Abstract

Government mandates, e-business and Internet growth requirements, and impending IPv4 address exhaustion concerns are prompting many enterprises to begin implementing an IPv6 transition strategy. The available techniques for supporting IPv6 transition strategies fall into three categories: IPv6 native (dual-stack), IPv6 tunneling, and IPv6 translation. Most of these techniques have limited capabilities, or are difficult or have constraints in deployment. The Locator/ID Separation Protocol (LISP) implements a new routing architecture that is designed for a much broader purpose than IPv6 transition. LISP was designed to be address family agnostic, and when incorporated LISP into an IPv6 transition strategy, LISP has demonstrated quick deployment time, low deployment and operational cost, little or no need for additional equipment or modifications, and high user-satisfaction. This document describes two common IPv6 transition use cases and shows how LISP can be incorporated to achieve success.

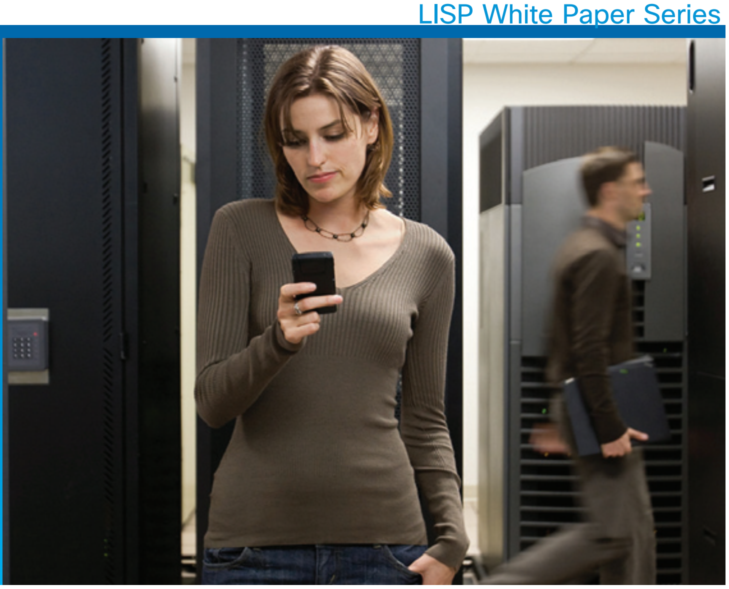

Enterprise IPv6 Transition Strategy

Using the Locator/ID Separation Protocol

Background

The ubiquity of the Internet and TCP/IP enables new opportunities for enterprises and consumers. However, as has been and continues to be well covered in both the media and technical forums, IPv4, the IP addressing scheme utilized by today's Internet, is approaching the point of address exhaustion. Some experts fear that this address exhaustion will have a limiting effect on new address allocations and thereby threaten Internet growth. The adoption of the IPv6 addressing scheme is now seen as the way forward.

To call IPv6 " new" would be incorrect, given its more-than-15 year history, but its adoption-rate has been slow. The reasons for this slow deployment vary, but the most often cited reasons are its lack of interoperability with IPv4, and its lack of vendor support. Because of these limitations, transitioning from IPv4 to IPv6 is not transparent or " free." It is well accepted that IPv6 resolves the IPv4 address-space depletion problem. However, this factor alone has not been sufficient, to date, to warrant wide-scale adoption. In addition, IPv6 does not rectify other " infrastructural" challenges such as route scalability and mobility (which might have been sufficient motivators to promote adoption).

During the intervening past 15 or more years, many " workarounds" have arisen in the IPv4 address-space such as Classless Inter-Domain Routing (CIDR), RFC 1918 addresses, and Network Address Translation (NAT), which have extended the availability of IPv4 addressing at much lower costs and technical requirements than those associated with converting or adopting IPv6. Until now these workarounds have offered sufficiently less expensive alternatives and have forestalled the compelling event required to justify the commitment of capital expenditures (CapEx) and operating expenditures (OpEx) to the IPv6 adoption process. When viewed from a business perspective, without a tangible payoff, the IPv6 transition has been hard to justify. Now, however, IPv4 address exhaustion and its potential threat to a growing Internet economy appear to be providing this compelling event.

IPv6 Deployment Solution Options

Many techniques exist to develop an IPv6 transition strategy. This section briefly describes at a very high level some of the network-based options available. (Host-based techniques also exists but are not discussed here.) Generally speaking, network-based IPv6 transition techniques fall into three categories:

- Dual-stack IPv4 andIPv6
- IPv6 tunneling
- IPv6 translation

Each approach has its features, benefits, and limitations. Understanding these parameters will lead to the selection of the most appropriate strategy for your situation. Be sure to note that the economics can be substantially different depending on the approach taken; these approaches are not equal in cost, complexity, or functional capabilities. Understanding the available options and deciding, based on your particular needs, which provides the necessary capabilities will help you select the most appropriate option–both as a transition strategy and as a long-term IPv6 strategy. As this document discusses, incorporating Locator/ID Separation Protocol (LISP) into your IPv6 transition strategy can be a compelling choice.

This document does not discuss in detail all the available approaches; existing literature provides many examples and comparisons for most of these IPv6 transition approaches. However, it does provide a brief, high-level overview of the three categories to provide the framework for the detailed LISP discussions that follow.

Dual-Stack IPv4 and IPv6

Dual-stack IPv4 and IPv6 techniques involve deploying IPv6 directly within the network infrastructure. In this case, IPv6 can be deployed by itself, or more commonly in a dual-stack (IPv4 and IPv6) implementation. Because IPv4 and IPv6 are literally different protocols, not unlike multiprotocol network environments of the past such as Internetwork Packet Exchange (IPX), AppleTalk, and IP, both can be run concurrently and each uses its own protocol stack independently of the other.

A native IPv6 deployment is considered the most direct deployment approach for IPv6 Internet connectivity. Dual-stack discussions usually focus on subscriber and front-end services, core infrastructure, and enterprise applications and back-end systems. Depending on the scale, a dual-stack approach can also be the most expensive and time-consuming to deploy, especially when considering all possible systems. It can potentially require new hardware or changes to existing hardware, considerable planning, a well-defined rollout plan, and a skilled operations staff. For example, if you need to address all three areas, and if your hardware is capable or can be upgraded, and if your network staff is capable or can be trained, and if you have the time and budget to roll out a native IPv6 deployment, then the dual-stack approach is a viable choice. However, if these are not requirements, then the dual-stack approach may not be the best first option. If your goal is simply to explore and gain experience with IPv6, or if you have a very limited time frame in which to roll out a system to meet an emerging market demand based on competitive requirements, then other options, or a combination of options such as combining dual-stack and LISP are better suited.

IPv6 Tunneling

Tunneling techniques involve transport through encapsulation of one protocol within another protocol. IPv6 tunneling encapsulates IPv6 packets within IPv4 packets and uses the existing IPv4 core to allow IPv6 end systems (islands) to communicate without the need to upgrade the intermediate IPv4 infrastructure between them. (This approach could also apply to the opposite situation; IPv4 packets carried over an IPv6 core.) Many such schemes have been described including: IPv6 to IPv4 (6to4), IPv6 rapid deployment (6RD), and IPv6 in IPv4 generic routing encapsulation (6in4-GRE).

Tunnels are typically manually or automatically configured. Manually configured tunnels are used primarily as stable links for regular communications, and are easily deployed over existing IPv4 infrastructures. However, because manually configured tunnels require configuration at both ends of the tunnel, they incur a large management overhead and have troubleshooting challenges. In addition, some tunneling mechanisms such as 6to4 and 6RD require the participation of the service provider; others, such as 6in4-GRE, can be deployed directly by the enterprise. The routers performing the tunneling are required to be dual-stack, but the core can run IPv4 only. Failure detection and resiliency can also be challenging for static configurations. In addition, some tunneling approaches are technically difficult to scale, making large deployments difficult to deploy and manage.

As you will see in this document, LISP is a special case of tunneling. All tunneling methods share the basic notion of the encapsulation of one protocol within another, but LISP's similarities with other tunneling mechanisms end here. LISP encapsulation is dynamic and requires no preconfiguration of tunnel endpoints. It is designed to work in a multi-homing environment and supports communications between LISP and non-LISP sites for simple interworking. Additional information about LISP, including information about its deployment in several common IPv6 transition use cases is presented in this document.

