



Locator ID Separation Protocol (LISP) Overview

Locator ID Separation Protocol (LISP) is a network architecture and protocol that implements the use of two namespaces instead of a single IP address:

- Endpoint identifiers (EIDs)—assigned to end hosts.
- Routing locators (RLOCs)—assigned to devices (primarily routers) that make up the global routing system.

Splitting EID and RLOC functions yields several advantages including improved routing system scalability, and improved multihoming efficiency and ingress traffic engineering.

LISP functionality requires LISP-specific configuration of one or more LISP-related devices, such as the LISP egress tunnel router (ETR), ingress tunnel router (ITR), proxy ETR (PETR), proxy ITR (PITR), map resolver (MR), map server (MS), and LISP alternative logical topology (ALT) device.

- [Finding Feature Information, on page 1](#)
- [Prerequisites for Configuring LISP, on page 1](#)
- [Restrictions for Configuring LISP, on page 2](#)
- [Information About Configuring LISP, on page 2](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Configuring LISP

Before you can configure Locator/ID Separation Protocol (LISP), you will need to determine the type of LISP deployment you intend to deploy. The LISP deployment defines the necessary functionality of LISP devices, which, in turn, determines the hardware, software, and additional support from LISP mapping services and proxy services that are required to complete the deployment.

LISP configuration requires the datak9 license.

Restrictions for Configuring LISP

- LISP is not supported on Tunnels.
- Management traffic generated on a LISP xTR with the source of a LISP EID interface does not work because management traffic such as SSH or telnet are not LISP aware. To make management protocols LISP aware, you need to create a static route pointing towards correct next hop. The static route should have a next hop of the LISP virtual interface and IP of the RLOC for the remote ETR. Management traffic generated on an xTR from a LISP EID interface needs this route inserted in the routing table as a workaround of this limitation.

Information About Configuring LISP

LISP Functionality Overview

Problem

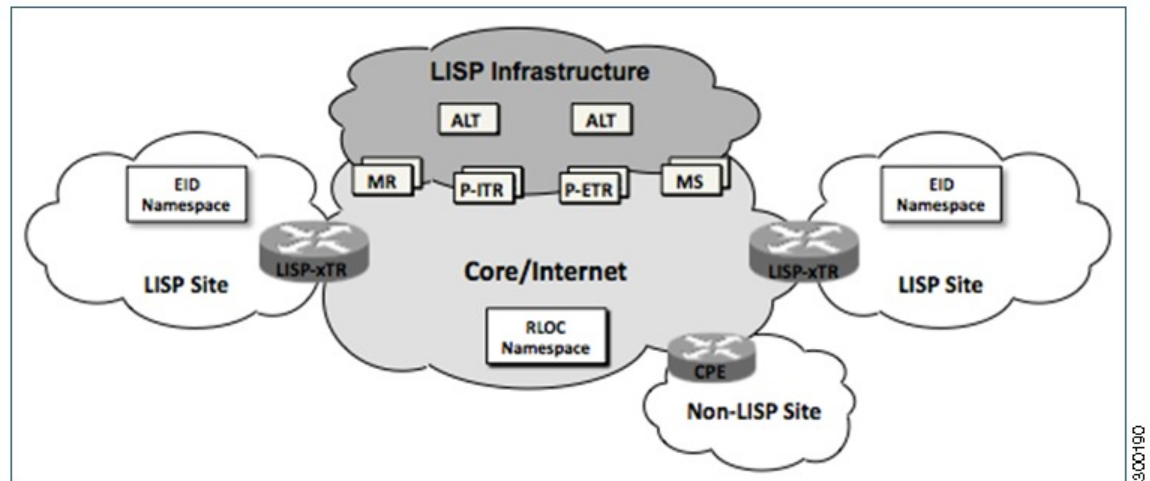
The continuous growth of the Internet presents a number of challenges. Among the most fundamental of these challenges is ensuring that the routing and addressing system continues to function efficiently even as the number of connected devices continues to increase. A basic observation during early network research and development work was that the single IP address, which includes both identity and location, leads to suboptimal route scaling and hinders multihoming and device mobility.

Solution

Locator ID Separation Protocol (LISP) provides improved routing scalability and facilitates flexible address assignment for multi-homing, provider independence, mobility, and virtualization. LISP offers an alternative to traditional Internet architecture by introducing two separate IP addresses: one to indicate routing locators (RLOCs) for routing traffic through the global Internet and a second address for endpoint identifiers (EIDs) used to identify network sessions between devices.

The figure below displays a general overview illustration of a LISP deployment environment, including the three essential environments that exist in a LISP environment: LISP sites (EID namespace), non-LISP sites (RLOC namespace), and LISP mapping service (infrastructure).

