

Deploying a Secure Hybrid Cloud Extension with Cisco CSR 1000V and LISP

White Paper

Contents

Executive Summary	3
Introduction	3
Market Needs for Hybrid Cloud	5
Demand for Layer 3-Based Hybrid Cloud	5
Scope	6
Terminology	6
Solution Overview	7
Security Considerations	8
Reference Topology.....	8
Solution Considerations	9
Requirements on the Underlying Network	10
Deployment at Amazon Web Services (AWS)	10
Routing Considerations.....	11
LISP Stretched Subnets Advertisement.....	11
Traffic Symmetry.....	11
No Changes to Virtual or Physical Servers	11
Support for Multiple Hypervisors	11
Deploying LISP on Cisco CSR 1000V for Secure Hybrid Cloud Extension	12
Reference Topology.....	12
Implementation Details on the Enterprise Data Center	14
Implementation Details on the Public Cloud.....	15
Configuration Example	15
Enterprise Cisco CSR	15
Cloud-Based Cisco CSR.....	20
Integration with Virtual WAAS	24
Packet Walk and Encapsulation Stack	25
Conclusion	26
Appendix	27
Overview of Key Technologies	27
Cisco CSR 1000V	27
Locator/Identifier Separation Protocol.....	28
LISP Is a Routing Architecture	28
LISP Basic Concepts	29
AppNav and Wide Area Application Services	30
High Availability and Redundancy.....	30
Manageability: Automated Deployment with BDEO and REST API.....	31

Executive Summary

Traditionally, outsourcing services was about a provider fully dedicating a data center or part of a data center to a customer. The cloud is all about changing this “all or nothing” approach. With Software-as-a-Service (SaaS), providers can offer application outsourcing. With Infrastructure-as-a-Service (IaaS), providers can offer partial data center hosting. It is important to note that the success of such approaches is linked to how easily the cloud provider can be integrated with the customer’s existing resources; this relationship is critical to success.

Hybrid Cloud is about connecting the cloud provider to the enterprise private cloud. Two types of traffic must be considered with the hybrid cloud: the inter-subnet-routed traffic that is mostly a VPN transported type of traffic and the intra-subnet traffic. Intra-subnet traffic is an interesting new paradigm that is facilitating subnet continuity for cloud insertion, allowing providers to insert their infrastructure and application hosted on it right into the heart of the enterprise data centers. IP routing cannot really provide such a subnet extension. In the last few years, a new design approach has risen in the market, called Data Center Interconnect (DCI), where VLAN is extended over a long-distance network to allow subnets to be extended. The big question: Is this architecture realistic in a cloud environment where the two edge parts of the DCI would not belong to the same owner? Can the broadcast domain of an enterprise be extended to the service provider’s domain? In response to this question, Locator/ID Separation Protocol (LISP; RFC 6830) is certainly appealing with its ability to extend subnets without extending VLANs.

Several hybrid cloud use cases would benefit from subnet extensions using the LISP approach:

- Data center migrations, with no need to rework the workload IP address or the firewall rules, ensuring subnet continuity in a routed fashion with no broadcast domain extension
- SaaS with insertion of the provider application right in the heart of the enterprise data center
- IaaS with insertion of the compute overflow right in the heart of the enterprise data center
- Cloud bursting with virtual insertion of a virtual machine in the enterprise server farm, while the virtual machine is running on the provider site
- Backup services, for partial disaster recovery and disaster avoidance with the capability to clone the failing area and reactivate the application space, including storage, because it is in the provider space

It is worthwhile to note that the subnet extension is a very powerful tool for hybrid cloud versus VLAN extension. However, in 2014 we still cannot provide server cluster peering without VLAN extension because, for some clusters, the heartbeat messages are not routable.

Introduction

Most enterprises are expected to pursue a hybrid cloud computing strategy within the next 3 to 5 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 IP address, subnet mask, and default gateway configurations. Also, enterprises would like to adopt their own IP addressing scheme in the cloud and not be limited by the addressing scheme of the cloud provider infrastructure.

You can use LISP 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 (identity of server), subnet mask, and default gateway configurations as the ones used inside their own private data centers. LISP updates on its database the endpoint ID-to-router locator (EID-to-RLOC) mapping of the server to reflect the new location that, in this example, is moved to the cloud. No changes are required to the end systems, users, or servers, because LISP handles the mapping between identity (server IP address) and location (enterprise data center or public cloud), transparently to the users trying to reach the server.

LISP operates as an overlay, encapsulating the original packet from the server into a User Datagram Protocol (UDP) packet along with an additional outer IPv4 or IPv6 header, which holds the source and destination RLOCs. This encapsulation 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, providing total isolation of broadcast (Layer 2) domains between the enterprise and the public cloud.

Non-LISP-enabled sites communicate to the servers moved to the cloud through the enterprise data center, where LISP is deployed.

The solution documented on this paper does not require LISP to be enabled globally, but can be deployed by enabling LISP on just the enterprise data center and the public cloud, with minimal impact on the operations of both the data center and the cloud.

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

The LISP service is provided in the network by the Cisco Cloud Services Router 1000V (CSR 1000V) to any virtual machine, independent 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 virtual machine in the public cloud without the need to modify the virtual-machine configuration before migration.

LISP is deployed in the Cisco Cloud Services Router 1000V in the cloud, and can be deployed in either a Cisco CSR 1000V or a Cisco ASR 1000 Aggregation Services Router in the enterprise data center. The LISP-enabled CSR 1000V or ASR 1000 deployed within the enterprise data center does not need to be the default gateway for the local servers (physical and virtual machines).

The communication between the LISP-enabled CSR 1000V or ASR 1000 deployed within the enterprise data center and the CSR 1000V deployed within the public cloud is secured by an IP Security (IPsec) tunnel established between them. LISP encapsulated traffic is protected with the IPsec tunnel that provides data origin authentication, integrity protection, anti-reply protection, and confidentiality between the cloud and the enterprise. A routing protocol - this solution uses Open Shortest Path First (OSPF) - is also enabled over the IPsec tunnel to advertise the LISP RLOC (location) of the CSRs.

This LISP-enabled hybrid cloud solution allows intra-subnet communication regardless of the location of the server, meaning that communication between two servers located on the same subnet can happen even when one server is located at the enterprise data center and another server is located at the cloud. Inter-subnet communication is also supported.

Another important point addressed by the solution documented herein is that when servers move to the cloud, users still expect the performance of the application to be the same or better than it was when the server was running on the enterprise data center. With the servers remotely located on the cloud, the need to accelerate applications and to optimize the WAN bandwidth between the enterprise and the cloud is evident.

As part of this solution, Cisco Virtual Wide Area Application Services (vWAAS) is integrated into the hybrid cloud solution to provide application acceleration and WAN bandwidth compression for the traffic between the enterprise and the cloud.

Market Needs for Hybrid Cloud

About half of United States and European enterprise IT decision makers have reported that their companies use cloud Infrastructure-as-a-Service (IaaS). IT decision makers reported the greatest interest in using IaaS in a hybrid cloud approach to complement on-premises capacity.

[Research by Forrester Consulting](#) shows that the hybrid model of implementing IaaS is particularly enticing for United States and European enterprises. Twenty-eight percent of hardware decision makers report using a hybrid cloud strategy for server resources today, and nearly half predict that they will have workloads in both traditional or on-premises and hosted or service provider cloud environments by 2016.

Demand for Layer 3-Based Hybrid Cloud

One of the requirements for a hybrid cloud solution is that it must allow the servers to be moved to the cloud without changing their IP address.

One option to achieve this requirement is to extend Layer 2 between the enterprise and cloud, but the requirement for a Layer 3-based solution that provides IP mobility is growing because enterprises and cloud providers alike would prefer not to extend the broadcast (Layer 2) domain between their environments.

The LISP on Cisco CSR 1000V solution is a Layer 3-based approach for hybrid cloud that provides IP mobility; it has the following characteristics:

- It allows a routed (Layer 3) connection between the enterprise and the cloud.
- It meets the requirements of customers who mandate a Layer 3 connection for long distance.
- It provides total isolation of broadcast (Layer 2) domains between the enterprise and the cloud.
- It natively provides a Layer 3 gateway in the cloud for optimal (localized) routing between servers moved to the cloud.
 - No hairpinning saves intercloud bandwidth.
- It supports hundreds of virtual machines in the cloud, depending on CSR 1000V throughput.
- It works with any standard virtual machine in the cloud; you do not need to modify the virtual machine before migration.
- It allows you to move virtual machines back from the cloud to the enterprise.
- It allows direct access to the virtual machines in the cloud from the user's site if LISP is deployed at the branch offices.