IPv6 Translation

IPv6 translation schemes implement some form of packet-header translations between the IPv6 and IPv4 addresses. The goal is to translate packets with IPv6 addresses to those with IPv4 addresses, so that IPv6-only hosts can talk to the IPv4-only Internet. On the surface process may sound simple, and in some cases it can be. For example, some server load balancers (SLBs) are capable of load balancing and translating packets with IPv6 addresses to IPv4 packets. This capability is advantageous in data center deployments in that the existing IPv4 infrastructure can remain unchanged, and only the SLB requires modification. This kind of technique requires other mechanisms (such as tunneling or dual-stack techniques) to get the IPv6 packets to the SLB. However, as a means for accessing the IPv4 Internet, the translation technique can actually be quite complex. Originally, NAT-Protocol Translation (NAT-PT: RFC 2766) was proposed for this purpose, but it has subsequently been deemed impractical to deploy (due to the need for application layer gateways (ALGs), the management of IPv6 Domain Name System (DNS) requests, and the need for session initiation ordering for connection state establishment). In addition, IPv6 translation approaches break the Internet end-to-end connectivity model (a common trait of all NAT implementations), and cannot handle fragmentation in the core.

Choosing an Option

Each technique potentially has a role to play in IPv6 transition strategies, but the best choice (or choices) should be based on the specific details of each use case. Thus, you should consider, at a high level, the questions that need to be addressed for any IPv6 transition strategy, as listed in Table 1 below. After you have reviewed these questions, and assuming that you have decided to go forward with IPv6, the next step is to select an IPv6 transition approach that makes sense both financially and technically.

Table 1. Some Important Questions to Help Guide an IPv6 Transition Strategy

Questions and Answers	LISP Advantages
Why is IPv6 being considered? 1. We are a government agency or an Enterprise doing business with the government and have a mandate to deploy IPv6 by the end of 2011. 2. We have an urgent need to establish an IPv6 web presence to maintain market leadership. 3. We have a long-term IPv6 strategy but are not sure where to begin. It seems so overwhelming to convert everything. Initially we just need to gain experience with IPv6 before committing to a substantial build-out.	Whether you are deploying IPv6 in a production environment to establish a web presence for example, or you just need to gain hands-on experience with the IPv6 protocol, LISP can help you achieve your goals. Using LISP, you can deploy IPv6 and retain your existing IPv4 WAN connectivity. Using LISP gives you the ability to rapidly deploy IPv6 without significant CapEx or OpEx. In addition, the broader features of LISP (inherent multi-homing, ingress traffic engineering, and mobility) can be significant benefits for long-term deployments.
What is the intent of the IPv6 deployment? 1. This IPv6 deployment is intended as a public deployment (e.g. IPv6 Internet web presence, IPv6 Internet access to corporate assets, etc.). 2. This is a private IPv6 deployment (e.g. a corporate IPv6 VPN, connection of IPv6 islands, etc.).	Establishing an IPv6 web presence or deploying a private IPv6 VPN can both be accomplished with few changes to your existing network infrastructure using the inherent address family-agnostic design of LISP. Using LISP allows you to retain your existing IPv4 WAN connectivity, allowing you to meet your goals quickly and efficiently.
Does your current hardware and infrastructure (routers, switches, load balancers, DNS, etc.) support IPv6? 1. My entire infrastructure is IPv6-capable. For example, my SLB can perform IPv6-to-IPv4 load balancing. 2. Some of my infrastructure is IPv6-capable. For example, my routers can support IPv6, but my SLB cannot perform IPv6-to-IPv4 load balancing.	Obviously, some components must support the IPv6 protocol for you to deploy any IPv6 transition solution, but which components and how much of the network must be converted to IPv6 depends on your choice of an IPv6 transition scheme. Incorporating LISP reduces the need for changes to your existing network infrastructure because of the inherent address family-agnostic design of LISP.
What is your IPv6 deployment CapEx and OpEx budget? 1. I have a large budget for both capital improvements (hardware upgrades, new circuits, etc.) and operational improvements (staff training, management support, etc.). 2. I have a limited budget and need to do as much as I can with my current resources.	Deploying IPv6 can be an expensive undertaking if you use the wrong approach. Dividing the complete process into manageable steps and considering the mandatory requirements first allows you to complete the process incrementally. Using LISP simplifies the initial steps by incorporating existing IPv4 infrastructure, allowing you to meet your goals with a minimal CapEx and OpEx.
How much time do you have to get your IPv6 deployment operational? 1. I need to have IPv6 deployed immediately; anything that saves time is useful. 2. My needs are longer term; I want to take my time and deploy IPv6 correctly.	Deploying IPv6 can take time. For example, just obtaining an IPv6 WAN connection can take months. Using LISP simplifies the initial steps by incorporating existing IPv4 WAN links and infrastructure, allowing you to deploy an IPv6 solution rapidly.
Is this a temporary (throw-away) deployment, or is it a long-term (permanent) solution? 1. This is a temporary installation. After we gain experience with IPv6, this deployment will likely be replaced with a permanent one. 2. This deployment is intended to be a permanent installation.	Deploying IPv6 can be an expensive undertaking, especially if the deployment is temporary. Incorporating LISP helps reduce the number of changes to your existing network and also reduces CapEx and OpEx, which is beneficial for temporary installations. The broader features of LISP (inherent multi-homing, ingress traffic engineering, and mobility) can also be significant benefits for long-term deployments.

When selecting an approach, accurately assessing these important parameters can help you define the best solution. For example, if the intended deployment is new (no existing networks or equipment), you may want to consider an IPv6-only deployment strategy. However, if the intended deployment augments an existing IPv4 network, then IPv4 and IPv6 coexistence must be assumed at the outset, and likely for many years.

The remainder of this document describes in detail the use of LISP as a means for implementing an IPv6 transition strategy. Two common use cases that incorporate LISP as the main IPv6 transition technique are described.

LISP Solutions for IPv6 Transition

LISP makes sense for many IPv6 transition situations, as been proven by real success stories in practical, production implementations (see references [1], [2], and [3] at the end of this document). These deployments have shown common themes including: quick deployment time, low deployment and operational cost, little or no need for additional equipment or modifications, and high user-satisfaction.

LISP is not a feature, nor was it invented as an IPv6 transition mechanism like most of the other choices listed in the preceding section. Rather, LISP is a new routing architecture that was designed for a much broader purpose. it was designed to transparently accommodate multiple address families, and so using LISP in IPv6 transition solutions is very natural. LISP implements a new semantic for IP addressing that creates two namespaces: Endpoint Identifiers (EIDs), which are the current addresses assigned to end-hosts today, and Routing Locators (RLOCs), which are the addresses assigned to devices (primarily routers) that make up the global routing system. Splitting EID and RLOC functions yields many benefits such as improved routing scalability, improved multi-homing efficiency, and IP mobility, and the capability to use the same or different address families for the EIDs and the RLOCs. Therefore, LISP is easy to use for IPv6 transition solutions.

This document does not cover the detailed inner workings of LISP itself, but instead focuses on several common IPv6 transition use cases supported by LISP. Additional material on LISP can be found in references [3] and [4] at the end of this document. However, a short overview of LISP, including the context for of operations, is provided to help you fully understand the use cases.

As mentioned earlier, LISP is a special case of tunneling that uses a dynamic encapsulation approach rather than requiring the preconfiguration of tunnel endpoints. It is designed to work in a multi-homing environment and supports communications between LISP and non-LISP sites for simple interworking. A LISP-enabled network includes some or all of the following components:

LISP Site Devices

- **Ingress Tunnel Router (ITR):** An ITR is a LISP Site edge device that receives packets from site-facing interfaces (internal hosts) and encapsulates them to remote LISP sites, or natively forwards them to non-LISP sites.
- **Egress Tunnel Router (ETR):** An ETR is a LISP Site edge device that receives packets from core-facing interfaces (the Internet) and decapsulates LISP packets and delivers them to local EIDs at the site.

Note: Customer Edge (CE) devices typically implement ITR and ETR functions at the same time. When this is the case, the device is referred to as an xTR.

LISP Infrastructure Devices:

- **Map-Server (MS):** An MS is a LISP Infrastructure device that LISP site ETRs register to with their EID prefixes. The MS advertises aggregates for the registered EID prefixes into the LISP mapping system. All LISP sites use the LISP mapping system to resolve EID-to-RLOC mappings.
- **Map-Resolver (MR):** An MR is a LISP Infrastructure device to which LISP site ITRs send LISP Map-Request queries when resolving EID-to-RLOC mappings.
- **Proxy ITR (PITR):** A PITR is a LISP Infrastructure device that provides connectivity between non-LISP sites and LISP sites by attracting non-LISP traffic destined to LISP sites and encapsulating this traffic to LISP sites. In the IPv6 transition case, the PITR can attract IPv6 non-LISP traffic and forward it to a LISP site using IPv4 as the transport.
- **Proxy ETR (PETR):** A PETR is a LISP Infrastructure device that allows IPv6 LISP sites that have only IPv4 RLOC connectivity to reach LISP and non-LISP sites that have only IPv6 RLOC connectivity.

