



LISP Proxy Tunnel Router Redundancy Deployment

July 23, 2012

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at <http://www.cisco.com/go/trademarks>. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

LISP Proxy Tunnel Router Redundancy Deployment

© 2011 Cisco Systems, Inc. All rights reserved.



PxTR LISP Redundancy Deployment


This document describes how to configure a Proxy Tunnel Router when deployed in a redundant model. The content of this document relies heavily on work previously done in the Cisco Systems Development Unit (SDU) DCI lab, in conjunction with the CVD validation effort. Where appropriate and helpful, product configuration and documentation links are provided for further reference and detail.

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 to non-LISP interworking. These LISP infrastructure devices include the Proxy Ingress Tunnel Router (PITR) and Proxy Egress Tunnel Router (PETR).

A Proxy ITR provides connectivity between non-LISP sites and LISP sites. The Proxy ITR functionality is a special case of ITR functionality whereby the router attracts native packets from non-LISP sites, for example, the Internet, that are destined for LISP sites, and encapsulates and forwards them to the destination LISP site. A Proxy Tunnel Router is ideally placed in the path between LISP and non-LISP sites. In the validation topology, the redundant PxTRs are connected to the core network, as in [Figure 1](#).

The DCI validation topology did not use a Proxy ETR since the return traffic is forwarded by the ITR natively (without LISP encapsulation).

Figure 1 DCI Validation Topology



Configuring a LISP Proxy Tunnel Router

This document assumes a working knowledge of LISP. The data center and remote LISP sites with xTR functionality, the Map-Server(MS), and Map-Resolver(MR), as well as core connectivity, shown in [Figure 1](#), should have been previously configured. LISP documentation, including LISP Command Reference and Configuration Guides can be found at: http://lisp.cisco.com/lisp_down.html.

A Cisco ISR 3845 running IOS Engineering Software Release 15.1(4)XB6 was deployed as the Proxy Tunnel Router for this exercise.

The following procedure enables and configures LISP PxTR functionality.

-
- Step 1** Enter global configuration mode.
- configure terminal**
- Step 2** Enable and enter into the router LISP configuration mode.
- router lisp**
- Step 3** Select eid-table to be default, meaning the global routing table for association with the configured instance-ID and enter router-lisp-eid-table configuration mode.
- eid-table** *VRF name* **instance-ID** *instance-ID*
- Step 4** Define the extent of the LISP “EID space” it is proxying. For example, tell PITR specifically what EID space to send map-requests for and populate the PITR IPv4 LISP map cache when it receives traffic.
- map-cache** *IPv4 EID-prefix for RLOC static mapping* **map-request**
- Step 5** Configure rloc-probing as the locator reachability algorithm used by LISP to determine the reachability status of locators cached in its map-cache.
- loc-reach-algorithm rloc-probing**
- Step 6** Exit router-lisp-eid-table configuration mode and enter router LISP configuration mode.
- exit**
- Step 7** Configure the IPv4 address to be used by the Ingress Tunnel Router (ITR) as the source address for LISP IPv4 map-request messages sent to a configured Map-Resolver.
- ipv4 map-request-source** *source-address*
- Step 8** Configure a router to act as an IPv4 LISP Proxy Egress Tunnel Router (PETR). (**Note:** The DCI topology did not use PETR functionality.)
- ipv4 proxy-etr**
- Step 9** Configure the router to act as an IPv4 LISP Proxy Ingress Tunnel Router (PITR). By default, the router does not provide PITR functionality. The IPv4 locator address is used as a source address for data packets encapsulation, or for Map-Requests. It is often the IP address of a Loopback interface on a PxTR.
- ipv4 proxy-itr** *ipv4-local-locator*
- Step 10** Configure the IP address of the LISP Map Resolver to which this router, in the same way that an ITR uses a Map-Resolver, sends LISP map-requests to obtain EID-to-RLOC mapping and build a map-cache.
- ipv4 itr map-resolver** *IPv4 Address of map resolver*
- Step 11** Exit the router LISP configuration mode.
- exit**
-

LISP PxTR Redundancy Configuration Example

The remote non-LISP site is configured like any regular enterprise site and hosts the client IP network. The LISP Map-Server need not have any configuration or knowledge regarding non-LISP sites.