Scope

This paper documents a solution that allows a secure hybrid cloud extension based on LISP running on the Cisco CSR 1000V, and optionally uses Cisco vWAAS to accelerate the connection between the enterprise and the cloud.

This solution requires the Cisco CSR 1000V running Cisco IOS[®] XE Software Release 3.11 or later with the premium technology package activated.

For the vWAAS, the solution was validated with Cisco Wide Area Application Services (WAAS) Software Release 5.3.3 or later.

Terminology

Note: More detailed terminology is discussed later in the document.

CSR 1000V: Cisco's virtual router offering that you can deploy in private, public, or hybrid cloud environments.

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.

Locator/ID Separation Protocol (LISP): A tunnel protocol that uses a central database in which endpoint location is registered. LISP enables an IP host-based mobility of endpoints.

LISP-enabled virtualized router: A virtual machine or appliance that supports routing and LISP functions, including host mobility.

Endpoint ID (EID): IPv4 or IPv6 identifier of the devices connected at the edge of the network. Used in the first (most inner) LISP header of a packet.

Routing locator (RLOC): IPv4 or IPv6 addresses used to encapsulate and transport the flow between LISP nodes.

Ingress tunnel router (ITR): A router that has two functions: it resolves the location of an EID by querying the database, and then it performs the encapsulation toward the remote RLOC.

Egress tunnel router (ETR): A router that has two functions: it registers the endpoint or location associated with the database, and then it decapsulates the LISP packet and forwards it to the right endpoint.

xTR: A generic appellation for a device performing both ITR and ETR functions.

Proxy-ITR (PITR): Acts like an ITR but does so on behalf of non-LISP sites that send packets to destinations at LISP sites.

Proxy-ETR (PETR): Acts like an ETR but does so on behalf of LISP sites that send packets to destinations at non-LISP sites.

PxTR: A generic appellation for a device that performs both PITR and PETR functions.

Solution Overview

The concept of the solution is to use two Cisco CSR 1000V virtual routers, one deployed in the enterprise data center and one in the cloud with a LISP encrypted tunnel in between. LISP provides the mobility between the two sites, allowing you to stretch subnets from the enterprise to the virtual private cloud.

Non-LISP-enabled sites communicate to the servers moved to the cloud through the enterprise data center, where LISP is deployed.

The solution proposed in this paper does not require LISP to be enabled globally; it can be deployed by enabling LISP on just the enterprise data center and the public cloud, with minimal impact on the operations of both the data center and the cloud.

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

The LISP service is provided in the network, to any virtual machine, independent 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 virtual machine in the public cloud without the need to modify the virtual-machine configuration before migration.

LISP is deployed in Cisco CSR 1000V virtual routers in the cloud and in the enterprise data center. Potentially the same solution could be used with physical Cisco ASR 1000 Series Aggregation Services Routers, but this paper focuses on the usage of virtual devices.

The LISP-enabled router deployed within the enterprise data center does not need to be the default gateway for the local servers (physical and virtual machines). The CSR is deployed "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. The CSR has at least two interfaces facing the data center, one as the routed core interface and one for Layer 2 collection of VLANs. The stretched subnets are locally either connected using sub-interfaces or, if they are not numerous, attached to CSR interfaces.

The LISP CSR in the cloud is deployed with one interface facing the Internet or the Multiprotocol Label Switching (MPLS) network and the other interfaces facing the local cloud VLANs. The stretched subnets are hosted by the interfaces that face the local cloud VLANs. If the cloud infrastructure supports dot1Q VLAN, then only one interface with sub-interfaces could be used. This condition is rare, however, and most probably the CSR will have one interface per subnet, limiting the number of subnets to be extended depending on the hypervisor used in the cloud. Facing the cloud, the attachment described in this document uses VLANs, however, since the CSR supports Virtual Extensible LAN (VXLAN) to encapsulate layer 2 ethernet frames, this technology can be used natively.

The LISP-enabled hybrid cloud solution allows intra-subnet communication regardless of the location of the server, meaning that communication between two servers located on the same subnet can happen even when one server is located at the enterprise data center and another server is located at the cloud. Inter-subnet communication is also supported.

Security Considerations

To preserve privacy, the communication between the LISP-enabled router deployed within the enterprise data center and the router within the public cloud must be secured by encryption.

Two approaches are possible. First, the LISP tunnel could be encrypted using IPsec in transport mode, where no additional header is added to LISP. The problem with this efficient approach is the support of Network Address Translation (NAT), and NAT is used in the cloud. So the second approach is analyzed in this document, and NAT uses an IPsec tunnel mode. 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, with, in addition, support for NAT.

In this paper the RLOCs for encapsulation are loopbacks, so it is necessary to enable a routing protocol to announce them over the IPsec tunnel.

Reference Topology

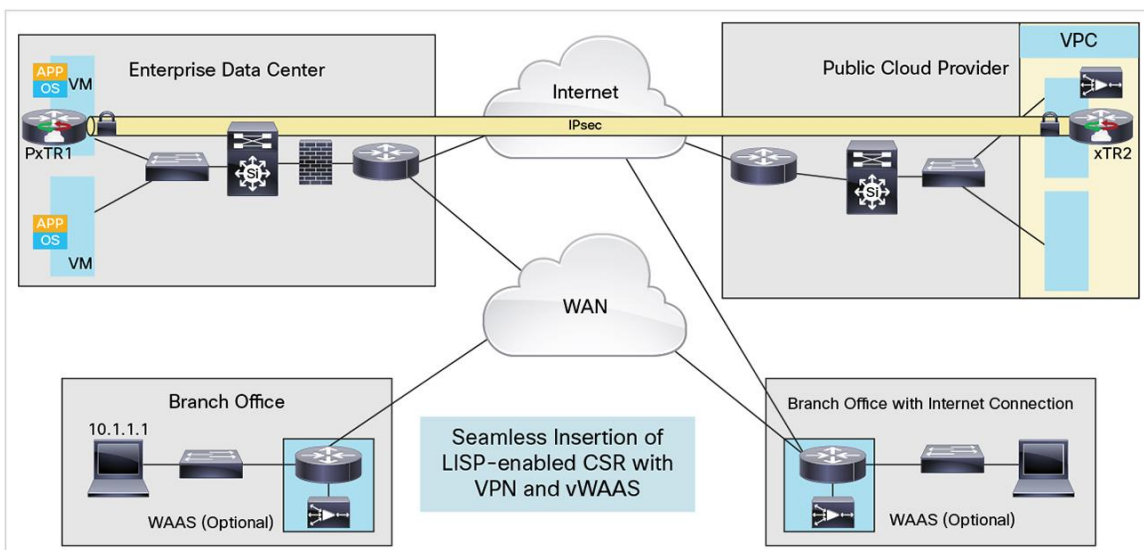
Figure 1 illustrates the use case described in this paper. Users' end devices, typically located at non-LISP sites, are connected to the enterprise WAN to access servers located either at the enterprise data center or at the public cloud.

The enterprise data center hosts some of the servers, and has a PxTR attached to the subnets that host mobile servers.

The PxTR deployed within the enterprise data center (PxTR-1) is configured with an IPsec tunnel toward the cloud 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 data center and the public cloud.

Figure 1. Enabling a Secure Hybrid Cloud with LISP on Cisco CSR 1000V



Solution Considerations

LISP routers are deployed at both the enterprise data center and the public cloud. At the enterprise data center, PxTR-1 does not need to be the default gateway for the local servers (physical and virtual machines), 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 to allow the LISP router to detect the servers and provide IP mobility for this subnet.

PxTR-1 can detect the server EIDs by various ways, including by listening to Address Resolution Protocol instances (ARPs) that may be sent by the servers, for example during bootup time, or by initiating traffic (Internet Control Message Protocol [ICMP] requests) to the servers. In order to ensure that the CSR on the enterprise site will always detect any host that is present, a liveness script is proposed in the configuration. In the future, the LISP code itself will force host detection, but until this final implementation, it is best to have a script to poll the different subnets under mobility right after the CSR boots. 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 through the enterprise data center. Because PxTR-1 is not the default gateway and is not on the regular data path (that is, the data path before there is any migration to the cloud), we refer to the deployment of LISP in the enterprise data center as nonintrusive. To redirect traffic from the enterprise data center to the cloud, the PxTR uses 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 data center or xTR-2 running within the cloud. If you enable the map-resolver and map-server functions on one of the LISP routers used to provide the hybrid cloud extension, you can deploy the solution without adding other infrastructure at the cloud provider or at the enterprise.