EID namespace is used within the LISP sites for end-site addressing for hosts and routers. These EID addresses go in DNS records, just as they do today. Generally, EID namespace is not globally routed in the underlying Internet. RLOC namespace, however, is used in the (Internet) core. RLOCs are used as infrastructure addresses for LISP routers and core (service provider) routers, and are globally routed in the underlying infrastructure, just as they are today. Hosts do not know about RLOCs, and RLOCs do not know about hosts. As shown in the IPv6 transition use cases that follow, EIDs are the IPv6 addresses assigned to hosts and servers within LISP sites, and RLOCs are the existing IPv4 infrastructure addresses.

The remainder of this document describes the two LISP deployments that support common IPv6 transition use cases.

Scenario 1: IPv6 Islands Connected Over an IPv4 Core

Scenario

Scenario 1 addresses the IPv6 transition use case in which an enterprise wants to gain basic experience with IPv6, but has no urgent or compelling need that would merit significant CapEx or OpEx outlays or changes to the existing infrastructure. To accomplish this goal, one approach is to create several IPv6 islands within the corporate network–one at headquarters (HQ), and one at each of two remote sites–and then to connect these IPv6 islands together. This deployment can be accomplished rapidly and easily with LISP by connecting these IPv6 islands over the existing IPv4 network without the need for any changes to the underlying network. This results in a cost-effective solution for this scenario.

Topology

Figure 1 illustrates a simplified view of the starting topology for Scenario 1. In this case, three sites of the enterprise are shown: the headquarters site (Headquarters), and two branch-office sites (Spoke 1 and Spoke 2). This initial topology is entirely an IPv4 infrastructure. Based on the requirements in this scenario, IPv6 islands are added at each site, also as illustrated in Figure 1. In this case, one router at each site is configured to run dual-stack to provide the connection to the existing IPv4 topology, as well as carrying the new IPv6 prefixes. These routers also perform the required LISP functions.

Figure 1. Scenario 1 Showing Conceptual Existing IPv4 Topology and Added IPv6 Islands

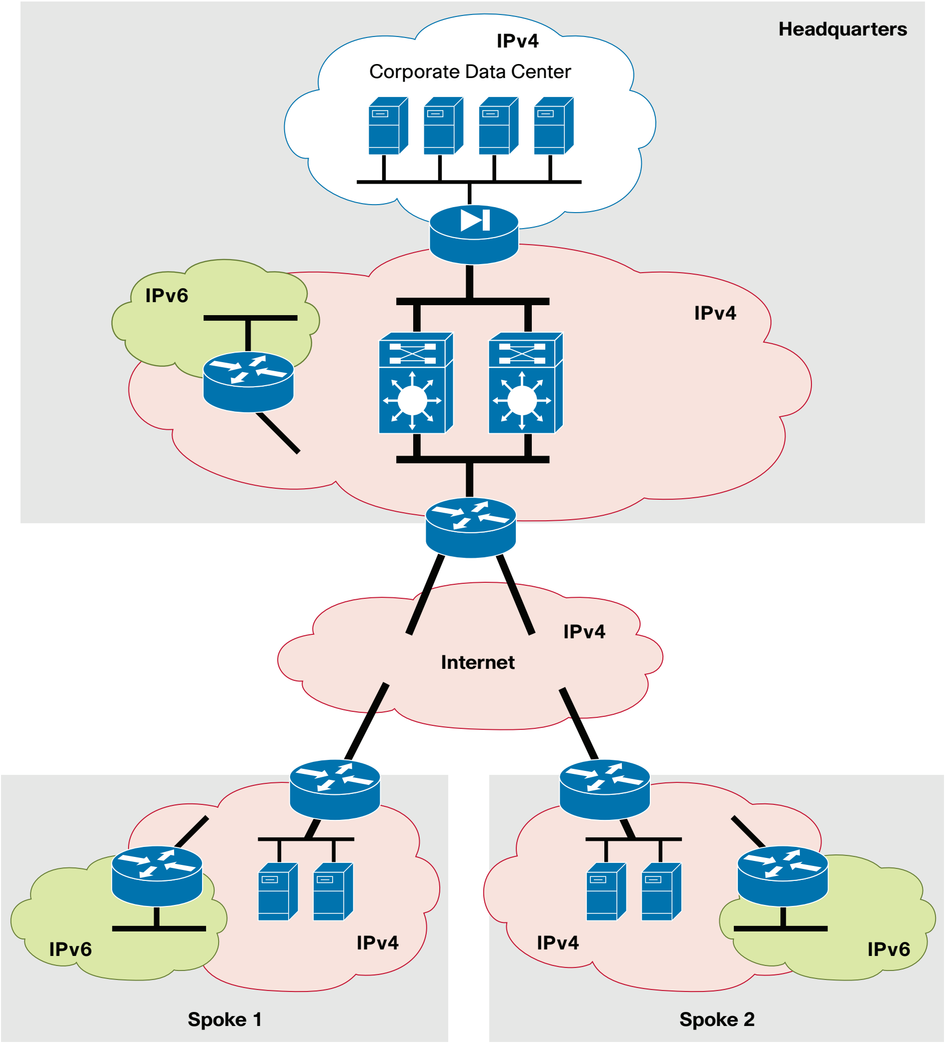

Note: Assuming that the existing routers have acceptable capabilities, new routers are not required to support this scenario. If existing routers can be reconfigured to add the necessary IPv6 subnets and LISP configurations, they can be used to enable this IPv6 solution. However, the addition of new routers to serve as LISP devices does provide a clear demarcation for this IPv6 trial, and reduces the possibility of disruptions of the existing production network.

LISP Configurations

Figure 2 illustrates a simplified topology that reflects the basic elements of Figure 1. This figure provides a clear view of the LISP components used to accomplish IPv6 island connectivity across an IPv4 network. The base requirement for the LISP routers is that they have IPv4 reachability between them. In this case, the locator addresses (LISP RLOCs) are assigned from the enterprise Internet address space (shown in the figures).

In this example, all LISP functions can be run internally in the enterprise. Thus, the following LISP functions are required:

- The HQ LISP router (RTR-A) will be configured to provide LISP mapping services and LISP Encapsulation services. Thus, this router will be configured as a Map-Resolver/Map-Server (MR/MS), and as an Ingress Tunnel Router/Egress Tunnel Router (xTR) concurrently.
- The Spoke-1 and Spoke-2 LISP routers (RTR-B and RTR-C) are configured only as LISP xTRs.

This is one of the most basic of all LISP deployments. (Note that the configurations shown here are based on a Cisco IOS® Software LISP deployment, and are fully functional for the given topology.)

Figure 2. Scenario 1: IPv4 and IPv6 Addressing for LISP Configurations

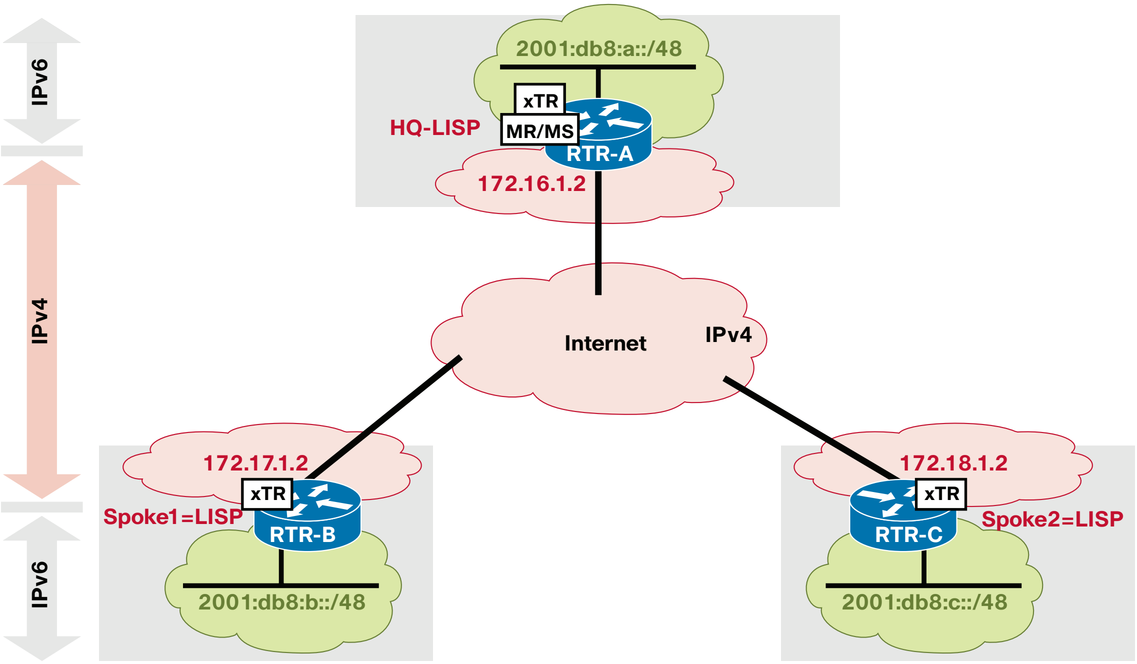

Note: All addresses used in all figures and configuration examples in this document are shown as IPv4 private (RFC 1918) and IPv6 documentation (RFC 3849) addresses as is standard practice in Cisco® documentation. In this example, the IPv4 RLOC addresses (shown as 172.16.1.2, 172.17.1.2, and 172.18.1.2) represent the appropriate routable IPv4 addresses for your environment (e.g. Internet routable addresses). However, since the IPv6 EID addresses are encapsulated within LISP, they can remain as shown (2001:db8:a::/48, 2001:db8:b::/48, and 2001:db:c::/48) if they are applicable only within this private deployment.