The redundant PxTRs are deployed with a Cisco ISR router configured as a LISP PITR as discussed above. The RLOC identifying each PxTR device is unique and represented by the IP address of a Loopback interface on each. When connected to the Internet, a PITR acts as a gateway between the Internet and the LISP-enabled network. To accomplish this, the PITR must advertise an aggregated prefix for the LISP EID namespace on behalf of LISP sites into the Internet, in order to attract to itself non-LISP-sourced traffic destined to those LISP EIDs. Network 10.25.0.0/16 is configured as the global LISP EID space on all the xTRs. In the following configuration example, any packet received by the PITR and destined to an IP address which is part of 10.25.0.0/16 will trigger a map-request. Traffic, for which EID-to-RLOC mapping is obtained and map-cache is built, will be LISP encapsulated and forwarded to the LISP site in the same way an ITR communicates with a LISP site.

Table 1 is a sample configuration based on this setup.

Table 1 Sample Device Configurations

PITR-1	PITR-2
<pre>interface GigabitEthernet0/0 description OSPF Neighbor link to Core mtu 9216 ip address 10.55.1.25 255.255.255.252 interface Loopback1001 description proxy-ITR ip address 77.77.77.77 255.255.255.255 router lisp eid-table default instance-id 0 map-cache 10.25.0.0/16 map-request exit loc-reach-algorithm rloc-probing ipv4 map-request-source 77.77.77.77 ipv4 proxy-itr 77.77.77.77 ipv4 itr map-resolver 5.6.6.100 exit router ospf 200 router-id 77.77.77.77 redistribute static metric 500 subnets route-map populate network 77.77.77.77 0.0.0.0 area 0 ip route 0.0.0.0 0.0.0.0 10.55.1.26 ip route 10.25.0.0 255.255.0.0 Null0 tag 123 route-map populate permit 123 match tag 123</pre>	<pre>interface GigabitEthernet0/0 description L3 OSPF Neighbor link to Core mtu 9216 ip address 10.55.1.29 255.255.255.252 interface Loopback1001 description proxy-itr ip address 88.88.88.88 255.255.255.255 router lisp eid-table default instance-id 0 map-cache 10.25.0.0/16 map-request exit loc-reach-algorithm rloc-probing ipv4 map-request-source 88.88.88.88 ipv4 proxy-itr 88.88.88.88 ipv4 itr map-resolver 5.6.6.100 exit router ospf 200 router-id 88.88.88.88 redistribute static metric 500 subnets route-map populate network 88.88.88.88 0.0.0.0 area 0 ip route 0.0.0.0 0.0.0.0 10.55.1.30 ip route 10.25.0.0 255.255.0.0 Null0 tag 456 route-map populate permit 456 match tag 456</pre>



Note

When a default route is configured, PITR may still have routing reachability to a failed RLOC. As a result, the locator state remains unchanged and traffic may get black-holed. To allow switching to using another locator from the cached locators in case of a failure scenario, enable the locator reachability algorithm rloc-probing. rloc-probing functionality is disabled by default by LISP.

In a redundant model, the two PITRs operate in a completely stateless fashion with each other. Both PITRs redistribute the same coarse route. As a result, the remote router will see two updates about network 10.25.0.0, however, only puts the route to network 10.25.0.0 with the lower metric. Packets entering the network will seek that PxTR or the one that is up, in case of a PxTR failure.

```
dca-branch2#show ip route 10.25.0.0
Routing entry for 10.25.0.0/16
  Known via "ospf 200", distance 110, metric 500
  Tag 123, type extern 2, forward metric 2
  Last update from 10.46.1.6 on GigabitEthernet0/1, 4d21h ago
  Routing Descriptor Blocks:
  * 10.46.1.6, from 77.77.77.77, 4d21h ago, via GigabitEthernet0/1
    Route metric is 500, traffic share count is 1
    Route tag 123
```

Or, when pxtr1 is down:

```
dca-branch2#show ip route 10.25.0.0
Routing entry for 10.25.0.0/16
  Known via "ospf 200", distance 110, metric 500
  Tag 456, type extern 2, forward metric 3
  Last update from 10.46.1.6 on GigabitEthernet0/1, 00:00:03 ago
  Routing Descriptor Blocks:
  * 10.46.1.6, from 88.88.88.88, 00:00:03 ago, via GigabitEthernet0/1
    Route metric is 500, traffic share count is 1
    Route tag 456
```

In addition to the configuration on each PITR as shown, it's necessary to filter the incoming routing updates on the other LISP-enabled routers. This is to ensure that remote LISP site traffic is not attracted to the PITR for LISP encapsulation but locally handled. Following is an example configuration on a LISP enabled branch router where external prefixes with the tag values of 123 and 456 are prevented from being installed in the routing table.

```
router ospf 200
  router-id 8.8.8.1
  network 8.8.8.1 0.0.0.0 area 111
  distribute-list route-map test5 in
!
route-map test5 deny 10
  match tag 123
!
route-map test5 deny 20
  match tag 456
!
route-map test5 permit 30

No native route should be present for the remote EID.
dcb-branch1#show ip route 10.25.1.50
% Subnet not in table
```

Verifying the PxTR

The following procedure verifies the PxTR.