Figure 1: LISP Deployment Environment



As illustrated in the figure, the LISP EID namespace represents customer end sites in the same way that end sites are defined in non-LISP environments with one difference: The IP addresses used within these LISP sites are not advertised within the non-LISP Internet (RLOC namespace). Instead, end-customer LISP functionality is deployed exclusively on customer endpoint routers, which perform both the egress tunnel router (ETR) and ingress tunnel router (ITR) functions of a LISP device (abbreviated as xTR in the figure).

To fully implement LISP with support for mapping services and Internet interworking may require additional LISP infrastructure components as part of the deployment. As displayed in the figure above, these additional LISP infrastructure components include devices that function in the LISP roles of map resolver (MR), map server (MS), proxy egress tunnel router (PETR), proxy ingress tunnel router (PITR), and LISP alternative logical topology (ALT) device.

LISP Network Element Functions

The LISP architecture defines seven LISP-specific network infrastructure components. In some cases, a single physical device can implement more than one of these logical components. For more information, refer to the descriptions of the LISP components described in the following sections:

LISP Alternative Logical Topology

An alternative logical topology (ALT) device (not present in all mapping database deployments) connects through generic routing encapsulation (GRE) tunnels and border gateway protocol (BGP) sessions, map resolvers, map servers, and other ALT routers. The only purpose of ALT routers is to accept EID (Endpoint Identifier) prefixes advertised by devices that form a hierarchically distinct part of the EID numbering space and then advertise an aggregated EID prefix that represents that distinct space to other parts of the ALT. Just as in the global Internet routing system, this aggregation is performed to reduce the number of prefixes that need to be propagated throughout the entire network. An MS or combined MR/MS may also be configured to perform the functions of an ALT router.

LISP Egress Tunnel Router

An ETR connects a site to the LISP-capable part of a core network (such as the Internet), publishes EID-to-RLOC mappings for the site, responds to Map-Request messages, and decapsulates and delivers

LISP-encapsulated user data to end systems at the site. During operation, an ETR sends periodic Map-Register messages to all its configured map servers. The Map-Register messages contain all the EID-to-RLOC entries for the EID-numbered networks that are connected to the ETR's site.

An ETR that receives a Map-Request message verifies that the request matches an EID for which it is authoritative, constructs an appropriate Map-Reply message containing its configured mapping information, and sends this message to the ingress tunnel router (ITR) whose RLOCs are listed in the Map-Request message. An ETR that receives a LISP-encapsulated packet that is directed to one of its RLOCs decapsulates the packet, verifies that the inner header is destined for an EID-numbered end system at its site, and then forwards the packet to the end system using site-internal routing.

The ETR function is usually implemented in the customer premises equipment (CPE) router and does not require hardware changes on software-switched platforms, such as a Cisco Integrated Services Router (ISR). The same CPE router will often provide both ITR and ETR functions and, when doing so, is referred to as an xTR.

LISP Ingress Tunnel Router (ITR)

An ITR is responsible for finding EID-to-RLOC mappings for all traffic destined for LISP-capable sites. When the ITR receives a packet destined for an EID, it first looks for the EID in its mapping cache. If the ITR finds a match, it encapsulates the packet inside a LISP header with one of its RLOCs as the IP source address and one of the RLOCs from the mapping cache entry as the IP destination. The ITR then routes the packet normally.

If no entry is found in the ITR's mapping cache, the ITR sends a Map-Request message to one of its configured map resolvers and then discards the original packet. When the ITR receives a response to its Map-Request message, it creates a new mapping cache entry with the contents of the Map-Reply message. When another packet, such as a retransmission for the original and, now, discarded packet arrives, the new mapping cache entry is used for encapsulation and forwarding.



Note Sometimes the Map-Reply message will indicate that the destination is not an EID. When this happens, a negative mapping cache entry is created, which causes packets to either be discarded or forwarded natively when the packets match that cache entry.

Like the ETR, an ITR is usually implemented in a LISP site's customer premises equipment (CPE) router, which is typically configured as an xTR (performs functions of both ETR and ITR components).

LISP Map Resolver

Like an MS, a LISP MR connects to the ALT. The function of the LISP MR is to accept encapsulated Map-Request messages from ingress tunnel routers (ITRs), decapsulate those messages, and then forward the messages to the MS responsible for the egress tunnel routers (ETRs) that are authoritative for the requested EIDs.

When an MR is implemented concurrently with an MS in a private mapping system deployment, the concurrent MS forwards the encapsulated Map-Request messages to the authoritative ETRs. When a LISP ALT is present in the deployment, the MR forwards the Map-Request messages directly over the ALT to the MS responsible for the ETRs that are authoritative for the requested EIDs. An MR also sends Negative Map-Replies to ITRs in response to queries for non-LISP addresses.