Based on the topology in Figure 2, the following configurations apply to the three LISP routers.

RTR-A

RTR-A is the HQ router that provides both LISP MS/MR and ITR/ETR (xTR) services. The configuration for RTR-A follows; the highlighted elements are those relevant to LISP for RTR-A.

```
hostname RTR-A
!
vrf definition lisp
 rd 1:1
 !
address-family ipv6
 exit-address-family
!
ip cef
ipv6 unicast-routing
ipv6 cef
!
lisp site HQ-LISP
 description LISP HQ Site
 authentication-key s3cr3t-hq
 eid-prefix 2001:db8:a::/48
lisp site Spoke1-LISP
 description LISP Spoke Site 1
 authentication-key s3cr3t-1
 eid-prefix 2001:db8:b::/48
lisp site Spoke2-LISP
 description LISP Spoke Site 2
 authentication-key s3cr3t-2
 eid-prefix 2001:db8:c::/48
!
interface LISP0
!
interface Loopback0
 no ip address
 ipv6 address 2001:db8:a::1/48
!
interface Ethernet0/0
 ip address 172.16.1.2 255.255.255.0
!
ipv6 lisp database-mapping 2001:db8:a::/48 172.16.1.2 priority 1 weight 100
ipv6 lisp itr map-resolver 172.16.1.2
ipv6 lisp itr
ipv6 lisp etr map-server 172.16.1.2 key s3cr3t-hq
ipv6 lisp etr
!
ipv6 lisp map-server
ipv6 lisp map-resolver
ipv6 lisp alt-vrf lisp
!
ip route 0.0.0.0 0.0.0.0 172.16.1.1
!
ipv6 route ::/0 Null0
!
```

In the preceding RTR-A configuration, the following LISP components are defined:

- First, since this device acts as a map-server and map-resolver, vrf lisp is defined to provide the container for storing LISP EID prefixes.
- Next, LISP sites are defined for the HQ-LISP (this router), and for each spoke site (Spoke1-LISP and Spoke2-LISP). This definition is required so that each site can register to the MS running on this router. (Note that RTR-A registers with itself for LISP services).
- Next, relevant IPv6 LISP services are enabled. Note the map-server and map-resolver configurations point to this router itself.
- Finally, the IPv6 EID prefix for this router (2001:db8:a::/48) is specified as being associated with the IPv4 locator (172.16.1.2), as referenced by the database-mapping command.
- Note that Loopback0 is configured with the IPv6 EID prefix (2001:db8:a::1/48) as a simple way of demonstrating connectivity in this scenario.

RTR-B and RTR-C

RTR-B and RTR-C are spoke site routers and provide only LISP xTR services. The configurations relevant to LISP for RTR-B follow; the highlighted elements are those relevant to LISP for RTR-B.

```
hostname RTR-B
!
ip cef
ipv6 unicast-routing
ipv6 cef
!
interface LISP0
!
interface Loopback0
 no ip address
 ipv6 address 2001:db8:b::1/48
!
interface Ethernet0/0
 ip address 172.17.1.2 255.255.255.0
!
ipv6 lisp database-mapping 2001:db8:b::/48 172.17.1.2 priority 1 weight 100
ipv6 lisp itr map-resolver 172.16.1.2
ipv6 lisp itr
ipv6 lisp etr map-server 172.16.1.2 key s3cr3t-1
ipv6 lisp etr
!
ip route 0.0.0.0 0.0.0.0 172.17.1.1
!
ipv6 route ::/0 Null0
!
```

In the preceding RTR-B configuration, the following LISP components are defined:

- First, relevant IPv6 LISP services are enabled. Only ITR and ETR services are required. Note that the map-server and map-resolver configurations point to the HQ-LISP router (RTR-A, 172.16.1.2), which is configured as the MR/MS.
- Finally, the IPv6 EID prefix for this router (2001:db8:b::/48) is specified as being associated with the IPv4 locator (172.17.1.2), as referenced by the database-mapping command.
- Note that Loopback0 is configured with the IPv6 EID prefix (2001:db8:b::1/48) as a simple way of demonstrating connectivity in this scenario.

The configurations relevant to LISP for RTR-C follow; the highlighted elements are those relevant to LISP for RTR-C.

```
hostname RTR-C
!
ip cef
ipv6 unicast-routing
ipv6 cef
!
interface LISP0
!
interface Loopback0
 no ip address
 ipv6 address 2001:db8:c::1/48
!
interface Ethernet0/0
 ip address 172.18.1.2 255.255.255.0
!
ipv6 lisp database-mapping 2001:db8:c::/48 172.18.1.2 priority 1 weight 100
ipv6 lisp itr map-resolver 172.16.1.2
ipv6 lisp itr
ipv6 lisp etr map-server 172.16.1.2 key s3cr3t-2
ipv6 lisp etr
!
ip route 0.0.0.0 0.0.0.0 172.18.1.1
!
ipv6 route ::/0 Null0
!
```

In the preceding RTR-C configuration, the following LISP components are defined:

- First, relevant IPv6 LISP services are enabled. In this case, only ITR and ETR services are required. As in the case of Spoke1, the map-server and map-resolver configurations point to the HQ-LISP router (RTR-A, 172.16.1.2), which is configured as the MR/MS.
- Finally, the IPv6 EID prefix for this router (2001:db8:c::/48) is specified as being associated with the IPv4 locator (172.18.1.2), as referenced by the database-mapping command.
- Note that Loopback0 is configured with the IPv6 EID prefix (2001:db8:c::1/48) as a simple way of demonstrating connectivity in this scenario.

Verifications

Figure 3 illustrates the LISP packet flow corresponding to the IPv6 island connectivity enabled for this scenario.

Figure 3. Scenario 1: LISP Packet Flows for IPv6 over IPv4 Connectivity

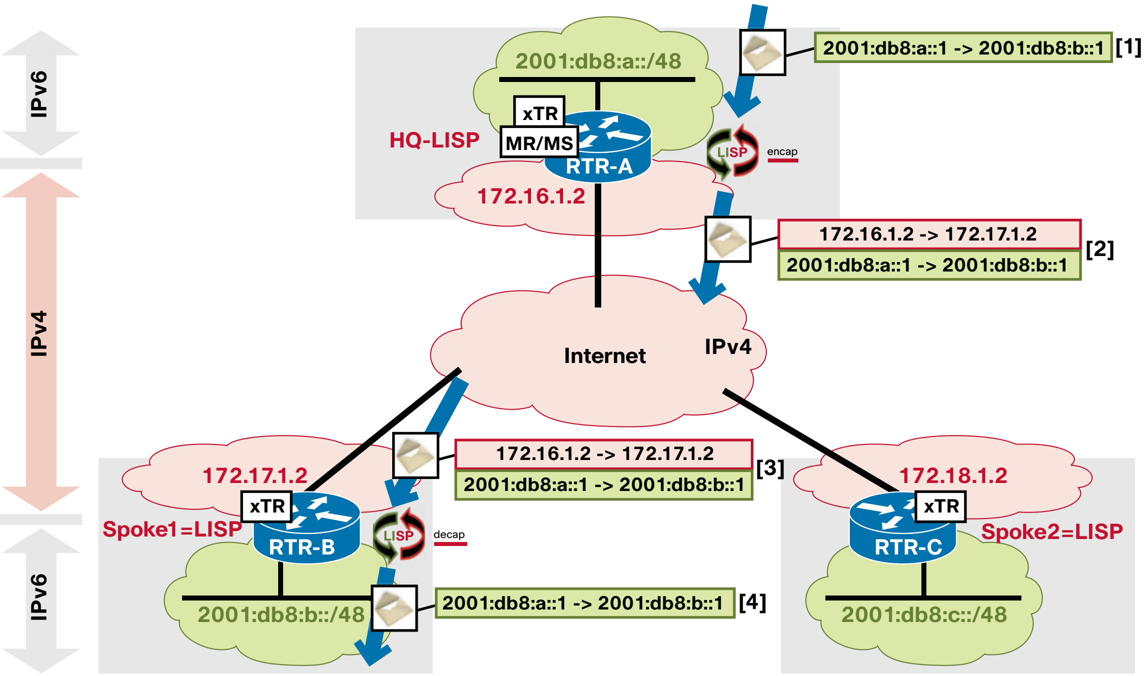

As illustrated in Figure 3, test packets are sourced from the HQ-LISP EID **2001:db8:a::1** and destined to the Spoke1-LISP EID **2001:db8:b::1**. The IPv6 packets are encapsulated by RTR-A using the source and destination IPv4 locator addresses 172.16.1.2 and 172.17.1.2 respectively.