Step 1 Verify that the PITR lisp map-cache has no entries initially.

```
dcb-PxTR1#show ip lisp map-cache
LISP IPv4 Mapping Cache for EID-table default (IID 0), 1 entries

10.25.0.0/16, uptime: 00:01:43, expires: never, via static send map-request
```

```
Negative cache entry, action: send-map-request
```

Step 2 Verify that the core routers see the IPv4 LISP EID aggregate prefix, as advertised by the PITR.

```
dcb-core2#show ip route 10.25.0.0
Routing entry for 10.25.0.0/16
  Known via "ospf 200", distance 110, metric 500
  Tag 123, type extern 2, forward metric 1
  Last update from 10.55.1.25 on GigabitEthernet3/14, 00:03:37 ago
  Routing Descriptor Blocks:
  * 10.55.1.25, from 77.77.77.77, 00:03:37 ago, via GigabitEthernet3/14
    Route metric is 500, traffic share count is 1
    Route tag 123
```

A host on the remote non-LISP site sends a packet to an EID destination. The packet arrives at a PITR where a LISP map-request gets triggered if there is no valid entry in the local map-cache. The PITR must then create a LISP map-cache entry for this EID prefix.



Note The **debug lisp control-plane all** command can be used on the PITR to verify the process of how the map-cache entry gets created.

Step 3 Verify that Map-cache now has the entry installed.

```
dcb-PxTR1#show ip lisp map-cache
LISP IPv4 Mapping Cache for EID-table default (IID 0), 2 entries

10.25.0.0/16, uptime: 00:04:51, expires: never, via static send map-request
  Negative cache entry, action: send-map-request
10.25.1.50/32, uptime: 00:00:12, expires: 23:59:40, via map-reply, complete
Locator  Uptime  State  Pri/Wgt
1.1.1.35  00:00:12  up      1/10
1.1.1.50  00:00:12  up      1/10
```

Step 4 To verify PITR redundancy, shut down PITR-1 interface connected to the core router. Verify that the non-LISP branch router and core routers see the IPv4 LISP EID aggregate prefix advertised by PITR-2.

```
dcb-branch2#show ip route 10.25.0.0
Routing entry for 10.25.0.0/16
  Known via "ospf 200", distance 110, metric 500
  Tag 456, type extern 2, forward metric 3
  Last update from 10.46.1.6 on GigabitEthernet0/1, 00:00:29 ago
  Routing Descriptor Blocks:
  * 10.46.1.6, from 88.88.88.88, 00:00:29 ago, via GigabitEthernet0/1
    Route metric is 500, traffic share count is 1
    Route tag 456

dcb-core2#show ip route 10.25.0.0
Routing entry for 10.25.0.0/16
  Known via "ospf 200", distance 110, metric 500
  Tag 456, type extern 2, forward metric 2
  Last update from 10.55.1.2 on TenGigabitEthernet1/2, 00:01:11 ago
  Routing Descriptor Blocks:
  * 10.55.1.2, from 88.88.88.88, 00:01:11 ago, via TenGigabitEthernet1/2
    Route metric is 500, traffic share count is 1
    Route tag 456
```

