Network-Based Protocol Innovations in Secure Encryption Environments

Using Locator/ID Separation Protocol (LISP) to Optimize Routing in IP Encryption Environments

Craig Hill
Distinguished Systems Engineer
U.S. Federal Area
crhill@cisco.com

Individual Contributors:
Tim Thomas – Sr. Systems Engineer, U.S. Federal
Dino Farinacci – Cisco Fellow

Abstract

This paper examines how recent network-based protocol innovations can be used to simplify the overall deployment, functionality, and operation of networks in which IPsec VPN devices (IVDs) are required—specifically, in Department of Defense (DoD), Intelligence Community (IC), and secure enterprise networks.

The main innovation addressed in this paper is the Locator/ID Separation Protocol (LISP) routing architecture framework, which, when deployed in an IVD environment, addresses some of the major deployment and operational challenges common in IVD networks today.

This paper compares the use of LISP with technologies currently used in IVD deployments today. It also highlights the advantages and differentiation that simplify operations and deployment, and offers current and future enhancements that could change the way IVD networks are deployed.

Please note that LISP has published drafts within the Internet Engineering Task Force (IETF) and will continue to target full standards-based approach moving forward.

Problem Statement

In secure federal and enterprise customers, IP encryption is a popular security protocol for core networks and for hub-and-spoke topologies in which there is a need to backhaul remote site traffic to a single aggregation site.

This IPsec VPN device (IVD) supports flexible packet-based encryption at the IP layer, allowing network designers to take advantage of IPv4/v6 packet-based transport, specifically on the "unsecure" side of the network transport. A typical IVD encrypts the received packet (from the secure/clear-text interface), encapsulates it into an IPv4/v6 transport packet, and forwards it to the destination IVD (which has already executed its key exchange process, trusted secure association, and IP route forwarding establishment).

Note: In this context, "secure" refers to routers and locations within the encryption boundary of the IVD. "Unsecure" refers to routers and locations that are outside the encryption boundary, and therefore not secure but the data is encrypted.

The receiving IVD then decrypts/de-encapsulates the received packet and forwards it to the attached receiving device. Because it is IP-based and offers the flexibility for any-to-any IP communications, as well as multiple options for transport services between IVDs (e.g., optical, Ethernet, or IP), the IVD packet-based encryption solution is extremely popular among DoD and various enterprise designers.

However, there are known industry challenges in trying to build an IP backbone with these devices as they fall short in delivering several key features users have come to expect in today's IP-based routers:

- Limited/no support for dynamic IP unicast routing protocols
- No support for dynamic IP multicast protocols
- No support for Virtual Route Forwarding (VRF), virtual LAN (VLAN), or Multiprotocol Label Switching (MPLS) functions (control and data plane)
- Limited capabilities for transporting IP type of service/differentiated services code point (ToS/DSCP) bits
- No 802.1Q/p support on the encryptor's Ethernet interfaces
- No means to rapidly detect host mobility between locations

Given these limitations, network operators instead most commonly deploy IP tunnel technology, namely generic routing encapsulation (GRE) (RFC 2784), between the secure routers inside the IVD boundary.

GRE tunnels, in combination with the IVD, allow the transport of IP service packets (e.g., IPv4/v6, multicast, or MPLS) through the IVDs. In its most generic form (and as stated in RFC 2784), GRE allows the transport of a "payload" packet (the packet needing to be encapsulated and delivered) within an outer header consisting of a GRE header plus an IP header (24 bytes total).

The result is the creation of an overlay IP topology between the GRE tunnel endpoints on the secure routers, which is transparent to the IVDs and the networks between IVDs. The IVDs route the traffic based on the destination IP address in the outer IP header of the GRE packets. By leveraging this "IP tunnel" overlay topology, the secure routers are able to support more sophisticated technologies than can be provided by the IVDs alone. These technologies include interior gateway protocols (IGPs) such as Open Shortest Path First (OSPF) or Enhanced Interior Gateway Routing Protocol (EIGRP); MPLS services such as IP Border Gateway Protocol (BGP) VPNs; Layer 2 VPN point-to-point or point-to-multipoint; IP multicast; and IPv6.