Within the cloud, the LISP-enabled CSR (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 virtual machines located within the cloud. Whenever a route to the destination is not found on the xTR-2 routing table, xTR-2 must route that traffic through PxTR-1 (at the enterprise data center). This function is useful to ensure that the traffic flow is symmetric between non-LISP-enabled sites and the cloud, and it must be used when firewalls or other stateful devices are located at the enterprise data center.

Being the default gateway on the cloud, xTR-2 can detect EIDs of servers by listening to ARP requests that are sent by the servers themselves, for example during bootup time or whenever the host needs to communicate outside its subnet, because the host will perform ARP on the packet or send it toward 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 virtual machine on the cloud performs ARP requests 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.

In the cloud, an important consideration is the placement of any firewalling service. This service clearly depends on the traffic pattern. In this solution, the main use case is having the session go first to the enterprise site, and then cross the IPsec tunnel towards the cloud using the LISP redirection of traffic. In the main case, the one analysed in this paper, the flow will have crossed the WAN edge firewall and the traffic is sent back to the enterprise to cross the firewall statefully, so a cloud firewall is unnecessary. But a cloud firewall must be considered in three cases. One is for the traffic coming directly from a branch office to the cloud using the LISP path optimization, another is the deployment of multitier applications in the cloud, and finally it is when the default gateway has been determined to be hosted by a cloud virtual firewall.

The option to support such a design is to have the CSR generate a default route that the firewall uses, pointing all egress traffic to the CSR after executing the security policy. For the ingress traffic the CSR will use Policy-Based Routing (PBR) to force the traffic to reach the firewall. In that case the CSR becomes a second hop, and is no longer able to receive ARP requests that the firewall will consume. In order to maintain a capability of detection, the CSR will have, in addition to its link to the firewall, an additional link connected straight to the VLAN where the virtual machine resides. The CSR has two links toward the cloud, one in Layer 3 and one in Layer 2. The CSR still needs a link toward the core routing to reach the outside.

Requirements on the Underlying Network

The transport service in the WAN side is usually of two types. It could be a plain public Internet network, or it could be a private MPLS VPN network. Both transport services are supported.

With the Internet, an important consideration is the NAT service. NAT is usually used between the cloud and the transport network, and most probably the same occurs on the enterprise side. Because IPsec supports NAT, the fact that this solution uses IPsec as a tunnel connection between the enterprise and the cloud solves any NAT burden. In the IPsec tunnel, as well as in the cloud, the private enterprise addressing is used, and LISP operates over these private addresses.

With the MPLS VPN transport, because all addresses are in a private space, no NAT considerations are necessary. With private MPLS VPN the only important consideration is that you may need to access a licence server that may not be on the private network, but on the public one.

Deployment at Amazon Web Services (AWS)

When the solution described on this whitepaper is deployed in a Virtual Private Cloud (VPC) at Amazon Web Services (AWS) the subnet mask of the Virtual Machines deployed at AWS must be changed to 255.255.255.255 (/32). This change is performed on the operating system of the VM, no change is required on AWS portal. The subnet created on AWS portal can still be a /24.

The reason for this change is that Amazon AWS does not flood ARP request from the VMs. When a VM located in AWS needs to communicate with other VMs part of the same subnet but the VMs are located on the enterprise, the VM located at AWS would ARP for the IP address of the remote VMs. Because AWS doesn't forward those ARP requests, the CSR deployed on AWS cannot proxy-ARP for those requests and traffic is dropped. To overcome this limitation, the only change required is that the VM deployed on AWS should have their subnet mask configured to 255.255.255.255 (/32). With this change the VMs located at AWS are able to communicate with VMs located in the enterprise even within the same subnet (intra-subnet traffic).

Routing Considerations

In any case, this solution requires only that each CSR have one unique outside address; then an IPsec tunnel is built between these addresses, and the IPsec tunnel has a private address.

In the hybrid cloud solution, an OSPF peering is enabled over this IPsec tunnel. The objective is to announce to the other CSR the loopback address that is used as the LISP location (RLOC). A simple static route could have done the same job. Now, enabling a routing protocol could allow you to advertise some subnet of the cloud that does not need mobility and could not be subject to LISP handling.

LISP Stretched Subnets Advertisement

This solution is a variant of the LISP across subnet mode (ASM), but is not fully compliant with it. It is called LISP stretched subnet mode (SSM). The difference with LISP ASM is that the same subnet is declared on both sides of the network, and the mobility is run between these two segments of the same subnet. It is not LISP extended subnet mode (ESM) either, because there is no Layer 2 extension technique at all in this solution.

LISP SSM stretches a subnet naturally, and the same subnet is present at the same time in the enterprise and the cloud.

The LISP map server will see a /24 or so associated with this segment on both sites, but only the host-based routing will effectively be used in LISP to identify where a host is (whether it is in the enterprise or in the cloud).

The subnets that will be stretched from the enterprise to the cloud already exist within the enterprise data center. Those subnets should already be advertised toward the enterprise WAN by the existing routing protocol to ensure that non-LISP remote sites have a route to the LISP-enabled subnets through the enterprise data center, where PxTR-1 attracts all the traffic directed to the subnets that have been "stretched" to the cloud.

Traffic Symmetry

For the cases where stateful devices (that is, firewalls and load balancers) are located within the enterprise data center, traffic symmetry is mandatory. To achieve traffic symmetry when the traffic is first attracted to the enterprise and then tunneled to the cloud, the iTR must ensure that the traffic on the return from the cloud toward non-LISP sites is first returned to the PeTR at the enterprise data center. To execute such a traffic path, a function called "use-petr" is applied on the xTR in the cloud, and all the traffic that is subject to default routing is encapsulated and forced back to the enterprise.

No Changes to Virtual or Physical Servers

Because the solution uses Proxy-ARP for traffic interception and VLAN or VXLAN as the encapsulation method, it can work with any virtual machine. You don't need to modify the server driver (physical or virtual) or the servers, whether they are eligible for mobility or not.

Support for Multiple Hypervisors

This solution requires the flexibility of supporting any hypervisor. The hypervisor type used on the enterprise network may be different from the hypervisor used on the cloud infrastructure. Thus, we choose the Cisco CSR 1000V, which is deployed as a virtual machine on top of a hypervisor. It is server-, virtual switch-, and hypervisor-agnostic. The Cisco CSR 1000V can run with any virtual switch such as the Cisco Nexus 1000V, vSwitch, or dVS. The Cisco CSR 1000V supports all hypervisors in the market today, such as VMware ESXi, Citrix XenServer, Red Hat KVM, and Microsoft Hyper-V. Table 1 lists the supported hypervisor versions.

Table 1. Support Matrix for Hypervisor Versions

Cisco CSR 1000V with Cisco IOS Software Release XE 3.12S	
ESXI	<ul style="list-style-type: none"> • 5.0 • 5.1 • 5.5
KVM	<ul style="list-style-type: none"> • Linux KVM based on Red Hat Enterprise Linux 6.31 • Red Hat Enterprise Virtualization 3.1 • Ubuntu 12.04.03 LTS Server 64 Bits
XEN	<ul style="list-style-type: none"> • 6.1
Hyper-V	<ul style="list-style-type: none"> • Windows Server 2012 R2

The minimum footprint required to run the Cisco CSR 1000V is 1 virtual CPU (vCPU) and 2.5-GB memory.

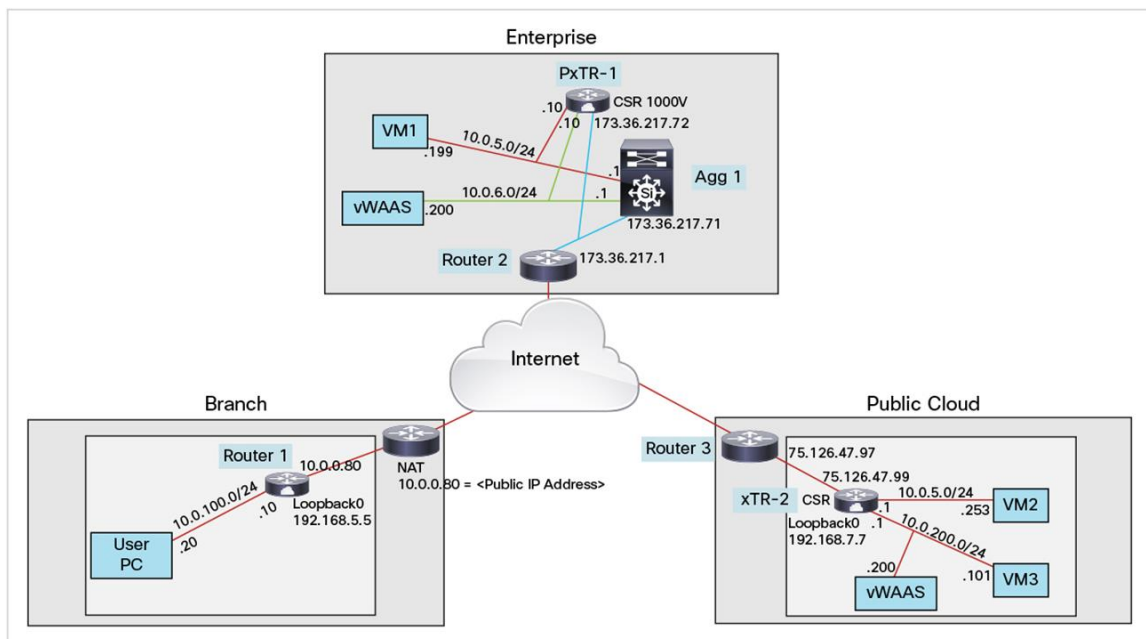
Deploying LISP on Cisco CSR 1000V for Secure Hybrid Cloud Extension

This section covers implementation details for the solution.