The following output illustrates this packet validation, as well as other useful LISP information.

RTR-A

- RTR-A is the Map-Server. The following output shows the status of the LISP registration of the sites in this example. Note that all sites are listed as up. Also note the EID Prefix that each site registers and the associated IPv4 locator.

```
RTR-A#show lisp site
LISP Site Registration Information

Site Name       Last       Up    Who Last              EID Prefix
                Register         Registered
HQ-LISP         00:00:09   yes   172.16.1.2            2001:DB8:A::/48
Spoke1-LISP     00:00:50   yes   172.17.1.2            2001:DB8:B::/48
Spoke2-LISP     00:00:52   yes   172.18.1.2            2001:DB8:C::/48
RTR-A#
```

- When the IPv6 map-cache on RTR-A is populated, a mapping is obtained for the packet flow shown in Figure 3, as illustrated in the following output. Note that the map-cache entry for the IPv6 EID prefix 2001:db8:b::/48 is associated with the IPv4 locator 172.17.1.2.

```
RTR-A#show ipv6 lisp map-cache
LISP IPv6 Mapping Cache, 3 entries
::/0, uptime: 00:05:07, expires: never, via static
  Negative cache entry, action: send-map-request
2001:DB8:B::/48, uptime: 00:01:03, expires: 23:59:24, via map-reply, complete
  Locator      Uptime    State      Pri/Wgt
  172.17.1.2   00:01:03  up           1/100
2001:DB8:C::/48, uptime: 00:00:24, expires: 23:59:28, via map-reply, complete
  Locator      Uptime    State      Pri/Wgt
  172.18.1.2   00:00:24  up           1/100
RTR-A#
```

- Finally, output is provided showing that a sourced-ping between the HQ-LISP router (RTR-A) and both Site1 and Site2 LISP routers (RTR-B and RTR-C) is successful.

```
RTR-A#ping 2001:db8:b::1 so 2001:db8:a::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:B::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:A::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/4 ms
RTR-A#ping 2001:db8:c::1 so 2001:db8:a::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:C::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:A::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/0 ms
RTR-A#
```

RTR-B

- When the IPv6 map-cache on RTR-B is populated, the output shown here is seen.

```
RTR-B#show ipv6 lisp map-cache
LISP IPv6 Mapping Cache, 3 entries

::/0, uptime: 00:09:47, expires: never, via static
  Negative cache entry, action: send-map-request
2001:DB8:A::/48, uptime: 00:02:51, expires: 23:57:01, via map-reply, complete
  Locator      Uptime    State      Pri/Wgt
  172.16.1.2   00:02:51  up           1/100
2001:DB8:C::/48, uptime: 00:01:13, expires: 23:58:39, via map-reply, complete
  Locator      Uptime    State      Pri/Wgt
  172.18.1.2   00:01:13  up           1/100
RTR-B#
```

- The next output shows that a sourced-ping between the Site1-LISP router (RTR-B) and the Site2-LISP router (RTR-C) is successful (spoke to spoke), as is one between Site1-LISP router (RTR-B) and the HQ-LISP router (RTR-A).

```
RTR-B#ping 2001:db8:c::1 so 2001:db8:b::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:C::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:B::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/4 ms
RTR-B#ping 2001:db8:a::1 so 2001:db8:b::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:A::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:B::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/3/8 ms
RTR-B#
```

RTR-C

- When the IPv6 map-cache on RTR-B is populated, the output shown here is seen.

```
RTR-C#show ipv6 lisp map-cache
LISP IPv6 Mapping Cache, 3 entries

::/0, uptime: 00:10:32, expires: never, via static
  Negative cache entry, action: send-map-request
2001:DB8:A::/48, uptime: 00:03:14, expires: 23:56:38, via map-reply, complete
  Locator      Uptime    State      Pri/Wgt
  172.16.1.2   00:03:14  up           1/100
2001:DB8:B::/48, uptime: 00:02:07, expires: 23:57:45, via map-reply, complete
  Locator      Uptime    State      Pri/Wgt
  172.17.1.2   00:02:07  up           1/100
RTR-C#
```

- The next output shows that a sourced-ping between the Site2-LISP router (RTR-C) and the Site1-LISP router (RTR-B) is successful (spoke to spoke), as well one between the Site2-LISP router (RTR-C) and the HQ-LISP router (RTR-A)

```
RTR-C#ping 2001:db8:b::1 so 2001:db8:c::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:B::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:C::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/2/8 ms
RTR-C#
RTR-C#ping 2001:db8:a::1 so 2001:db8:c::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:A::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:C::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/8 ms
RTR-C#
```

Scenario 1 Recap

In Scenario 1, LISP was used to connect IPv6 islands at the Headquarters site and two spoke sites. The existing IPv4 routable infrastructure was used without any changes. One existing router at each site was modified to include the IPv6 subnet and a few lines of LISP configuration. LISP mapping services were collocated with LISP xTR services on the HQ router, and the spoke sites ran only LISP xTR services. With these small modifications, accomplished in just a few minutes, LISP was able to fully connect the IPv6 islands. This same LISP solution can be extended to other private LISP deployments (IPv4 to IPv4 over IPv4, IPv4 to IPv4 over IPv6, and IPv6 to IPv6 over IPv6). In addition, the rich set of Cisco IOS Software features (access control lists (ACLs), NAT, quality of service (QoS), NetFlow, encryption, etc.) can also be added if desired, to further extend this solution.

Scenario 2: IPv6 Internet Web-Presence

Scenario

Scenario 2 focuses on the IPv6 Transition use case in which an enterprise needs to establish an IPv6 Internet web presence. The enterprise has a compelling need to establish this IPv6 Internet web presence to demonstrate its continued leadership in the market. In addition, as IPv4 address depletion occurs, the enterprise is concerned that new users who can obtain only IPv6 access will be unable to utilize the enterprise's web-based services. CapEx and OpEx are not concerns; however, the solution must be implemented quickly and without disrupting the existing IPv4 web presence.

To accomplish this scenario, the enterprise might consider deploying native IPv6 Internet connectivity. This could be an option, but currently it is difficult to get new IPv6 circuits installed quickly, with installation times of 90 days or more being common. Since this time frame is unacceptable for this scenario due to the urgency of the desired deployment, the use of LISP is considered. Using LISP and a public LISP mapping service provider (reference [7] at the end of this document) allows the enterprise to establish an IPv6 web presence with its existing IPv4 WAN connectivity, and with few modifications to its current data center infrastructure. This LISP solution seems optimal in the business environment of the scenario. Even if it becomes a temporary solution, at least it meets the business goals of getting the enterprise IPv6 Internet web presence to market quickly and with minimal cost and impact on the existing network. However, LISP provides other benefits (such as ingress traffic engineering and mobility) that could be useful in making this a long-term solution as well.

Topology

Figures 4 and 5 illustrate simplified views of the starting topology considered in Scenario 2. In this case, the corporate site is illustrated, with its IPv4 WAN connectivity and data center that supports the existing IPv4 web presence. Note the current DNS A record and data center infrastructure. Figure 4 also shows the IPv6 Internet, and the public LISP mapping services that support the interworking between non-LISP and LISP sites in both the IPv4 and IPv6 domains. These components were briefly described earlier. (See the reference section at the end of this document for links to current LISP mapping service providers). What LISP accomplishes in this case is the interworking between the IPv6 Internet user-community (whether non-LISP or LISP), and the corporate site, using the existing IPv4 WAN links. Note that Figures 4 and 5 show a simplified view of the public LISP infrastructure components (MS/MR, PITR, and PETR). In reality, many of these devices are already deployed, and more are planned. (See reference [6] for additional information about the current LISP infrastructure.)