The use of GRE tunnels over IVDs has become common practice in certain customer deployments. It should be noted that enhancements to GRE performance (up to 40 Gbps of GRE packet encapsulation and forwarding on the Cisco CRS-1 Carrier Routing System, with future support of 140 Gbps of GRE on the Cisco CRS-3) and increased flexibility through the use of multipoint GRE tunnels look to overcome operational burdens historically associated with point-to-point GRE tunnels, especially in environments requiring larger number of sites (N – 1 GRE tunnels for full mesh, where N is the number of locations). However, using GRE tunnels in larger-scale environments has proven complex and troublesome, both operationally and in terms of the hardware required to forward GRE packets and manage the maximum transmission unit (MTU) implications GRE introduces.

The use of dynamic discovery of the routes to the secure networks the IVDs are protecting could increase demand on hardware resources and IVD functionality in order to process and hold a potentially larger number of IP prefixes being received from the protected network. This could prove challenging, particularly if the IVD hardware design was not originally intended to hold a large amount of IP prefixes.

Proposed Solution

This document details a solution to simplify the overall deployment and operations of IVDs, including using the Locator ID Separation Protocol (LISP) routing architecture framework to simplify the connectivity, forwarding, and operations between secure router endpoints when IVDs are required. The paper specifically addresses the challenges described prior, which could significantly impact how these types of networks are designed.

LISP is not a feature, but rather a new routing architecture that is gaining traction for the broad range of uses and applications with which it can integrate. LISP implements a new semantic for IP addressing that creates two name spaces: 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) comprising the global routing system.

Splitting EID and RLOC functions yields many benefits, including improved routing scalability, superior multi-homing efficiency, IPv6 transition, and virtual machine (VM) and IP mobility. Additionally, given the level of indirection incorporated into its forwarding scheme, LISP has been identified as a method for simplifying IP network deployments in customer networks when the use of external IVDs is required.

LISP can greatly simplify the overall IP routing paradigm in environments requiring IVDs, eliminating the need for full-mesh GRE tunnels running end-to-end routing protocols. LISP uses a "pull" model similar to DNS, which only requests endpoint host addresses when needed for communication. Because LISP inherently uses IP/UDP for forwarding (e.g., IP encapsulation), the data plane works seamlessly over IVDs without the need to manually configure GRE tunnels between each pair of secure routers. Further, LISP significantly limits the potential number of IP prefixes the IVD might be required to hold (e.g., RLOC addresses) to simplify the operational aspect of these networks. In turn, this will reduce the prefix memory and other related resources needed in the IVD design.

For a detailed description on how LISP operates, see http://tools.ietf.org/id/draft-farinacci-lisp-12.txt. That site also provides details about LISP and its control plane components that are outside the scope of this paper.

Solution Description of LISP in IVD Environments

This section addresses using LISP as the IP routing framework in an IVD environment, and assumes the reader has a basic understanding of the various LISP components, including data planes and control planes.

Figure 1 depicts the topology and components of a typical IP architecture in using IVDs and GRE tunnels. In this topology, point-to-point GRE tunnels are established between the secure router endpoints (at each campus/data center site), thus allowing full- or partial-mesh communications over each GRE tunnel. This communications overlay is transparent to the IVDs and the IP transport of the encrypted IVD traffic.