Reference Topology

Figure 2 shows the reference topology.

Figure 2. Secure Hybrid Cloud Extension Using LISP and Cisco CSR 1000V



The reference topology has three environments: An enterprise branch office (remote site), an enterprise data center, and the public cloud.

The branch-office subnet where users connect is 10.0.100/24. Users receive an IP address within this range through Dynamic Host Configuration Protocol (DHCP). The switches and wireless access points that certainly exist within a branch office are not represented on the diagram for simplicity. The users' default gateway is the router (Router 1) with IP address 10.0.100.10. Router 1, usually a Cisco Integrated Services Router (ISR), is connected to the Internet, in this case through a provider router that performs NAT to translate from the private IP address

10.0.0.80 to a public address. Router 1 does not run LISP, so the branch office is a non-LISP site. Branch-office users access servers located either at the enterprise data center or at the public cloud.

The enterprise data center is connected to the Internet through the data center WAN edge router (Router 2). Router 2 has the IP address 173.36.217.1 on its interface facing the enterprise data center and has a public address not shown on the topology on its connection to the Internet. The switch represented on the topology (Agg 1) is the aggregation layer switch within the enterprise data center. Agg 1 has one interface connected to Router 2 with IP address 173.36.217.71 and has an interface with IP address 10.0.5.1 on the subnet 10.0.5.0/24. Agg 1 is the default gateway for all the servers located in the enterprise data center; on the diagram 10.0.5.0/24 is the only subnet represented for simplicity. Agg 1 also has an interface on subnet 10.0.6.0/24, where the virtual WAAS appliance exists. The Cisco CSR 1000V installed on the enterprise, PxTR-1, has three interfaces. The blue link is the interface with IP address 173.36.217.72 and it is the PxTR-1 connection to the Internet. The green link is an interface with the IP address 10.0.6.10 located on the same segment where the vWAAS appliance exists. The red link is the PxTR-1 connection to the 10.0.5.0/24 subnet where mobility is required.

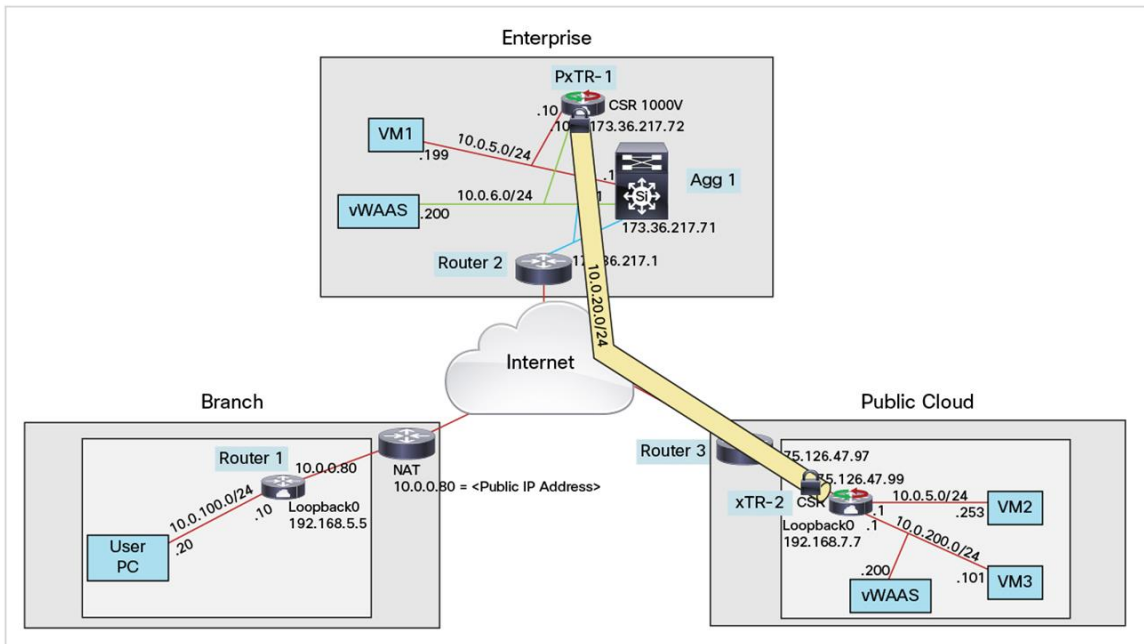
On the enterprise data center is a virtual machine with IP address 10.0.5.199/24 that is configured to use Agg 1 (IP 10.0.5.1) as its default gateway.

PxTR-1 is the only device within the enterprise data center running LISP, and it is configured as a LISP PxTR. PxTR-1 is also configured with an IPsec tunnel toward xTR-2, the CSR in the cloud. OSPF is enabled over the IPsec tunnel to advertise the IP address of the loopback interfaces of PxTR-1 and xTR-2 that are used as the LISP RLOCs.

A Cisco CSR 1000V is also deployed in the cloud. Router 3, with IP address 75.126. 47.97, is part of the cloud provider infrastructure and is not controlled by the enterprise; no changes are required for Router 3. The xTR-2 (CSR) deployed on the cloud has one public IP address, 75.126. 47.99, and two interface interfaces with private addresses, one interface configured with IP 10.0.5.1 and the other with 10.0.200.1. The IP 10.0.5.1 is the same as the IP used on Agg 1 on the enterprise data center so that when a server is moved to the cloud no change is required on the server configuration. A virtual machine with IP address 10.0.5.253 has been moved to the cloud; it uses the IP address 10.0.5.1 (xTR-2) as its default gateway.

xTR-2 is the only device in the cloud running LISP, and it is configured as a LISP xTR. xTR-2 is also configured with an IPsec tunnel toward PxTR-1, the enterprise data center CSR, and it has OSPF enabled over its IPsec tunnel to advertise its loopback interface address that is used as its RLOC. In Figure 3, the subnet 10.0.5.0/24 is stretched between the enterprise data center and the public cloud. The IPsec tunnel between the enterprise and cloud CSRs is used to secure communication between the enterprise and the cloud. OSPF is enabled over this tunnel to advertise the loopback used as RLOC for each CSR.

Figure 3. IPsec Tunnel between Enterprise and CSR1000V in the Cloud



Implementation Details on the Enterprise Data Center

At the enterprise data center, PxrTR-1 does not need to be the default gateway for the local servers (physical and virtual machines), but it must be directly connected to the subnet where IP mobility will be provided, in this case 10.0.5.0/24. PxrTR-1 must have an interface (physical or virtual sub-interface) connected to the same subnet of the servers that are eligible for moving to the cloud. This setup allows the PxrTR-1 (LISP router) to detect the servers and to provide IP mobility for this subnet.

PxrTR-1 can detect server EIDs by various ways, including by listening to ARP requests that may be sent by the servers, for example during bootup time, or by initiating traffic (ICMP requests) to the servers. PxrTR-1 must perform both LISP proxy-itr and proxy-etr functions, so that non-LISP-enabled sites can reach the servers moved to the cloud through the enterprise data center. Because PxrTR-1 is not the default gateway and is not on the regular data path (that is, the data path before there is any migration to the cloud), the deployment of this solution in the enterprise data center is nonintrusive. To redirect traffic from the enterprise data center to the cloud, the PxrTR-1 uses Proxy-ARP for both intra-subnet and inter-subnet communication.

PxrTR-1 is also configured with an IPsec tunnel toward xTR-2; OSPF is enabled on this tunnel to advertise the loopback address used as LISP RLOC by PxrTR-1.