The final step is adding the IPv6 web presence, which is outside the LISP process. This step can be accomplished in several ways after the IPv6 packets are brought into the network using LISP. First, a DNS AAAA record is needed for IPv6 Internet users to find this new service. Next, the IPv6 packets delivered by LISP must be forwarded appropriately to the web service. Several options are available for handling this traffic, and the best approach will likely depend on the capabilities of the existing infrastructure. The following approaches are commonly used:

- **SLB IPv6-to-IPv4 load balancing:** SLBs are heavily used today for stateful IPv4-to-IPv4 load balancing. Many SLBs can run a dual-stack configuration and perform IPv6-to-IPv4 stateful load balancing of incoming IPv6 packets into an existing IPv4 web services infrastructure. Return packets flow from the web server to the SLB, and the reverse operation is applied. When LISP is used, this approach literally requires the deployment of just a single IPv6 VLAN between the LISP xTR and the SLB. An example of this type of deployment is illustrated in Figure 4.
- **IPv6 server deployment:** Another straightforward approach is to deploy a new server in IPv6 space with mirrored content and behind the dual-stack router running LISP. Again, this deployment requires very little IPv6 infrastructure, as illustrated in Figure 5. This is the example is the one discussed here in this scenario because it is the most generic.

- **Web server proxy support:** Many web server operating systems can perform proxying services that translate URLs in real-time. Assuming a dual-stack web server, this approach simply requires proxying the IPv6 URL into the IPv4 URL and then using the existing IPv4 web content. New web servers are not required. This approach is simply a variation of the IPv6 service deployment approach and mirrors the topology in Figure 5.

Figure 4. Scenario 2: Conceptual Enterprise Topology with SLB Option for IPv6 Web-Presence

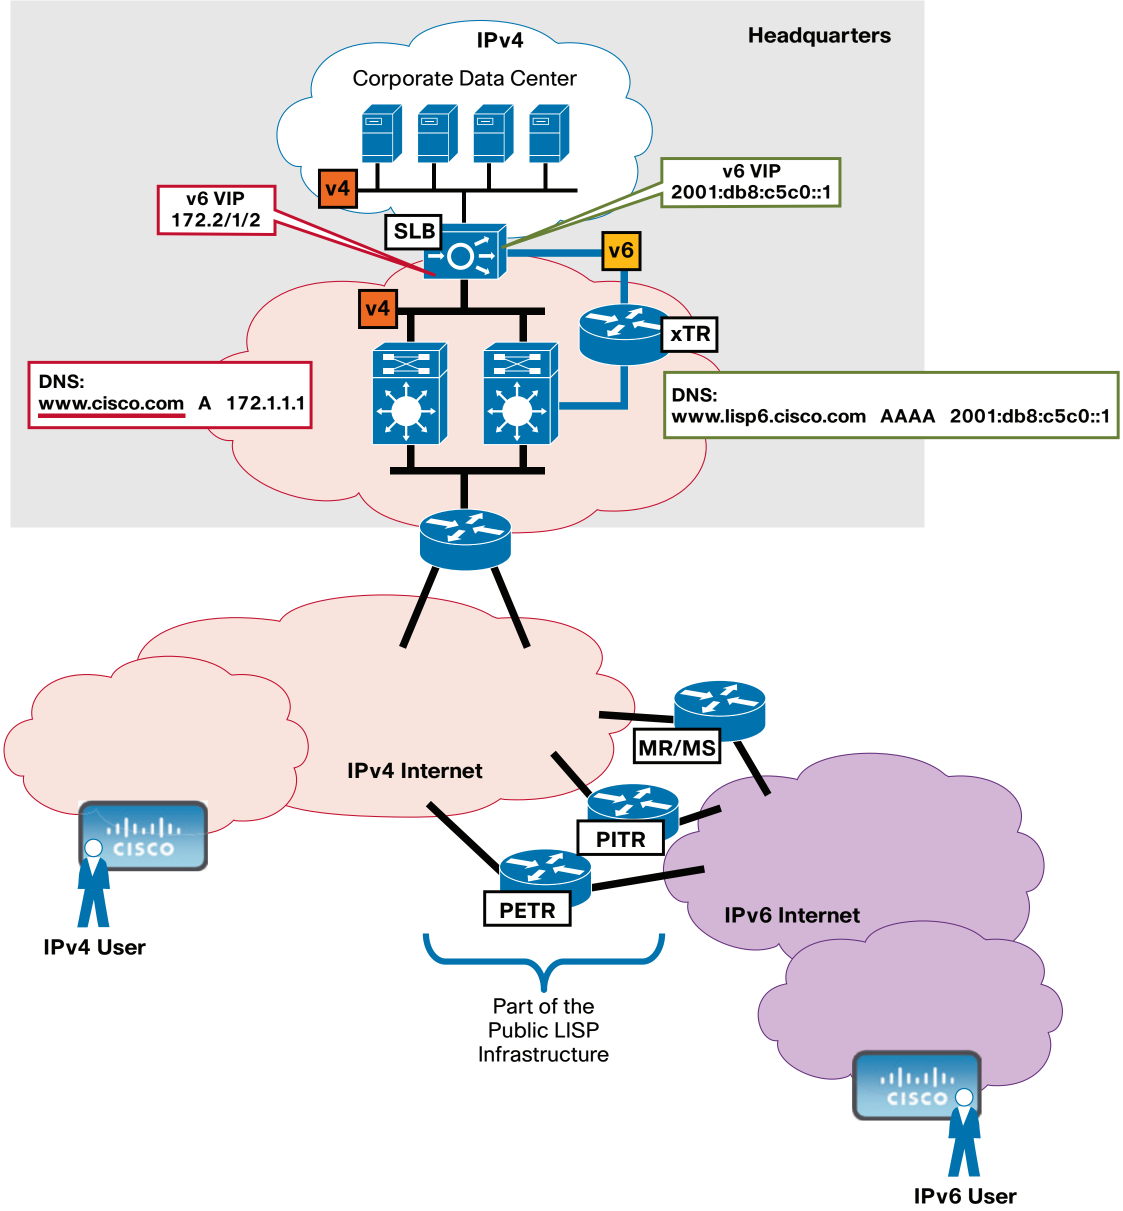

Figure 5. Scenario 2: Conceptual Enterprise Topology with IPv6 Server Deployment Option for IPv6 Web-Presence