Figure 1 – Typical IVD Architecture Leveraging GRE Tunnel Overlay

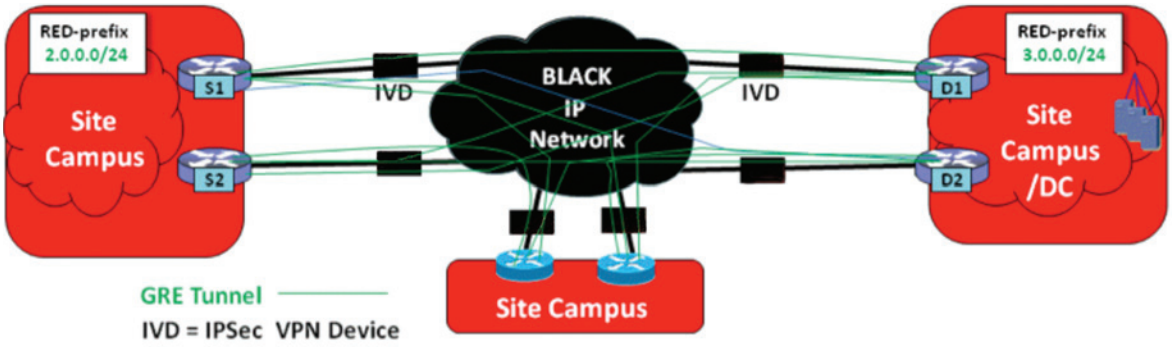

In the GRE deployment model, each IVD is responsible for holding the IP address prefix that the outer IP header of the GRE tunnel uses for communicating with each endpoint. It is important to note that this model hides the secure plaintext prefixes within each campus or data center site from the IVDs, thus limiting the number of prefix entries in each IVD to only those needed for GRE tunnel endpoint communications (vs. holding each secure prefix in the site campus or data center). In this model, GRE tunnels could be configured manually, or solutions such as Dynamic Multipoint VPN (DMVPN) could be used.

Figure 2 depicts the topology and components of a typical IP architecture using IVDs, this time with the use of a LISP framework and its associated components. The secure routers S1/S2 and D1/D2 will function as ingress tunnel routers (ITRs) and egress tunnel routers (ETRs) in the architecture. (Note: An "xTR" correlates to a router functioning as both an ITR and ETR.) The map resolver (MR) and map server(MS) ("MR/MS" in the figure) will be redundant and accessible only in the secure address space. If the need exists to communicate to non-LISP locations (this would be normal, even if only in a transition stage), one or more proxy ITRs/ETRs (PxTRs) will be provisioned; again, only accessible within the secure address space.

Figure 2 – Typical IVD Architecture Leveraging LISP

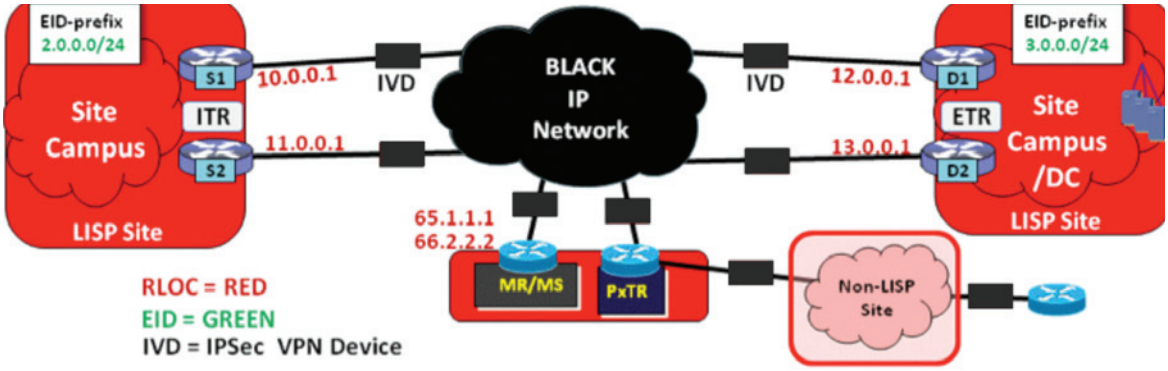

It should be noted that the IP network transport for the RLOC in this solution is immaterial to the function of the LISP architecture and can be assumed to use any of the network solutions typically found today in any IVD environment (e.g., IP, Ethernet, serial, optical).

LISP Operation in an IVD Environment

To use LISP in this secure environment, the secure xTRs will be directly connected to the IVDs. The RLOC addresses (shown as 10.0.0.1/32, 11.0.0.1/32, 12.0.0.1/32, and 13.0.0.1/32) would normally be manually entered into the IVD but could be advertised if an IGP was supported to allow this. The IVD will be responsible for distributing these RLOC addresses to all IVDs throughout the network as part of the normal IVD discovery process (details for IVD prefix discovery are outside the scope of this document). Lastly, the IP addresses for the MR/MS must also be advertised and reachable by all secure xTRs for the map registration and request functions to operate.

As in any LISP design, the EID address space will be hidden from the RLOC address space and, in this proposal, hidden from the IVDs as well. LISP operation in an IVD environment does not change from that found in the service provider/commercial space, meaning both data plane (ITR talking to an ETR) and control plane (ETR registering to a MS and an ITR requesting the RLOC-to-EID mapping from the MR) remain the same. However, in this secure environment, all communications will be over the IVD. Because LISP natively uses an IP encapsulation (IP/UDP) for forwarding, the operator is not required to manually configure any IP/GRE tunnels between any secure routers (e.g., ITRs/ETRs) over the IVDs.