The LISP map-resolver and map-server functions can be enabled on either PxrTR-1 in the enterprise data center or xTR-2 running within the cloud. For this configuration, PxrTR-1 is the map resolver and map server. If you enable the map-resolver and map-server functions on one of the LISP routers used to provide the hybrid cloud extension, you can deploy the solution without adding other infrastructure at the cloud provider or at the enterprise sites.

Implementation Details on the Public Cloud

Within the cloud, the LISP-enabled Cisco CSR 1000V (xTR-2) is the default gateway for the virtual machines on those subnets that require IP mobility. xTR-2 is 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 virtual machines located within the cloud. For traffic leaving the cloud, whenever a route to the destination is not found on the xTR-2 routing table, xTR-2 must route that traffic through PxTR-1 at the enterprise data center. This function, known as use-petr, is useful to ensure that the traffic flow is symmetric between non-LISP-enabled sites and the cloud, and it must be used when firewalls or other stateful devices are located at the enterprise data center.

Being the default gateway on the cloud, xTR-2 can detect EIDs of the servers by listening to ARP requests sent by the servers, for example during bootup time or whenever the host needs to communicate outside its subnet, because the host will perform ARP or send the packet toward 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 virtual machine on the cloud ARP requests another IP address 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. xTR-2 uses PxTR-1 as its map resolver and map server.

xTR-2 is also configured with an IPsec tunnel toward PxTR-1; OSPF is enabled on this tunnel to advertise the loopback address used as LISP RLOC by the xTR-2.

Configuration Example

This section provides the full configuration for the Cisco CSR 1000V routers located at the enterprise and in the cloud and highlights the relevant configurations.

Enterprise Cisco CSR

```
CSR-ENTERPRISE#sh run
Building configuration...

Current configuration: 6340 bytes
!
! Last configuration change at 17:22:20 UTC Fri Apr 11 2014 by cisco
!
version 15.4
service timestamps debug datetime msec
service timestamps log datetime msec
service call-home
no platform punt-keepalive disable-kernel-core
platform console virtual
!
hostname CSR-ENTERPRISE
!
boot-start-marker
boot system flash bootflash:csr1000v-universalk9.03.11.00.S.154-1.S-std.SPA.bin
boot-end-marker
!
!
```

```

enable secret <PASSWORD>
!
aaa new-model
!
!
aaa authentication login default local
!
aaa session-id common
!
subscriber templating
!
multilink bundle-name authenticated
!
crypto pki trustpoint TP-self-signed-3974101357
  enrollment selfsigned
  subject-name cn=IOS-Self-Signed-Certificate-3974101357
  revocation-check none
  rsakeypair TP-self-signed-3974101357
!
!
crypto pki certificate chain TP-self-signed-3974101357
  certificate self-signed 01
    3082022B 30820194 A0030201 02020101 300D0609 2A864886 F70D0101 05050030
    31312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
    69666963 6174652D 33393734 31303133 3537301E 170D3133 31313235 31373131
    31375A17 0D323030 31303130 30303030 305A3031 312F302D 06035504 03132649
    4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D33 39373431
    30313335 3730819F 300D0609 2A864886 F70D0101 01050003 818D0030 81890281
    8100B03C 40129DFA 4D301B8E B084C8E8 D34A147B D4F4B2BC BD8A1315 720E83BE
    02BF9CA9 C91D980D BD384DEA F1985F20 78C32225 846375FF C9E1F069 1E6F6943
    6483CA6D A0EC1279 83453D87 75D76B93 B6B2069C 7DCB1D26 8A90C589 C27CEEA3
    FFAF4F25 2E08E724 B02BC66F 03C34150 B7E48D79 1C951616 03678AAF 03260F36
    5C850203 010001A3 53305130 0F060355 1D130101 FF040530 030101FF 301F0603
    551D2304 18301680 146724A7 4726D601 C0AABD5D 1163F7CA 8F75307B D1301D06
    03551D0E 04160414 6724A747 26D601C0 AABD5D11 63F7CA8F 75307BD1 300D0609
    2A864886 F70D0101 05050003 81810088 940D6E45 9C8BB146 6EE006C6 EBFACFFF
    98B96120 7CA20596 38842C7F 4092A149 2CEF3E24 E87CB4C8 D4CA6EA9 DF386BE3
    90D80BE7 61DCC710 DD46FC69 E07BF09D 6B04F83F 4F61736A B1960982 4084F927
    F3DE6A0C 471249AC 3909184C D5F69EE8 1B97A942 DBEC3DCF F6D6B055 04ED0F5F
    A60A9B79 53CB6C5B 3DACA195 D45622
  quit
license udi pid CSR1000V sn <SERIAL NUMBER>
license accept end user agreement
license boot level premium
spanning-tree extend system-id

```

Activate Premium feature set.


```

!
username <USERNAME> privilege 15 secret <PASSWORD>
!
redundancy
 mode none
!
class-map type appnav match-all sl-waas
 match access-group name sl-waas
!
policy-map type appnav sl-waas
 class sl-waas
  distribute service-node-group sl-waas
!
!
crypto isakmp policy 1
 encr aes
 authentication pre-share
 group 14
crypto isakmp key <KEY VALUE> address <PUBLIC ADDRESS OF CLOUD CSR>
!
crypto ipsec transform-set VPN esp-aes esp-sha-hmac
 mode tunnel
!
crypto ipsec profile VPN-Profile
 set transform-set VPN
!
!
service-insertion service-node-group sl-waas
 service-node 10.0.6.200
!
service-insertion appnav-controller-group sl-waas
 appnav-controller 10.0.6.10
!
service-insertion service-context waas/1
 appnav-controller-group sl-waas
 service-node-group sl-waas
 service-policy sl-waas
 vrf default
 enable
!
!
interface Loopback0
 ip address 192.168.6.6 255.255.255.255
 ip ospf 1 area 0
!

```

AppNav configuration used to redirect traffic toward vWAAS node.
Traffic indicated by Access List "sl-waas" is redirected to service group node "sl-waas".

IPsec VPN configuration.

Define the AppNav service node group "sl-waas" and specify the IP address of the vWAAS node.

Define the AppNav controller group "sl-waas" and specify the IP address of the appnav controller. In this case the CSR itself is the controller, so it points to an IP address locally configured on the Cisco CSR.

Enable service redirection to vWAAS node previously defined.

Loopback used as LISP RLOC, with OSPF enabled on it to advertise it to OSPF neighbors.

```
interface Tunnel2
```

```
description VPN TO PUBLIC CLOUD
ip address 10.0.20.2 255.255.255.0
ip ospf network point-to-point
ip ospf 1 area 0
load-interval 30
tunnel source GigabitEthernet1
tunnel mode ipsec ipv4
tunnel destination <PUBLIC ADDRESS OF CLOUD CSR>
tunnel protection ipsec profile VPN-Profile
```

Tunnel interface used to connect to cloud Cisco CSR. OSPF and IPsec are enabled over it. Source interface is the local interface with connectivity to the cloud.

```
!
```

```
interface LISP0
```

```
!
```

```
interface GigabitEthernet1
```

```
ip address <PUBLIC IP ADDRESS AND MASK>
negotiation auto
```

Interface with connectivity to the cloud. It could be any of the interfaces on the Cisco CSR.

```
!
```

```
interface GigabitEthernet2
```

```
ip address 10.0.5.10 255.255.255.0
negotiation auto
```

```
lisp mobility LISP1
```

```
service-insertion waas
```

Interface connected to the subnet where mobility is required. LISP mobility is enabled on the interface. WAAS can also be enabled.

```
!
```

```
interface GigabitEthernet3
```

```
ip address 10.0.6.10 255.255.255.0
negotiation auto
```

Interface with connectivity to the vWAAS node.

```
!
```

```
interface AppNav-Compress1
```

```
ip unnumbered GigabitEthernet3
no keepalive
```

```
!
```

```
interface AppNav-UnCompress1
```

```
ip unnumbered GigabitEthernet3
no keepalive
```

```
!
```

```
router lisp
```

```
locator-set DC1
```

```
192.168.6.6 priority 1 weight 100
```

```
exit
```

LISP enabled and IP 192.168.6.6 selected as RLOC.

```
!
```

```
eid-table default instance-id 0
```

```
dynamic-eid LISP1
```

```
database-mapping 10.0.5.0/24 locator-set DC1
```

```
map-notify-group 239.0.0.1
```

```
exit
```

LISP mobility enabled for subnet 10.0.5.0/24. LISP1 name is used under interface config to associate this config to the interface.

```

!
  ipv4 itr map-resolver 192.168.6.6
no ipv4 itr
  ipv4 etr map-server 192.168.6.6 key cisco
  ipv4 etr
exit
!
  site DATA_CENTER
  authentication-key cisco
  eid-prefix 10.0.5.0/24 accept-more-specifics
  eid-prefix 10.0.100.0/24
  exit
!
  ipv4 map-server
  ipv4 map-resolver
  ipv4 proxy-etr
  ipv4 proxy-itr 192.168.6.6
  ipv4 itr map-resolver 192.168.6.6
  ipv4 etr map-server 192.168.6.6 key cisco
  exit
!
router ospf 1
  router-id 192.168.6.6
!
ip forward-protocol nd
!
no ip http server
ip http secure-server
ip route 0.0.0.0 0.0.0.0 GigabitEthernet1 <IP ADDRESS TO REACH THE INTERNET>
!
  ip access-list extended sl-waas
  permit tcp any any
!
control-plane
!
!
line con 0
  stopbits 1
line aux 0
  stopbits 1
line vty 0 4
  privilege level 15
!
!
end

```

Map resolver and map server configured and pointing to local address because enterprise Cisco CSR is the map resolver and map server.

