

Locator/ID Separation Protocol (LISP)

Q. What is LISP?

A. LISP stands for Locator/ID Separation Protocol and is a next-generation IP routing feature that creates a new paradigm in how IP addressing is assigned and interpreted by splitting the device identity, known as an endpoint identifier (EID), and its location, known as its routing locator (RLOC), into two different namespaces. Creating separate IP addresses for EID and RLOC functions yields several advantages, including improved scalability of the routing system through greater aggregation of RLOCs and improved multihoming efficiency and ingress traffic engineering.

Q. Why was LISP developed?

A. The current Internet routing and addressing architecture uses a single namespace, the IP address, to simultaneously express two functions about a device: its identity, and its location within the network. One very visible and detrimental result of this single namespace has been manifested in the rapid growth of the Internet's DFZ (default-free zone) as a consequence of multihoming, traffic engineering, nonaggregatable address allocations, and business events such as mergers and acquisitions. Separating the locator and the identifier has been discussed for many years as a way to greatly reduce the size of the Internet DFZ. LISP comprises the development of protocol specifications for the IETF and several implementations that run on different routers and/or hosts.

Q. Is LISP a standard? Does Cisco own LISP?

A. LISP standards are currently being developed within the IETF LISP Working Group. In addition to leading this standards effort, Cisco is also developing LISP software for both Cisco IOS® and NX-OS platforms, as well as gaining working experience with LISP through the deployment and operation of a public LISP network. The Cisco® contributions to LISP are being developed as open standards, with no Cisco intellectual property rights (IPR). We feel it is in the best interest of the overall Internet community to resolve the issues facing the Internet today. A healthy and growing Internet benefits everyone, including Cisco. Cisco has constantly sought to stimulate outside interest and development efforts from the outset. Other LISP development is being pursued by the open source community, several competitive companies, and many researchers and universities.

Q. How is LISP deployed?

A. Cisco's philosophy for the development of LISP has been to minimize end-customer changes and deployment complexities. From the LISP perspective, there are essentially two areas to be considered:

- **LISP sites:** LISP sites use IP addresses in the EID namespace to address hosts and in Domain Name System (DNS) in exactly the same way they are currently used. These addresses are not advertised within the non-LISP RLOC namespace (that is, the Internet), but instead are advertised by the LISP mapping services. The LISP site router supports the LISP functionality of Ingress Tunnel Router (ITR) and Egress Tunnel Router (ETR).
- **LISP infrastructure:** To fully implement LISP with Internet scale and interoperability between LISP and non-LISP sites, additional LISP Infrastructure components are required to support the LISP EID-to-RLOC mapping services (much like DNS provides the URL-to-IP address mapping function today) and LISP to non-LISP interworking. These LISP infrastructure devices include the Map-Server (MS), Map-Resolver (MR), Proxy Ingress Tunnel Router (PITR), and Proxy Egress Tunnel Router (PETR) devices. These infrastructure services are most efficiently provided as a mapping service in a fashion similar to how DNS services are provided today.

Q. What are the various logical components of LISP?

A. The specific devices used by LISP and their functions are as follows.

- **LISP site devices:**

- **ITR:** An Ingress Tunnel Router receives packets from site-facing interfaces and either LISP-encapsulates packets to remote LISP sites, or natively forwards packets to non-LISP sites. An ITR also sends map requests and processes received map replies in order to resolve EID-to-RLOC mappings.
- **ETR:** An Egress Tunnel Router receives packets from core-facing interfaces, de-encapsulates them, and delivers them to local EIDs at the site. An ETR also registers its EID prefixes and RLOCs with the Map-Server, and responds to map requests received from the Map-Server.

In most cases, a LISP site edge device will implement both ITR and ETR functions, in which case it is commonly referred to as an **xTR**.