One of the key concerns in an IVD architecture in which an IGP will be used for secure IP route discovery and exchange is the impact that a large amount of prefixes could have on the IVD devices, particularly when not using GRE tunnels. Consider that when using IVD discovery with an IGP, the IVD device itself would potentially need to learn all secure prefixes found within the agency's secure network and distribute those prefixes to every other IVD. The prefix count in each IVD could get very large, potentially impacting the overall performance and scale of the IVD devices and network performance.

Using the LISP framework for this type of network can greatly reduce the potential for route explosion in the IVD, as LISP inherently hides the end-user network prefixes (EIDs) from the IVDs through the RLOC/EID separation. No matter how large the secure routing tables become in the EID space, the IVDs will not be impacted and will only require knowledge of the RLOC address of each xTR when using the LISP framework.

The IVD routing table size-scaling factor is based only on the scale of the RLOC address space in the network, which will be minimal. Consider that the number of RLOC addresses will equal the number of secure router (xTR) interfaces, plus/minus the addresses of the MR/MS/PxTR, regardless of how much the prefix count increases behind each xTR (i.e., EID address space). This is an enormous benefit of LISP for scaling large IVD environments.

Aside from basic IP routing requirements in these networks, applications can be deployed that will increase this explosion of prefixes in the agency networks. One key application is the rapid addition of virtual machines (VMs), in which each VM host will have a /32 address (IPv4). The same can be said for mobile or tactical networks, where /32 addresses are much more frequently seen to identify each endpoint. (LISP has other methods for simplifying this VM mobility challenge, which is described briefly in the use case section.)

The LISP framework offers unlimited potential and should continue to be evaluated and considered in these complex IVD environments. Combining the "pull" method for host-to-host communication, dynamic prefix discovery in the IVD using IGPs, and native IP encapsulation, LISP has the potential to dramatically simplify overall network operation, setup, and scale for IVD network deployments and operation.

Key Advantages for LISP in an IVD Environment

Using the LISP routing architecture in combination with IVD deployments addresses several key challenges common in IVD networks today. Highlighted below are key advantages of how the LISP + IVD solution could benefit network operators and designers who are either already running IVDs in their networks, or planning to deploy them:

1. **Native IP forwarding:** Using the LISP framework in an IVD environment eliminates almost all the manual setup and change management required in IVD networks today, including GRE tunnel establishment, IGP/BGP peering over GRE, the impact of adding new locations, and address moves within a location.
2. **IP-encapsulated data plane:** The LISP data plane natively uses IP/UDP encapsulation (verses stateful tunnel technology) for forwarding, eliminating the need for the network operator to manually configure GRE tunnels between secure router endpoints.
3. **IVD discovery option:** In combining LISP functionality with optional IP prefix discovery options in the IVD, LISP xTRs can dynamically advertise their RLOC address to the IVD. This eliminates the need for establishing IGPs over static GRE tunnels between xTRs (the typical discovery process), and also simplifies–or eliminates–the configuration of static entries in each IVD device.
4. **Conversational learning (on-demand "pull" model):** Using LISP, each secure host (and/or xTR) only requests (i.e., "pulls") communications to the specific host with which it needs to communicate, on demand. This can be thought of as "conversational learning" in that it only requests specific information (/32 IP address) to talk to a specific host, verses inefficiently pushing routes, even where they are not needed. This creates an EID-to-RLOC cache entry in the sending router (ITR), which will maintain the cache for a period of time while the flow is active. In turn, this eliminates the need for all secure routers (and IVDs depending upon the prefix discovery method chosen) to hold all of the routes for the entire customer routing domain. It also eliminates the need to configure a full or partial mesh of GRE tunnels that require IGP neighbor establishment for route exchange.
5. **Reduction of routing table prefix entries in the IVD (assuming the use of IGPs for prefix discovery):** Because LISP uses a pull model and the IVD can use IGPs for secure IP prefix discovery, the IVD is no longer required to hold all the prefixes for each subnet in the secure routers on the secure side of the IVD. Instead, the IVD will only require knowledge of the RLOC prefixes, which will be /32 addresses and will equate to the number of secure router interfaces connecting to the secure side of the IVD. In contrast, a standard routing solution would use the push model, in which each secure router at each site would advertise its entire routing table to the IVD, which would then distribute it fully to the other IVDs and give each IVD and secure router an identical copy of the agency-wide routing table and topology.
6. **Controlling traffic in multi-homing topologies:** For locations running LISP that offer multiple entry points (two IVDs and/or two secure routers), operators have the option to control how the traffic load is sent to the receiving location (ETR) on a per-prefix basis. LISP ETRs have the ability to set a priority/weight on a per-prefix basis to dictate how traffic is sent to them by the originating ITRs.
7. **Seamless mobility of host/node/VM:** By leveraging the RLOC and EID separation capability inherent within LISP, node/host mobility between xTRs is available to the level that TCP connections can be maintained during the move. In this case, the host with an EID address would move (e.g., a virtual machine mobility requirement) and, while the same EID prefix follows the mobile host, the LISP infrastructure would dynamically discover the new RLOC now associated with the EID host, allowing communications to be seamlessly maintained.