Enterprise Cisco CSR configured as map server with an authentication key. Map server configured to accept specific (/32) registrations for subnet 10.0.5.0/24, where mobility is enabled. Any subnet where mobility is enabled should be listed.

Enterprise Cisco CSR configured as map server and map resolver. Also configured as Proxy-ETR and Proxy-ITR.

OSPF enabled.

Access list used for classifying traffic that should be redirected to vWAAS node.

Cloud-Based Cisco CSR

The configuration of the Cisco CSR 1000V located on the cloud is very similar to the configuration of the Enterprise CSR; therefore, only the differences are highlighted, although the full configuration is provided.

```
CSR-PUBLIC-CLOUD#sh run
Building configuration...

Current configuration: 6417 bytes
!
! Last configuration change at 09:26:55 UTC Tue Apr 8 2014 by cisco
!
version 15.4
service timestamps debug datetime msec
service timestamps log datetime msec
service password-encryption
no platform punt-keepalive disable-kernel-core
platform console virtual
platform hardware throughput level MB 50
!
hostname CSR-PUBLIC-CLOUD
!
boot-start-marker
boot-end-marker
!
!
vrf definition Mgmt-intf
!
  address-family ipv4
  exit-address-family
!
  address-family ipv6
  exit-address-family
!
enable secret 5 <PASSWORD>
enable password <PASSWORD>
!
!
aaa new-model
!
!
aaa authentication login default local
!
aaa session-id common
!
```

```

no ip domain lookup
!
subscriber templating
!
multilink bundle-name authenticated
!
crypto pki trustpoint TP-self-signed-2164836809
  enrollment selfsigned
  subject-name cn=IOS-Self-Signed-Certificate-2164836809
  revocation-check none
  rsakeypair TP-self-signed-2164836809
!
!
crypto pki certificate chain TP-self-signed-2164836809
certificate self-signed 01
  3082022B 30820194 A0030201 02020101 300D0609 2A864886 F70D0101 05050030
  31312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 32313634 38333638 3039301E 170D3133 31323239 32313231
  34305A17 0D323030 31303130 30303030 305A3031 312F302D 06035504 03132649
  4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D32 31363438
  33363830 3930819F 300D0609 2A864886 F70D0101 01050003 818D0030 81890281
  8100DB83 699A0F42 FA8E63EF E000758B 8D047DE3 37255457 E05243D6 67593E0F
  FCB45482 E74FD330 8192C76C C4A571AB DB2161C0 A3CF3260 7DE110A2 219FA685
  D4CDF7ED F7CCF7D2 12001E4E 8B1BA5C8 A6719744 70E3F5C6 46114A9F 835CFBCB
  7FFD5C97 3C9EB547 BA004E07 1797A807 FA076AC8 51701F83 01EBCCBC C54353B8
  21010203 010001A3 53305130 0F060355 1D130101 FF040530 030101FF 301F0603
  551D2304 18301680 149F8BA3 2A9534E3 A33C9944 57033C2B FA40249C 6E301D06
  03551D0E 04160414 9F8BA32A 9534E3A3 3C994457 033C2BFA 40249C6E 300D0609
  2A864886 F70D0101 05050003 81810055 AF9CFFD5 792A393B 07677123 6401BDB7
  B6AA2C8C 7025DCAB 3EE47A14 0375BFC4 DB576CF5 802C793C 35AE1831 6F7EE458
  5D977BE3 00922F11 4501E3F8 AE7F3EBD 47B88F13 CFA164AA B1C739E2 5C220BFD
  FB054B9F 7D71685B 42B513B5 10B326F4 CAD07670 1492448E EEFDD25C 4C9A3F3D
  2F08681B EA46658D FAB36513 94B4B5
quit
license udi pid CSR1000V sn <SERIAL NUMBER>
license accept end user agreement
license boot level premium
spanning-tree extend system-id
!
username admin password <PASSWORD>
!
redundancy
mode none
!
!

```

```
class-map type appnav match-all waas
  match access-group name waas
!
policy-map type appnav waas
  class waas
    distribute service-node-group waas
!
!
crypto isakmp policy 1
  encr aes
  authentication pre-share
  group 14
crypto isakmp key cisco address <PUBLIC IP ADDRESS ENTERPRISE CSR>
!
crypto ipsec transform-set VPN esp-aes esp-sha-hmac
  mode tunnel
!
crypto ipsec profile VPN-Profile
  set transform-set VPN
!
service-insertion service-node-group waas
  service-node 10.0.200.200
!
service-insertion appnav-controller-group waas
  appnav-controller 10.0.200.1
!
service-insertion service-context waas/1
  appnav-controller-group waas
  service-node-group waas
  service-policy waas
  vrf default
  enable
!
!
interface Loopback0
  ip address 192.168.7.7 255.255.255.255
  ip ospf 1 area 0
!
interface Tunnel2
  description VPN TO ENTERPRISE
  ip address 10.0.20.1 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1 area 0
  load-interval 30
  tunnel source GigabitEthernet2
```

```

tunnel mode ipsec ipv4
tunnel destination <PUBLIC IP ADDRESS ENTERPRISE CSR>
tunnel protection ipsec profile VPN-Profile
!
interface LISP0
!
interface GigabitEthernet1
description CONNECTED TO PRIVATE NETWORK
ip address 10.0.5.1 255.255.255.0
negotiation auto
lisp mobility LISP1
service-insertion waas
!
interface GigabitEthernet2
description CONNECTED TO INTERNET
ip address <PUBLIC IP ADDRESS AND MASK>
negotiation auto
!
interface GigabitEthernet3
ip address 10.0.200.1 255.255.255.0
ip ospf 1 area 0
negotiation auto
!
interface AppNav-Compress1
ip unnumbered GigabitEthernet3
no keepalive
!
interface AppNav-UnCompress1
ip unnumbered GigabitEthernet3
no keepalive
!
router lisp
locator-set DC1
192.168.7.7 priority 1 weight 100
exit
!
eid-table default instance-id 0
dynamic-eid LISP1
database-mapping 10.0.5.0/24 locator-set DC1
map-notify-group 239.0.0.1
exit
!
exit
!

```

IP address is the same as on default gateway on the enterprise router.

```

ipv4 use-petr 192.168.6.6
ipv4 itr map-resolver 192.168.6.6
ipv4 itr
ipv4 etr map-server 192.168.6.6 key cisco
ipv4 etr
exit
!
router ospf 1
  router-id 192.168.7.7
!
!
!
ip forward-protocol nd
!
no ip http server
no ip http secure-server
ip route 0.0.0.0 0.0.0.0 <IP ADDRESS TO REACH THE INTERNET>
!
ip access-list extended waas
  permit tcp any any
!
!
control-plane
!
!
line con 0
  stopbits 1
line aux 0
  stopbits 1
line vty 0 4
  privilege level 15
  transport input ssh
!
end

```

Cloud Cisco CSR uses enterprise CSR as its Proxy-ETR, so whenever its routing table points to the default route (0.0.0.0/0), it sends the traffic toward PETR. Cloud CSR is a LISP ITR and ETR.

Integration with Virtual WAAS

Cisco Virtual WAAS (vWAAS) is used on the enterprise and the cloud to accelerate applications and to optimize the WAN bandwidth between the enterprise and the cloud. vWAAS is integrated in the hybrid cloud solution by using AppNav on the Cisco CSR 1000V and vWAAS service nodes. The Cisco CSR 1000V is the AppNav controller, and it distributes traffic from the CSR to the vWAAS. A WAAS Central Manager (WCM) should be deployed to configure and monitor the vWAAS service nodes, but those details are outside the scope of this document.