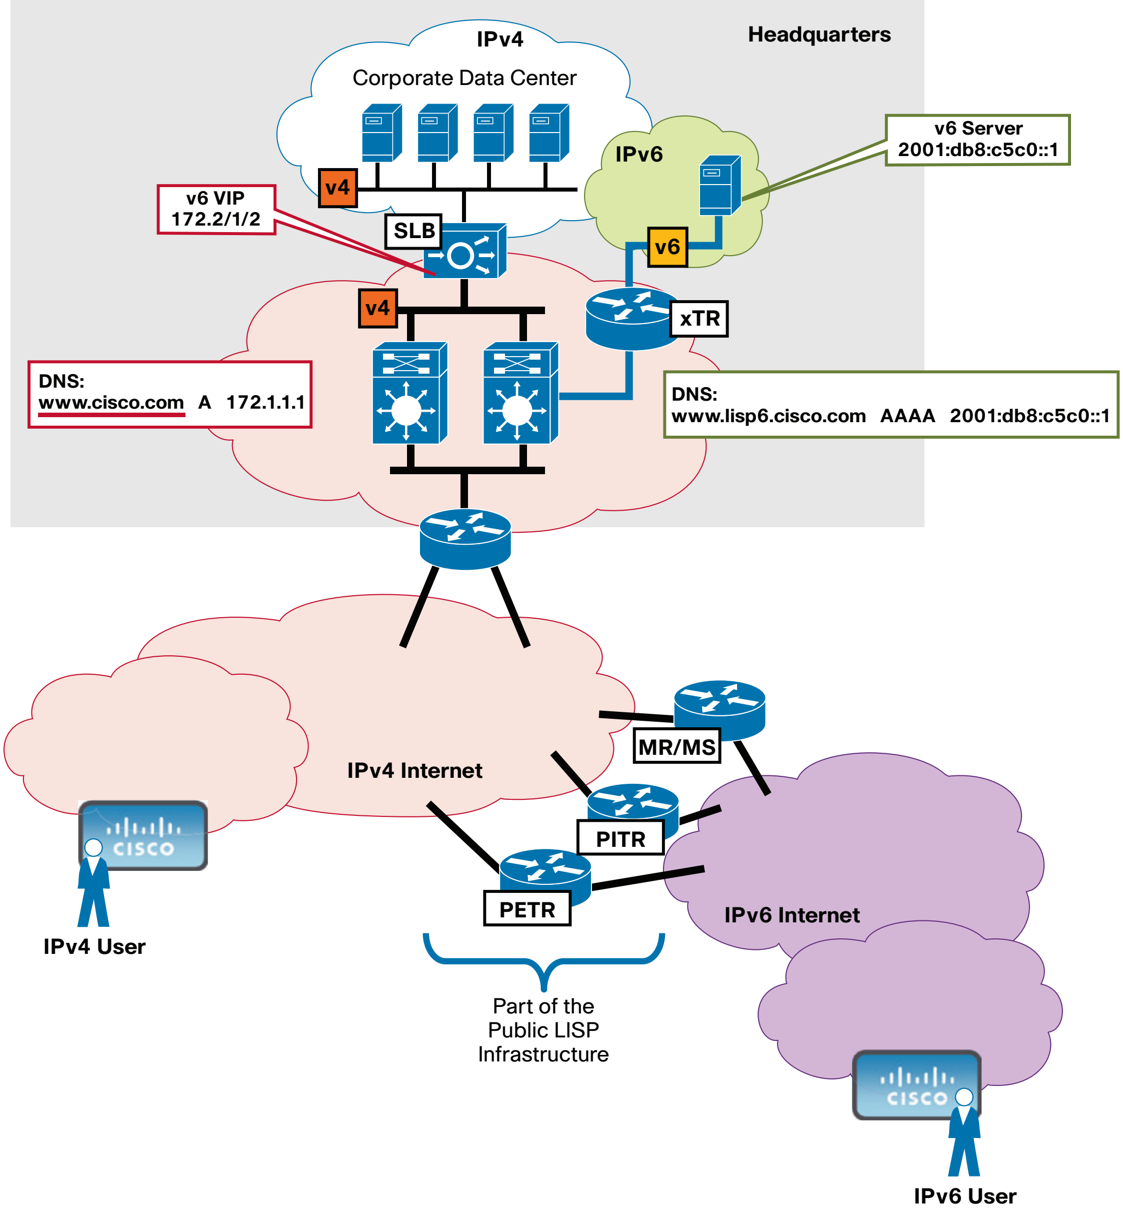

LISP Configurations

Figure 6 illustrates a simplified topology that reflects the basic elements of Figure 5. This figure provides a clear view of the LISP components used to accomplish IPv6 Internet connectivity to the data center web servers using the existing IPv4 WAN (Internet) connectivity. The base-requirement for the LISP router is that it has IPv4 Internet reachability (that is, its locator addresses (LISP RLOC) is assigned from the enterprise Internet address space), and that it registers to and uses public LISP mapping services.

Note: Just as in the first scenario, all addresses used in all figures and configuration examples are shown as IPv4 private (RFC 1918) and IPv6 documentation (RFC 3849) addresses as is standard practice in Cisco documentation. In this example, the IPv4 RLOC address (shown as 172.16.1.2) and PETR and PITR locator addresses (shown as 172.18.1.1 and 172.18.1.2 respectively) represent the appropriate routable IPv4 addresses for your environment (that is, the Internet routable addresses). In addition, the IPv6 EID address of the hypothetical client (shown as 2001:db8:c001::1) also represents an IPv6 routable address. The IPv6 address for the web server (shown as 2001:db8:c5c0::1) represents the real, DNS-advertised IPv6 address. This IPv6 address, assumed to be part of the LISP sites' IPv6 EID prefix, is registered with the LISP Map-Server (shown in Figure 5 but not in Figure 6), which in turn advertises this IPv6 EID prefix to the public LISP mapping system. As part of this process, the LISP PITR, which is also connected to the LISP mapping system, advertises a coarse-aggregate prefix covering this and other IPv6 LISP EID prefixes to the IPv6 Internet to attract non-LISP traffic that is destined to LISP sites.

In this example, from the Enterprise perspective, the following LISP functions are required:

- The data center LISP router (DC-XTR) is configured to provide LISP xTR services.
- The data center LISP router (DC-XTR) is configured to register with the public LISP mapping system (not shown, but assume it is located at 172.18.1.3) and to use the LISP interworking services provided by the PETR and PITR.
- A DNS AAAA record is added to advertise the IPv6 EID address associated with the web service URL (in this case, 2001:db8:c5c0::1).

Note: DNS replies can return both IPv4 and IPv6 addresses for a given URL, making it possible to simply add an AAAA record for an existing URL; however, this is typically not done. Instead, a new URL is used for the new IPv6 service, for several important reasons. First, the use of a new URL keeps the existing IPv4 service intact and limits the potential for disrupting the existing service. Second, when a DNS reply contains both IPv4 and IPv6 addresses, RFC 2874 specifies that hosts must try the IPv6 address first, and only when this process times out, try the IPv4 address. Some popular OSs are automatically enabled for IPv6, even when the underlying network is not. This causes long delays in DNS resolution (30 seconds or longer) while the client waits for a time out–an undesirable delay for end-users.

Figure 6. Scenario 2: IPv4 and IPv6 Addressing for LISP Configurations

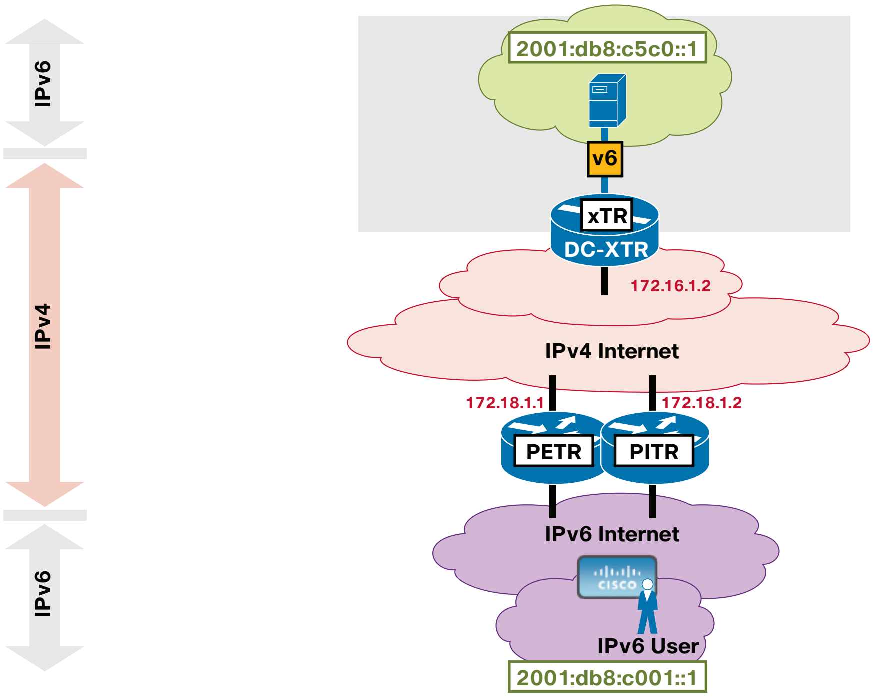

DC-XTR

The data center router DC-XTR provides LISP xTR services, and the relevant configuration elements are as shown here.

```
!
hostname DC-XTR
!
ip cef
no ip domain lookup
ipv6 unicast-routing
ipv6 cef
!
interface LISP0
!
interface Loopback0
 no ip address
 ipv6 address 2001:db8:c5c0::1/48
!
interface Ethernet0/0
 ip address 172.16.1.2 255.255.255.0
!
ipv6 lisp database-mapping 2001:db8:c5c0::/48 172.16.1.2 priority 1 weight 100
ipv6 lisp itr map-resolver 172.18.1.3
ipv6 lisp itr
ipv6 lisp etr map-server 172.18.1.3 key dc-s3cr3t
ipv6 lisp etr
ipv6 lisp use-petr 172.18.1.1
!
ip route 0.0.0.0 0.0.0.0 172.16.1.1
!
ipv6 route ::/0 Null0
!
```

As shown in the preceding DC-XTR configuration, the following LISP services are enabled:

- First, a loopback interface is configured with the IPv6 EID address 2001:db8:c5c0::1 simply as a way to represent the web services address.
- Next, the IPv6 EID prefix for this router (2001:db8:a::/48) is defined as being associated with the IPv4 locator (172.16.1.2), as referenced by the database-mapping command.
- Next, relevant IPv6 LISP ITR and ETR services are enabled.
- Then the LISP router is configured to register with the Map-Server (located at 172.18.1.3).
- Finally, the LISP router is configured to use the PETR (located at 172.18.1.1

The simplicity of the preceding LISP configuration is noteworthy given the capabilities that it enables. Even so, the configurations are fully functional for the topology shown in Figure 6.