8. **IPv6 transition:** Another key benefit of LISP is that it has the ability to use IPv6 addresses in the EID space while maintaining IPv4 addresses in the RLOC space, thus allowing the IVDs and the secure transport to remain at IPv4. In an IVD environment, IPv6 transition can begin in the secure router domain (EIDs) without the IVD and/or core transporting the RLOC addresses to also require a simultaneous transition to IPv6. This could be thought of as a "6 over 4" transition mechanism.

LISP and Competing Technologies

Although this paper highlights the advantages LISP offers in terms of simplifying IVD environments that are deployed in a variety of secure network topologies, it should be noted that LISP has relevant advantages for the secure encrypted traffic (e.g., RLOC addresses) network designs as well. For example, a standard IP transport for the encrypted traffic might have security requirements mandating the need for an additional layer of IPsec encryption for encrypted packets traversing between IVDs. LISP could be deployed on this network for reasons discussed in this paper, but would also provide the ability to leverage technologies such as Group Encrypted Transport (GET) VPN, which would hide (i.e., encrypt) the IVD address space. This would also include hiding the unsecure xTR EID addresses in the encrypted payload. In this scenario, the IVD source address would be an EID address relative to a xTR on the encrypted side of the IVD. Adding GET VPN encryption, the EID address would be encrypted, leaving only the RLOC space intact while transiting the IP transport, adding another level of security to the deployment.

In the areas where Layer 3 virtualization (VRFs) is required, there are proven MPLS VPN solutions using dynamic multipoint GRE technology that are tailored specifically to IVD environments. LISP is, however, also targeting virtualization deployment capabilities that will complement Cisco's suite of network virtualization options.

LISP is a key emerging technology for IP and is completely open standard. As such, LISP and MPLS VPN over IP feature enhancements and use cases continue to evolve. Cisco will continue to drive these innovations into features and capabilities for secure network communities.

Why Cisco

Cisco offers innovative products and solutions together with a wide range of services programs to accelerate customer success. These are delivered through a combination of people, processes, tools, and partners, resulting in high levels of customer satisfaction. LISP is Cisco® innovation that is being promoted as an open standard. Through its participation in standards bodies such as the IETF LISP Working Group, Cisco is committed to the development of the LISP architecture.

For More Information

Full details on these IPv6 transition strategies using LISP can be found in an IPv6 transition white paper located on the download page of the LISP website located at lisp.cisco.com. For more information about LISP, including information about the protocol itself, LISP deployment, LISP component descriptions, and LISP interworking, please visit www.cisco.com/go/lisp or lisp.cisco.com.

For general LISP solution questions, including deployment guidance, contact your local Cisco account representative or send an email to lisp-support@cisco.com.

References

Glen Nakamoto, Lisa Higgins, Justin Richer: MITRE Corporation, Scalable HAIPE Discovery Using a DNS-Like Referral Model

LISP Reference Source
http://tools.ietf.org/id/draft-farinacci-lisp-12.txt

LISP Overview
http://tools.ietf.org/id/draft-farinacci-lisp-12.txt

Informative References from "draft-farinacci-lisp-12"

[AFI] IANA, "Address Family Indicators (AFIs)," ADDRESS FAMILY NUMBERS http://www.iana.org/numbers.html, February 2007.