The configuration details for AppNav are presented in the “Configuration Example” section of this document. After you have finished the initial configuration, for any future subnet where WAN acceleration service is required use the command `service-insertion` under the interface configuration mode to enable AppNav interception and traffic redirection to the vWAAS node.

Packet Walk and Encapsulation Stack

This section explains the detailed communication flows referring to the diagram shown in Figure 1 earlier in this document.

PxTR-1, at the enterprise data center, is placed on a stick. PxTR-1 proxy replies on behalf of all nonlocal servers, inserting its own MAC address for any EID that is not locally detected. PxTR-1 can be either a physical router (Cisco ASR 1000), or a virtual router (Cisco CSR 1000V). PxTR-1 is enabled as a LISP PiTR so you can manage both locally sourced traffic and traffic coming from the WAN. PxTR-1 is also set up as a LISP PeTR so you can receive traffic back from the cloud and deliver it to the WAN. In summary, this node is a PxTR with respect to its role for handling LISP traffic.

At the cloud LISP site, xTR-2 is a standard LISP xTR for the locally attached subnets. xTR-2 is the default gateway for the extended subnets. xTR-2 plays 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 PxTR-1 deployed at the enterprise site.

LISP-Enabled Intra-Subnet Communication between Enterprise and Cloud

Virtual machine VM-1 at the enterprise site sends an ARP request for the IP address of VM-2 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) because VM-2 is not detected locally. Traffic is then sent to PxTR-1, and then it issues a map request to the map server to find the location of VM-2. Finally, it encapsulates the traffic toward xTR-2 at the cloud site because VM-2 is identified in the map server as belonging to that site. Traffic is then delivered toward the cloud-attached subnet. On the return path, the flow is handled in a symmetrical reverse way compared to the inbound one described previously.

Communication from Non-LISP-Enabled Remote Sites to Enterprise and Cloud

Traffic from user PCs, which are located at a remote site that is not LISP-enabled, is naturally attracted toward the enterprise data center by IP routing. When it reaches the enterprise site, it crosses the site's security and other services to reach the local subnet that is supposed to host the destination server VM-2. When the local default gateway sends an ARP request to find VM-2, PxTR-1 responds to this ARP request using the Proxy-ARP function as described previously. Traffic is sent LISP-encapsulated to the cloud site, where it is delivered to VM-2.

xTR-2, which is the default gateway for VM-2, handles the return traffic that is sent by VM-2. Because this traffic is intended for a non-LISP site, (that is, it is targeted to a user PC), xTR-2 sends the traffic to the PeTR configured on it (PxTR-1).

Communication from Enterprise Local LISP-Enabled Subnet to the Cloud LISP-Enabled Subnet

Traffic originated from a LISP-enabled subnet (from server B) intended for another LISP-enabled subnet and further extended to the cloud (to VM-2) first reaches the local gateway (Agg 1) and then will be routed locally to the extended subnet where PxTR-1 will respond to the ARP request issued by Agg 1. On the return path, the traffic will hit xTR-2, which is the default gateway, and then be routed by LISP toward the enterprise site.

Communication from Enterprise Local Non-LISP-Enabled Subnet to Cloud LISP-Enabled Subnet

In this case, traffic will first reach the default gateway (Agg 1), from which it will follow the same path as the traffic originated from a remote site. The PxTR function will be used in both directions.

Communication from LISP-Enabled Subnets to Non-LISP-Enabled Subnets

When traffic is sourced from a LISP-enabled subnet at the enterprise site (VM-1) toward 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.

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.

Inter-Subnet Communication between Servers in the Cloud

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

Communication from LISP-Enabled Remote Sites to Enterprise and Cloud

All previous considerations of traffic flows assume that the only LISP-enabled devices are PxTR-1 and xTR-2. If one remote site needs 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.

Conclusion

This whitepaper described a solution that enables enterprises and cloud providers to deploy a secure hybrid cloud extension with CSR 1000V using LISP.

This solution described can be deployed by enterprises regardless of the cloud service provider or can be offered as a service by cloud service providers to enterprise customers. It provides the ability to move servers between an enterprise and the cloud while keeping the same IP address and without extending the failure domain. It works with any standard VM in the cloud, with no need to modify the VMs before migration. The solution is hypervisor agnostic given that the CSR 1000V supports multiple hypervisors. It can also allow direct access to the VMs in the cloud from remote user sites if LISP is deployed on them. As demonstrated in this whitepaper, it also provides firewall and vWAAS services.

Last but not least, the solution described on this whitepaper has been submitted to IETF on the draft "[Using LISP for Secure Hybrid Cloud Extension](#)" as part of Cisco's commitment to promote and support an open, standards-based hybrid cloud environment.

For more information, please visit:

- [Cisco Cloud Services Router 1000V Web Page](#)
- [Cisco Cloud Services Router 1000V Free Trial](#)
- Learn more about [LISP](#)

Appendix

Overview of Key Technologies


This section provides a detailed overview of the key technologies that have been described in this white paper.

Cisco CSR 1000V

The Cisco CSR 1000V is Cisco's virtual router offering that you can deploy in private, public, or hybrid cloud environments. It offers a rich, tested feature set and solid performance that enables enterprises and service providers to migrate to the cloud while still maintaining security, application experience, and business continuity that they have come to expect in the physical world.

The Cisco CSR 1000V has its origins in the Cisco ASR 1000 Series of physical platforms. The Cisco ASR 1000 is one of Cisco's most successful platforms; it is widely deployed at enterprises and service providers worldwide. The Cisco CSR was developed by taking the Cisco ASR 1001 model, removing the hardware, and embedding the resulting container in a virtual machine that runs on a hypervisor. The Cisco CSR is built to run on general-purpose x86 hardware with a hypervisor providing the abstraction layer. The CPU, memory, hard disk, and network interfaces are generalized and presented to the guest OS (Cisco IOS XE Software) by the hypervisor (Figure 4).

Figure 4. Cisco CSR 1000V: Virtualized Cisco IOS XE Software Router



Key aspects of the Cisco CSR 1000V architecture:

- The guest OS is Cisco IOS XE Software, which runs a 64-bit MonteVista Linux kernel.
- The Cisco CSR 1000V retains the multithreaded nature of Cisco IOS XE Software. The threads are mapped to Linux.
- The route processor and forwarding processor are implemented as processes and mapped to virtual CPUs.
- The network interfaces are mapped to virtual network interface cards (vNICs). VMXNET3 para-virtualized drivers are the default.
- There is no dedicated crypto engine - the Intel Advanced Encryption Standard New Instructions (AES)-NI instruction set is used to provide hardware-based crypto assist.

- The same Cisco IOS XE Software command-line interface (CLI) is available - the Cisco CSR can be configured just like a physical router.

Locator/Identifier Separation Protocol

With the emergence of the cloud architecture, innovation is required in networking to allow IP to gain two mandatory missing features: IP mobility and VPNs.

The Locator Identity Separation Protocol (LISP) is a routing architecture that creates a new paradigm by splitting the device identity, known as an endpoint identifier (EID), and its location, known as its routing locator (RLOC), into two different numbering spaces. This capability brings renewed scale and flexibility to the network in a single protocol, enabling mobility, scalability, and security.

LISP is an overlay routing protocol in the sense that it allows decoupling of the core transport, which can be any IP transport method, from the edge, which can be any IP application site. LISP is intended for the cloud because it allows dynamic resource provisioning independently from the network infrastructure. In short, any IP address can be positioned anywhere it is needed.

LISP Is a Routing Architecture

LISP is a routing architecture, not a feature; it gives IP a full set of capabilities that it does not natively have.

LISP enables IP address portability, which can be seen in two ways. First, it allows the mobility of a host anywhere without changing the host IP address. Second, it allows defining an edge host IP address independently from the site IP structure it will reside on. The decoupling of the application definition from the network is critical for cloud flexibility.

LISP enables network VPN, allowing interconnecting Virtual Route Forwarding (VRF) instances over any IP network, giving to IP a similar capability as MPLS, but not limited to the core of the network, because virtualization is extended to any edge.

Cloud clearly requires huge scalability, and LISP also differentiates itself in this domain. LISP is based on a “pull” model. Similar to Domain Name System (DNS), LISP has one central mapping system where any node registers. When somewhere else in an organization an association between an edge identity versus routing location is required, then it is pulled from this mapping database. This feature differentiates LISP from OSPF or Border Gateway Protocol (BGP), which are “push” models where the full routing information is stored in the forwarding plane of every node. Like DNS, LISP is massively scalable.