Verification

Figure 7 illustrates the LISP packet flow corresponding to the IPv6 web services connectivity enabled for this scenario. As illustrated in Figure 7, test ping packets are sourced from the non-LISP IPv6 user located at 2001:db8:c001::1 and destined to the IPv6 web service located at the IPv6 LISP EID 2001:db8:c5c0::1. In this case, the following events occur:

- IPv6 packets sourced from the client are attracted to the PITR, which is dual-stacked, connected to the LISP mapping system, and advertising a coarse-aggregate prefix covering the data center LISP EID prefix into the IPv6 Internet to attract non-LISP traffic that is destined to LISP site.
- The PITR encapsulates these packets to DC-XTR using an IPv4 locator, as specified by the DC-XTR LISP policy.
- The LISP router DC-XTR decapsulates the IPv4 LISP header and forwards the IPv6 packet to the web server.
- Return traffic from the web server flows back through the DC-XTR, which is configured to LISP-encapsulate these IPv6 packets to the PETR using its IPv4 locator.
- The PETR decapsulates these (IPv4) packets, and forwards them natively within the IPv6 Internet to the IPv6 user.

Figure 7. Scenario 2: LISP Packet Flow for IPv6 over IPv4 Connectivity

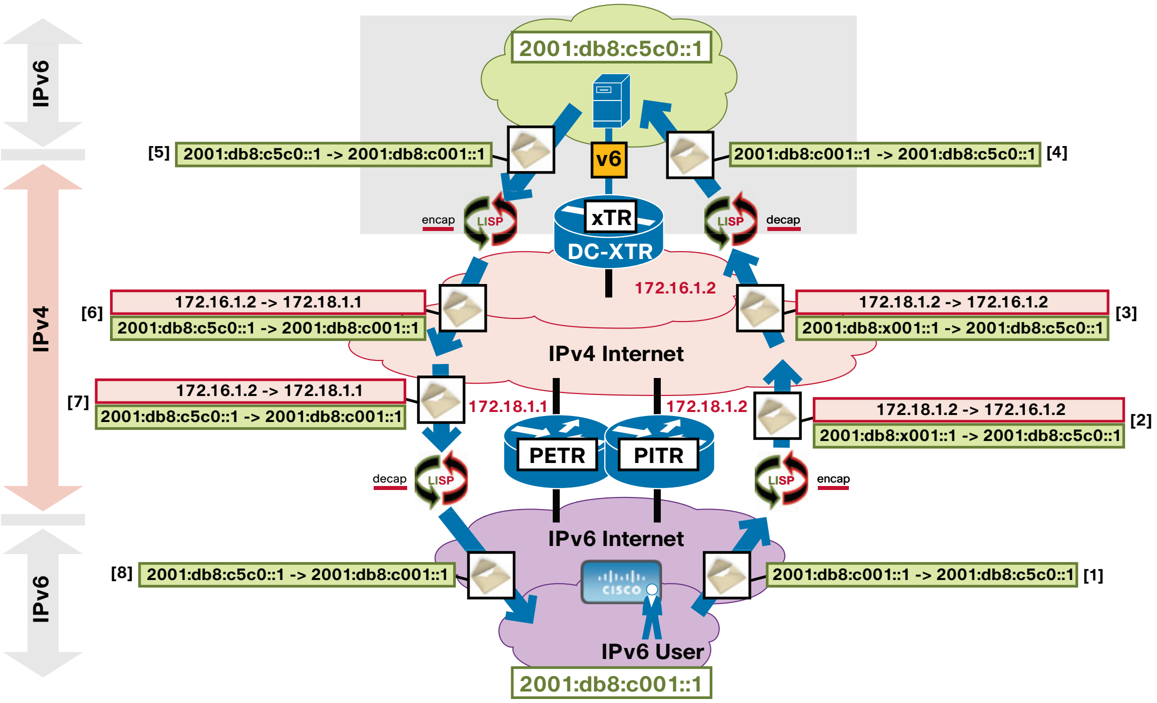

IPv6
IPv4
IPv6
2001:db8:c5c0::1
[5] 2001:db8:c5c0::1 -> 2001:db8:c001::1
v6
2001:db8:c001::1 -> 2001:db8:c5c0::1 [4]
encap
LISP
xTR
DC-XTR
LISP
decap
172.16.1.2
[6]
172.16.1.2 -> 172.18.1.1
2001:db8:c5c0::1 -> 2001:db8:c001::1
172.18.1.2 -> 172.16.1.2
2001:db8:x001::1 -> 2001:db8:c5c0::1
[3]
IPv4 Internet
[7]
172.16.1.2 -> 172.18.1.1
2001:db8:c5c0::1 -> 2001:db8:c001::1
172.18.1.1
172.18.1.2
172.18.1.2 -> 172.16.1.2
2001:db8:x001::1 -> 2001:db8:c5c0::1
[2]
PETR
PITR
decap
LISP
LISP
encap
IPv6 Internet
[8] 2001:db8:c5c0::1 -> 2001:db8:c001::1
2001:db8:c001::1 -> 2001:db8:c5c0::1 [1]
IPv6 User
2001:db8:c001::1

The output shown here illustrates this packet validation, as well as other useful LISP information.

IPv6 User

- As verification of end-to-end connectivity, this output shows a sourced-ping from the (representation of the) IPv6 User to the data center LISP IPv6 EID address.

```
IPv6User#ping 2001:db8:c5c0::1 so 2001:db8:c001::1 repeat 50

Type escape sequence to abort.
Sending 50, 100-byte ICMP Echos to 2001:DB8:C5C0::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:C001::1
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 96 percent (48/50), round-trip min/avg/max = 0/1/4 ms
IPv6User#
```

- In addition, **show** output verifies the LISP map-cache on DC-XTR. You can see in this case that the forwarding action for packets destined to non-LISP sites is to encapsulate using the PETR.

```
DC-XTR#show ipv6 lisp map-cache
LISP IPv6 Mapping Cache, 2 entries

::/0, uptime: 00:03:33, expires: never, via static
  Negative cache entry, action: send-map-request
2001:DB8:8000::/33, uptime: 00:01:06, expires: 00:13:50, via map-reply, forward-native
  Encapsulating to proxy ETR
DC-XTR#
```

Scenario 2 Recap

In Scenario 2, LISP was used to connect non-LISP IPv6 Internet users to corporate web services. This connection was achieved by deploying IPv6 on one internal subnet and enabling IPv6 on a web sever (alternatively, the SLB and web proxy methods could be used), and by using LISP to carry IPv6 user-traffic over the existing IPv4 core network. The existing IPv4 routable infrastructure and WAN connectivity was used without any changes. One new router was deployed, which included the IPv6 subnet and with a few lines of LISP configuration. Public LISP mapping services (including a MS and PETR and PITR services) were also used. With these few configurations, all accomplished within a few hours, LISP was able to fully connect non-LISP IPv6 Internet users to the corporate web services. This same LISP solution supports other traffic types as well (not just web traffic). For example, employees with native IPv6 connectivity at home or on the road can access normal corporate IT services (email, calendaring, etc.) using their normal SSL VPN client through this same LISP solution.

Conclusions

LISP implements a new routing architecture that is designed for a much broader purpose than IPv6 transition. Because it was designed from inception to be address family agnostic, it can be used to transparently support many IPv6 transition needs in a very natural manner. This capability has been proven by real success stories (such as those detailed in references [1], [2], and [3] at the end of this document) in practical, production implementations to be ideal for many IPv6 transition strategies. Scenarios incorporating LISP into an IPv6 transition strategy have demonstrated quick deployment time, low deployment and operational cost, little or no need for additional equipment or modifications, and high user-satisfaction.

The use-cases described in this document are summaries based on these success stories. They demonstrate that the economics of incorporating Locator/ID Separation Protocol (LISP) into your IPv6 transition strategy can be compelling. The broader features of LISP (inherent multi-homing, ingress traffic engineering, and mobility) add to the benefits of using LISP.

References

[1] "LISP Deployments at Facebook," NANOG50, Donn Lee, Facebook, at http://nanog.org/meetings/nanog50/presentations/Tuesday/NANOG50.Talk9.lee_nanog50_atlanta_oct2010_007_publish.pdf

[2] "IPv6 at Facebook", Google IPv6 Implementors Conference 2010, Donn Lee, Facebook, at http://sites.google.com/site/ipv6implementors/2010/agenda/06_Lee_IPv6atFacebookgoogleconference2010.pdf

[3] "Cisco IPv6 strategy and roadmap", Google IPv6 Implementors Conference 2010, Mark Townsley, Cisco, at http://sites.google.com/site/ipv6implementors/2010/agenda/01_townsley-google-2010-Final.pdf

[4] LISP Documentation: LISP Command Reference Guide, LISP Configuration Guide, and LISP Lab Test Guide, at http://lisp4.cisco.com and http://lisp6.cisco.com

[5] Cisco marketing information about LISP: http://www.cisco.com/go/lisp

[6] LISP Beta Network information: http://www.lisp4.net and http://www.lisp6.net

[7] Current list of LISP mapping service providers: http://lisp4.cisco.com and http://lisp6.cisco.com

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.**