[ALT] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "LISP Alternative Topology (LISP-ALT)," draft-fuller-lisp-alt-03.txt (work in progress), February 2009.

[APT] Jen, D., Meisel, M., Massey, D., Wang, L., Zhang, B., and L. Zhang, "APT: A Practical Transit Mapping Service," draft-jen-apt-01.txt (work in progress), November 2007.

[CHIAPPA] Chiappa, J., "Endpoints and Endpoint Names: A Proposed Enhancement to the Internet Architecture," Internet-Draft http://www.chiappa.net/~jnc/tech/endpoints.txt, 1999.

[CONS] Farinacci, D., Fuller, V., and D. Meyer, "LISP-CONS: A Content Distribution Overlay Network Service for LISP," draft-meyer-lisp-cons-03.txt (work in progress), November 2007.

[DHTs] Ratnasamy, S., Shenker, S., and I. Stoica, "Routing Algorithms for DHTs: Some Open Questions," PDF file http://www.cs.rice.edu/Conferences/IPTPS02/174.pdf.

[GSE] "GSE - An Alternate Addressing Architecture for IPv6," draft-ietf-ipngwg-gseaddr-00.txt (work in progress), 1997.

[INTERWORK] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking LISP with IPv4 and IPv6," draft-lewis-lisp-interworking-01.txt (work in progress), January 2009.

[LISA96] Lear, E., Katinsky, J., Coffin, J., and D. Tharp, "Renumbering: Threat or Menace?," Usenix, September 1996.

[LISP-MS] Farinacci, D. and V. Fuller, "LISP Map Server," draft-fuller-lisp-ms-00.txt (work in progress), March 2009.

[LISP1] Farinacci, D., Oran, D., Fuller, V., and J. Schiller, "Locator/ID Separation Protocol (LISP1) [Routable ID Version]," Slide set http://www.dinof.net/~dino/ietf/lisp1.ppt, October 2006.

[LISP2] Farinacci, D., Oran, D., Fuller, V., and J. Schiller, "Locator/ID Separation Protocol (LISP2) [DNS-based Version]," Slide set http://www.dinof.net/~dino/ietf/lisp2.ppt, November 2006.

[LISPDHT] Mathy, L., Iannone, L., and O. Bonaventure, "LISP-DHT: Towards a DHT to map identifiers onto locators," draft-mathy-lisp-dht-00.txt (work in progress), February 2008.

[LOC-ID-ARCH] Meyer, D. and D. Lewis, "Architectural Implications of Locator/ID Separation," draft-meyer-loc-id-implications-01.txt (work in progress), January 2009.

[MLISP] Farinacci, D., Meyer, D., Zwiebel, J., and S. Venaas, "LISP for Multicast Environments," draft-farinacci-lisp-multicast-01.txt (work in progress), November 2008.

[NERD] Lear, E., "NERD: A Not-So-Novel EID to RLOC Database," draft-lear-lisp-nerd-04.txt (work in progress), April 2008.

[OPENLISP] Iannone, L. and O. Bonaventure, "Open LISP Implementation Report," draft-iannone-openlisp-implementation-01.txt (work in progress), July 2008.

[RADIR] Narten, T., "Routing and Addressing Problem Statement," draft-narten-radir-problem-statement-00.txt (work in progress), July 2007.

[RFC3344bis] Perkins, C., "IP Mobility Support for IPv4, revised," draft-ietf-mip4-rfc3344bis-05 (work in progress), July 2007.

[RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network Without a Flag Day," RFC 4192, September 2005.

[RPFV] Wijnands, I.J., Boers, A., and E. Rosen, "The RPF Vector TLV," draft-ietf-pim-rpf-vector-08.txt (work in progress).

[RPMD] Handley, M., Huici, F., and A. Greenhalgh, "RPMD: Protocol for Routing Protocol Metadata Dissemination," draft-handley-p2ppush-unpublished-2007726.txt (work in progress), July 2007.

[SHIM6] Nordmark, E. and M. Bagnulo, "Level 3 multi-homing shim protocol," draft-ietf-shim6-proto-06.txt (work in progress), October 2006.

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**.