- **LISP infrastructure devices:**

- **MS:** A Map-Server advertises EID prefixes into the LISP–Alternative Topology (ALT) for ETRs that are registered to it. The MS also receives map requests over the ALT from ITRs seeking mapping resolutions for EID prefixes and forwards them to the registered ETR that is authoritative for the EID prefix being queried.
- **MR:** A Map-Resolver receives map requests from ITRs and forwards map requests to the ALT. The MR also sends negative map replies to ITRs in response to queries for non-LISP addresses.
- **ALT:** An Alternative Topology device is deployed to build out an overlay network that provides a mechanism for managing EID prefix aggregation. Other LISP infrastructure devices (MS, MR, PITR, PETR) connect to the ALT to participate in EID-to-RLOC mapping services:
- **PITR:** A Proxy ITR provides connectivity between non-LISP sites and LISP sites by advertising coarse-aggregate prefixes for the LISP EID namespace into the Internet DFZ (RLOC namespace) and forwarding this non-LISP traffic to LISP sites.
- **PETR:** A Proxy ETR allows IPv6 LISP sites without native IPv6 RLOC connectivity to reach LISP sites that only have IPv6 RLOC connectivity. In addition, the PETR also allows LISP sites with Unicast Reverse Path Forwarding (URPF) restrictions to reach non-LISP sites.

Q. Which LISP components are supported by Cisco?

A. Cisco supports all of the LISP components at this time. Due to the variations in requirements imposed by each device's needs, certain LISP devices are supported by different Cisco platforms. Always refer to the LISP Release Notes for current information on supported platforms and LISP functions, located at http://www.cisco.com/en/US/docs/ios/15_1/release/notes/1511xbrn.html.

Q. How does a LISP-enabled site communicate with another LISP-enabled site?

A. When an ITR receives a packet from a site-facing interface, it first does a destination lookup, with the possibility of one of two outcomes. If the lookup returns a nondefault destination, the ITR forwards the packet to that destination natively since the destination is not another LISP site. If the lookup returns a match against the default route, the ITR then checks the LISP map cache for an existing EID-to-RLOC mapping. If a mapping exists, the packet is encapsulated using that map-cache policy and forwarded. If no mapping exists, the ITR sends a map request for the destination EID in question to its configured Map-Resolver. This map request is forwarded over the ALT to the Map-Server to which the destination ETR is registered, and finally to the destination ETR. The ETR responds directly to the ITR with a map reply, which the ITR adds to its map cache.

The ITR can now forward LISP packets between its EID and the destination EID. When the LISP ETR receives the LISP packet, it de-encapsulates it, and then forwards it to the original (EID) destination IP address.

Q. How does a LISP-enabled site communicate with a non-LISP site?

- A.** When an ITR receives a packet from a site-facing interface, it first does a destination lookup, and because the destination is for a non-LISP site, the lookup returns a nondefault destination and the ITR forwards the packet natively (without LISP encapsulation). Because the destination IP address is known within the Internet DFZ, the packet is routed natively to the destination. Return traffic must be handled by a Proxy Ingress Tunnel Router. A PITR advertises coarse-aggregate prefixes for the LISP EID namespace into the Internet DFZ, thus attracting non-LISP traffic destined to LISP sites. This traffic is then encapsulated and forwarded to the LISP site in just the same way an ITR communicates with a LISP site.

Q. How does LISP work in a mixed IPv4/IPv6 environment?

- A.** LISP provides support for both IPv4 and IPv6 EIDs and RLOCs. Because of this, it can function in any of four ways:
- IPv4-to-IPv4 over IPv4
 - IPv4-to-IPv4 over IPv6
 - IPv6-to-IPv6 over IPv6
 - IPv6-to-IPv6 over IPv4

LISP does not translate between the IPv4 and IPv6 EID namespaces. Interworking between LISP and non-LISP sites requires the PITR/PETR infrastructure.

Q. How does LISP handle the maximum transmission unit (MTU)?

- A.** LISP encapsulation adds 36 bytes (IPv4) or 56 bytes (IPv6) to the original packet size. As with other protocols that encapsulate or tunnel traffic (for example, IP Security [IPsec] and generic routing encapsulation [GRE]), it is possible that the resultant encapsulated packets could exceed the permissible MTU somewhere along the transit path. LISP provides both stateful and stateless mechanisms for handling potential MTU issues, with the main goals being (1) preventing packets from being dropped, and (2) preventing the need for ETRs to perform packet reassembly prior to LISP de-encapsulation.