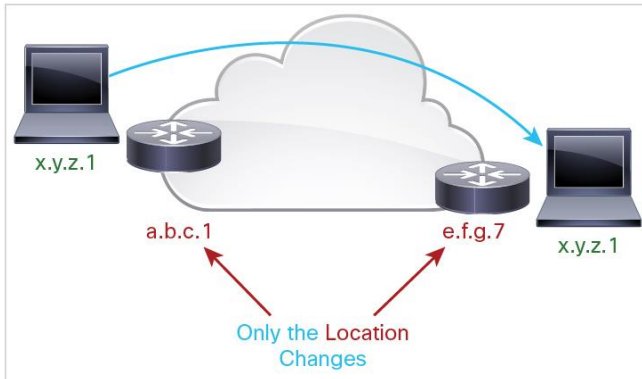
The overlay aspect of LISP should also be emphasized. LISP is an over-the-top technology, tunneling over the IP core and all the IP edge flows. This tunneling allows LISP to use any type of transport, meaning the Internet, a multi-Autonomous System (multi-AS) network, a private infrastructure, IPv4 or IPv6, as well as IP VPN services provided through MPLS. LISP can be encrypted natively in point-to-point or multipoint IPsec. This isolation of the edge world that becomes agnostic from the transport is critical for the cloud architecture.

Finally, LISP is an open standard service with no intellectual property rights; LISP is also very active at the IETF (lisp@ietf.org) and the object of several multivendor implementations.

LISP Basic Concepts

To understand LISP, it is important to understand the concept of “Location to Identity Separation” (Figure 5).

Figure 5. Location to Identity Separation



In traditional IP, the IP edge routing subnets are advertised all over the network using either an Interior Gateway Protocol (IGP) or an Exterior Gateway Protocol (EGP); it is very rare to advertise any host address (subnet mask /32). Most of the time, subnets larger or equal to /24 are used. In IP because all routes are advertised everywhere and installed in the forwarding plane, it is important to limit the amount of entries. To do so, IP subnets are strictly limited to a geographical area, and a subnet is managed by only one pair of routers - the default gateway, implying that if a node moves location, then its IP address must be updated accordingly to the local default gateway and subnet. This constraint is very cumbersome; in order to escape from it, we see more and more cross-site VLAN extensions with all the drawbacks this approach can raise.

With LISP, such constraints disappear; LISP splits the host ID (EID) from the RLOC, allowing any host to move from location to location while keeping its unique identity.

LISP architecture is composed of several elements, including the following:

1. ETR: Egress tunnel router
 - It registers the EID address space it is authorized for
 - It is identified by one (or more) RLOCs
 - It receives and decapsulates the LISP frames
2. Map server:
 - This server is the database where all EID and RLOC associations are stored
 - It can be deployed simply on a pair of devices
 - Or it can be a hierarchy of devices, organized like a DNS system for massive scale implementation (LISP-DDT)
3. ITR: Ingress tunnel router
 - Sends requests toward the map resolver
 - Populates its local map cache with the learned association
 - Responsible to perform the LISP encapsulation
4. Map resolver:
 - Receives the request and selects the appropriate map server

5. Proxy xTR:

- The point of interconnection between an IP network and a LISP network, playing the role of ITR and ETR at this peering point

An ETR is authoritative for a subnet, and registers it using a “map-register” message to the map server.

When triggered on the data plane by a packet destined to a remote EID, the ITR performs a “map request” toward the map resolver, which forwards it to the right map server, which then forwards it to the authoritative ETR.

This ETR replies to the requesting ITR using a “map-reply” message that contains the list of the RLOCs that can reach the requested EID, with their characteristics in terms of priority of usage and weighted load repartition.

AppNav and Wide Area Application Services

Maintaining user experience and application performance is a key consideration when making the transition to the cloud. The Cisco CSR 1000V can act as a traffic control point in the cloud by offering a feature such as AppNav. AppNav allows the Cisco CSR to intelligently redirect TCP traffic to virtual WAAS instances capable of optimizing traffic, reducing redundant flows and ultimately reducing costs for the customer by saving on precious bandwidth.

The AppNav-XE solution for the Cisco CSR 1000V includes the following:

- AppNav Controller: Component that intelligently distributes traffic from a router to services. It can intelligently redirect new flows based on the load on each service node, including loads of individual Layer 7 application accelerators. For flows that do not require any optimization, service nodes can inform the AppNav Controller to pass directly through the packets, thereby minimizing the latency and resource usage.
- AppNav service node auto-discovery feature: With this feature, the Cisco CSR 1000V automatically discovers the service nodes within the same Layer 2 connectivity of the AppNav-XE router (Cisco CSR 1000V) and adds them to the service node cluster.
- Cisco WAAS Central Manager (WCM): WCM monitors and configures the AppNav-XE component.

In addition, Cisco CSR 1000V offers several features that enable customers to retain control over the traffic and apply policies and value-added services to enable a better user experience. The Cisco CSR 1000V supports the full QoS stack, thereby enabling customers to mark, classify, and prioritize high-priority traffic. With support for the Cisco Application Visibility and Control (AVC) solution powered by the deep-packet inspection Network-Based Application Recognition (NBAR2) feature, more than 1500 applications can be detected and appropriate QoS policies can be applied.

High Availability and Redundancy

The Cisco CSR 1000V offers different high-availability services: Virtual machine-level high availability and network-level high availability. The Cisco CSR 1000 supports all the VMware high-availability features such as fault tolerance, dynamic resource scheduling (DRS), vMotion, NIC teaming, and NIC load balancing. These features will help cloud providers achieve virtual machine-level high availability.

In addition, the Cisco CSR 1000 supports several network high-availability features such as Hot Standby Router Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), and the Gateway Load Balancing Protocol (GLBP). These features provide high network availability by providing first-hop redundancy for IP hosts. Also the Cisco CSR 1000 supports Bidirectional Forwarding Detection (BFD), which provides fast-forwarding path-failure detection times for all media types, encapsulations, topologies, and routing protocols. Network administrators can use BFD to have consistent and predictable reconvergence time.

Manageability: Automated Deployment with BDEO and REST API

The cloud is expected to offer tremendous capacity with flexibility and elasticity. A consistent, scalable, and reliable management platform is required to deliver on these expectations. For a virtual router such as the Cisco CSR 1000, this requirement could translate to either dynamically spinning up a new CSR instance or applying a higher throughput license in real time. Deployment and provisioning workflows are also important. A common use case would be to bring up a preconfigured instance of the CSR, create a template, and then deploy additional CSR instances (from the template) with further feature configuration and customization depending on the service level chosen by the end customer. When the instance is up, it has to be continuously monitored, audited, and perhaps even torn down at some point.

The Cisco CSR integrates with the standard hypervisor vendor deployment platforms such as vCenter or vCloud Director from VMware. It also integrates with openstack for non ESXI hypervisors. You can also use the Cisco Prime™ Network Services Controller to provision, manage, and monitor the Cisco CSR 1000V. Cisco Prime Network Services Controller is a single, integrated solution for lifecycle management of virtual networking devices.

The Cisco CSR 1000V also offers the Cisco “Build, Deploy, Execute OVF (BDEO)” tool, which is included in the Cisco CSR 1000V software package. This Linux-based application enables you to create a custom OVA image with customized Cisco IOS XE Software configuration for your Cisco CSR 1000. The Cisco Build, Deploy, Execute OVF (BDEO) tool provides a simple command-line interface to enter the virtual-machine attributes into the .ova file. This tool can speed the process of deploying the Cisco CSR 1000V on multiple virtual machines. For more information about how to use the BDEO tool, please visit:

<http://www.cisco.com/c/en/us/td/docs/routers/csr1000/software/configuration/csr1000Vswcfg/swinstallcsr.html#wp1277790>.

The Cisco CSR 1000 offers traditional Cisco IOS CLI-based configuration management. Several service providers are CLI-knowledgeable and comfortable using scripts that configure and manage the Cisco CSR 1000. Other service providers and enterprises require an application programming interface (API)-directed solution that abstracts the configuration complexity from the end user. The Cisco CSR 1000 offers a Representational State Transfer (REST) API that provides an alternative method to the Cisco IOS XE CLI to provision selected functions on the Cisco CSR 1000V. REST is HTTP-like and based on a client-server model. The Cisco CSR 1000 implementation supports the JSON format. Being REST-capable enables the CSR to be compatible with several different cloud orchestration platforms. Today, Cisco CSR 1000V REST APIs support many Cisco IOS XE Software features and technologies such as NAT, VPN, firmware, access control list (ACL), VRF, routing protocols, DNS, DHCP, and others. LISP REST API support is coming on the Cisco CSR 1000 in the Cisco IOS XE Software Release 3.13.




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)