LISP Map Server

An MS implements part of the distributed LISP mapping database by accepting registration requests from its client egress tunnel routers (ETRs), aggregating the successfully registered EID prefixes of those ETRs, and advertising the aggregated prefixes into the alternative logical topology (ALT) with border gateway protocol (BGP).

In a small private mapping system deployment, an MS may be configured to stand alone (or there may be several MSs) with all ETRs configured to register to each MS. If more than one, all MSs have full knowledge of the mapping system in a private deployment.

In a larger or public mapping system deployment, an MS is configured with a partial mesh of generic routing encapsulation (GRE) tunnels and BGP sessions to other map server systems or ALT routers. For these deployments, ETRs need to register to only one MS (or a few if redundancy is desired) and an ALT device is used to ensure that the entire LISP mapping system is available to all MS and MR devices.

Because an MS does not forward user data traffic—it handles only LISP control plane traffic—it does not require high performance switching capability and is well suited for implementation on a general purpose router, such as a Cisco Integrated Services Router (ISR). Both MS and MR functions are typically implemented on the same device, which is referred to as an MR/MS device.

LISP Proxy ETR

A LISP PETR implements ETR functions on behalf of non-LISP sites. A PETR is typically used when a LISP site needs to send traffic to non-LISP sites but the LISP site is connected through an access network of a service provider that does not accept nonroutable EIDs as packet sources.

When dual-stacked, a PETR may also serve as a way for EIDs and RLOCs to communicate in a LISP site that contains EIDs in one address family and RLOCs in a different address family. A dual-stacked PETR also provides multiaddress family support for LISP EIDs within one address family to be able to communicate with non-LISP destinations in the same address family over a core network within a different address family.

Example

A LISP site with IPv4-only RLOC connectivity can send IPv6 EIDs within an IPv4 LISP header across the IPv4 Internet to a dual-stacked PETR where the packets are decapsulated and then forwarded natively to non-LISP IPv6 Internet sites.

The PETR function is commonly configured on a device that also functions as a PITR. A device that functions as both a PETR and a PITR is known as a PxTR. Additionally, a PETR carries LISP data plane traffic and can be a high packet-rate device. To take advantage of this high packet-rate capability, deployments typically include hardware-switched platforms or high-end Cisco Integrated Services Routers (ISRs).

LISP Proxy ITR

A LISP PITR implements ITR mapping database lookups and LISP encapsulation functions on behalf of non-LISP-capable sites. PITRs are typically deployed near major Internet exchange points (IXPs) or in ISP networks to allow non-LISP customers from those networks to connect to LISP sites. In addition to implementing ITR functionality, a PITR also advertises some or all of the non-routable EID prefix space to the part of the non-LISP-capable Internet that it serves so that the non-LISP sites will route traffic toward the PITR for encapsulation and forwarding to LISP sites.



Note PITR advertising of nonroutable EID prefix space is intended to be highly aggregated with many EID prefixes represented by each prefix that is advertised by a PITR.

Like the PETR, when dual-stacked, the PITR also provides multiple-address family support. But the PITR supports transport of non-LISP traffic from one address family to LISP sites in the same address family over a core network within a different address family.

Example

A LISP site with IPv4-only RLOC connectivity can take advantage of a dual-stacked PITR to allow non-LISP IPv6 Internet users to reach IPv6 EIDs across the IPv4 Internet.

The PITR function is commonly configured on a device that also functions as a PETR. A device that functions as both a PETR and a PITR is known as a PxTR. Additionally, a PITR carries LISP data plane traffic and can be a high packet-rate device. To take advantage of this high packet-rate capability, deployments typically include hardware-switched platforms or high-end Cisco® Integrated Services Routers (ISRs).

Feature Information for LISP Overview

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for LISP Overview

Feature Name	Releases	Feature Information
LISP Overview	15.1(4)M Cisco IOS XE Release 3.3.0S	The LISP Overview feature provides a general overview of LISP and its components. The following LISP components are supported: <ul style="list-style-type: none"> • Egress tunnel router (ETR) • Ingress tunnel router (ITR) • LISP alternative logical topology (ALT) device • Map resolver (MR) • Map server (MS) • Proxy ETR (PETR) • Proxy ITR (PITR)

Feature Name	Releases	Feature Information
LISP, SHA-2 support for site registration	15.3(2)T Cisco IOS XE Release 3.9S	<p>LISP can be configured to use SHA2-based HMAC algorithm for integrity-checking LISP site registration messages. Prior to this release, only SHA1-based HMAC algorithm was supported.</p> <p>The following commands were modified:</p> <ul style="list-style-type: none">• ipv4 etr map-server• ipv6 etr map-server