Q. Is LISP compatible with Network Address Translation (NAT)?

- A.** Yes, LISP can be deployed within a NAT environment, and LISP and NAT can be performed either on the same platform or on separate platforms. In either case, the same process takes place. When private addresses are used on hosts behind NAT, these private addresses are first translated into global addresses within the LISP EID namespace. These packets then follow the normal LISP encapsulation process, with the outer (LISP) header using the RLOC address. Thus, the outside world only sees these global addresses in both namespaces: the EIDs in the inner header, and the RLOCs in the outer header. The original, private addresses are held only within the NAT state.

Q. What are some key use cases for LISP?

- A.** Although LISP was initially conceived to address Internet DFZ scalability issues, once the concept of splitting the device identity and location is available, many potential benefits are possible. Some of the important use cases being developed include:
- Low operating-expense (OpEx) multihoming with site-based policy control (ingress tunnel engineering)
 - Low OpEx site independence with renumbering avoidance while scaling the Internet DFZ
 - Data center virtual machine mobility

- Fast mobility (LISP mobile node client)
- Connecting IPv6 islands (IPv6 transition)

Q. What benefits does LISP bring to multihomed environments?

A. Multihomed sites traditionally require the use of external Border Gateway Protocol (eBGP) to advertise site prefixes to multiple upstream providers. The OpEx is often high due to configuration complexity and significant BGP parameter tuning requirements (with a considerable expertise level) to achieve basic traffic manipulation. By employing LISP in a multihomed environment, the need for eBGP peering with upstream providers is eliminated, and administrators can set the ingress traffic policy directly through the mapping parameters. Thus, LISP simplifies multihoming and allows for greater control of the ingress traffic policy.

Q. What benefits does LISP bring to single-homed environments?

A. The deployment of LISP on customer premises equipment (CPE) router can benefit even devices with a single connection to the Internet. The most obvious and direct benefit is that in this case LISP gives provider independence. If the site wishes to change providers, only the RLOC need change. All other site addresses (EIDs) and DNS entries remain unchanged, eliminating the need for renumbering. In addition, this change can be made seamlessly by adding the new RLOC to the site, adding it to the LISP policy, and then removing the old RLOC from the LISP policy. Thus, no traffic disruptions need occur during the transition. Since the EID prefix is never in the Internet DFZ, this also helps scale the core.

Q. What benefits does LISP bring to the data center?

A. The deployment of LISP in the data center enables an architectural paradigm shift by using local instances of the LISP mapping database and its infrastructure components for server reachability instead of traditional, topologically dependent, hop-by-hop routing protocols. This lowers the network OpEx by making the core much simpler by allowing it to just move packets. In addition, and more importantly, because the servers' identifiers (EIDs) are separated from their location (RLOC), "movement" can now occur by simply changing the locator while keeping the EID fixed. This same feature of LISP that enables you to change service providers also allows you to change the location of a device or cloud while keeping the identifiers fixed. This, for example, allows TCP connections to remain established during these changes. Virtual machine mobility is enabled by this same concept. All this can be done because the identifier is not in the underlying routing system.

Q. What about fast mobility? Does Cisco plan to implement a LISP mobile node client?

A. Yes. We are currently working on a prototype implementation of a LISP mobile node client for an open-standards handset. Cisco has approached several handset vendors about enabling LISP in their products. More information will be coming during 2010.

Q. How can LISP help me in my transition to an IPv6 environment?

A. LISP can help in a number of ways. First, LISP provides mechanisms for encapsulating IPv6 host packets using IPv4 locators (or IPv4 host packets and IPv6 locators). Thus, you can build IPv6 islands and connect them using your existing IPv4 Internet connectivity. This may allow you to gain experience with IPv6 before investing in additional infrastructure. In addition, when the LISP infrastructure also includes the use of a PETR that has both IPv4 and IPv6 connectivity into the Internet, your LISP IPv6 site can also connect with non-LISP IPv6 sites.

Q. Where can I find more information about LISP?

A. Visit the Cisco LISP product information page at <http://www.cisco.com/go/lisp>.



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.

CCDE, CCENT, CCSI, Cisco Eos, Cisco Explorer, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco TrustSec, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco.Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLYNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (1002R)