the Cloud

All above considerations of traffic flows are assuming that the only LISP enabled devices are PxTR-1 and xTR-2.

If one remote site need to access directly a non-LISP enabled resource in the cloud, meaning a subnet that is strictly local to the cloud, then pure routing can be used.

If a remote site needs path optimization to directly reach the servers that are part of a LISP stretched subnet at the cloud site, LISP can be enabled on this remote site. An xTR deployed on this remote site would consult the map-server to receive the correct location of the server.

7. Performance and Scalability Considerations

TBD.

8. Management and Automation Considerations

TBD.

9. Acknowledgements

The authors would like to thanks Michael Nolan and Fabio Maino for their review, and Fabio Maino for his encouragement to develop this draft.

10. IANA Considerations

This memo includes no request to IANA.

11. Security Considerations

The connection from the enterprise to the public cloud provider is usually done over the internet, although it may happen via a dedicated private circuit on some cases. This draft strongly suggests that an IPsec tunnel SHOULD be established between the enterprise and the cloud provider and that the data sent within this tunnel SHOULD be encrypted.

The IPsec tunnel SHOULD be established between the LISP-enabled routers located on the enterprise and on the cloud. By establishing the IPsec tunnels and ensuring that the RLOC from the routers are reachable via those IPsec tunnels then the LISP encapsulated traffic between the enterprise and the cloud will also be encrypted as it will flow over the IPsec tunnels.

Also see [I-D.ietf-lisp-sec] for a list of security considerations for LISP.

12. References

12.1. Normative References

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

12.2. Informative References

[I-D.ietf-lisp-sec] Maino, F., Ermagan, V., Cabellos-Aparicio, A., Saucez, D., and O. Bonaventure, "LISP-Security (LISP-SEC)", draft-ietf-lisp-sec-05 (work in progress), October 2013.

[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.

Authors' Addresses

Santiago Freitas
Cisco Systems
New Square Park, Bedfont Lakes
Feltham TW14 8HA
UK

Phone: +44 20 8824 8429
Email: safreita@cisco.com

Patrice Bellagamba
Cisco Systems
L'Atlantis, 11, Rue Camille Desmoulins
Issy Les Moulineaux 92782
France

Phone: +33 15 804 6235
Email: pbellaga@cisco.com

Yves Hertoghs
Cisco Systems
De Kleetlaan 6a
Diegem 1831
BE

Email: yhertogh@cisco.com