Step 5 PITR-2 must then create a LISP map-cache entry for that EID prefix.

```
dcb-PxTR2#show ip lisp map-cache
LISP IPv4 Mapping Cache for EID-table default (IID 0), 2 entries

10.25.0.0/16, uptime: 01:26:16, expires: never, via static send map-request
```

```

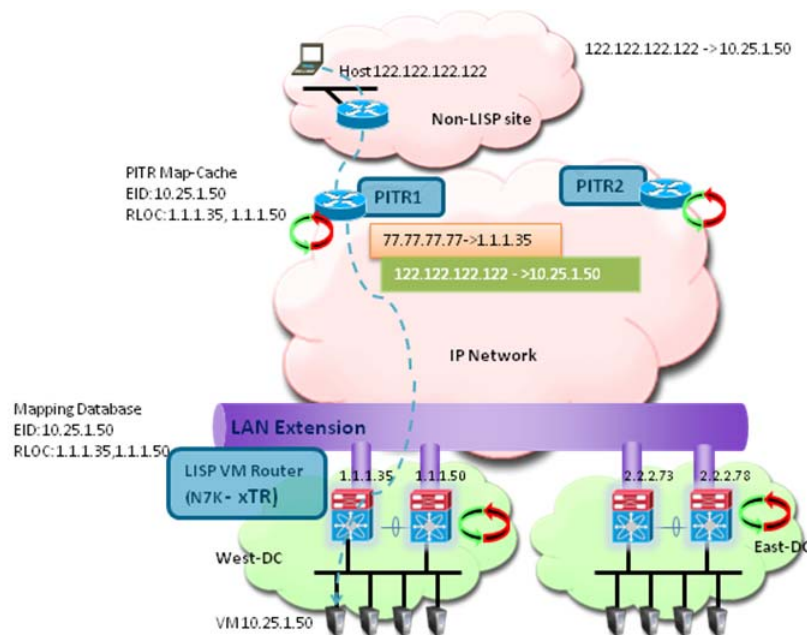
Negative cache entry, action: send-map-request
10.25.1.50/32, uptime: 00:00:13, expires: 23:59:39, via map-reply, complete
Locator  Uptime  State  Pri/Wgt
1.1.1.35 00:00:13 up      1/10
1.1.1.50 00:00:13 up      1/10

```

Step 6 After creating a LISP map-cache entry, when the PITR receives EID-destined packets, it checks the LISP map-cache first. If there is a valid entry, it will be used for encapsulating the packet and then forwarded to the ETR.

Figure 2 shows how the client to server traffic path would look.

Figure 2 *Non-LISP to LISP: Client to Server Traffic Flow using LISP Routing*



The following example shows that 10 packets were sent from a host on the non-LISP site to the EID destination 10.25.1.50.



Note

There is no **debug ip lisp** command in IOS however, you can use CEF debugging for the data plane. Following are some helpful show commands to check LISP encapsulation counters (underlined in the before and after output below) on the PITR.

Before

```

dcb-PxTR2#show ip lisp forwarding eid remote 10.25.1.50
Prefix          Fwd action  Locator status bits
10.25.1.50/32   encap       0x00000003
  packets/bytes 389/32676
  path list C184F630, flags 0x49, 4 locks, per-destination
  ifnums:
    LISP0(15): 1.1.1.35, 1.1.1.50
  2 paths
    path 69F7464C, path list C184F630, share 10/10, type attached nexthop, for IPv4
    nexthop 1.1.1.35 LISP0, adjacency IP midchain out of LISP0, addr 1.1.1.35 7003BB60

```



```

path 69F7448C, path list C184F630, share 10/10, type attached nexthop, for IPv4
nexthop 1.1.1.50 LISP0, adjacency IP midchain out of LISP0, addr 1.1.1.50 7003BA00
1 output chain
chain[0]:
loadinfo 6984897C, per-session, 2 choices, flags 0083, 5 locks
flags: Per-session, for-rx-IPv4, 2buckets
2 hash buckets
< 0 > IP midchain out of LISP0, addr 1.1.1.35 7003BB60 IP adj out of
GigabitEthernet0/0, addr 10.55.1.30 7003BCC0
< 1 > IP midchain out of LISP0, addr 1.1.1.50 7003BA00 IP adj out of
GigabitEthernet0/0, addr 10.55.1.30 7003BCC0
Subblocks:
None

```

dcb-PxTR2#show adjacency lisp 0 detail

```

Protocol Interface      Address
IP      LISP0          1.1.1.35(5)
                                0 packets, 0 bytes
                                epoch 0
                                sourced in sev-epoch 9
                                Encap length 36
                                450000000000400000116D760A37011D
                                01010123000010F5000000008002D30E
                                00000000
                                LISP
                                Next chain element:
                                IP adj out of GigabitEthernet0/0, addr 10.55.1.30
IP      LISP0          1.1.1.50(5)
                                389 packets, 46680 bytes
                                epoch 0
                                sourced in sev-epoch 9
                                Encap length 36
                                450000000000400000116D670A37011D
                                01010132000010F5000000008003A653
                                00000000
                                LISP
                                Next chain element:
                                IP adj out of GigabitEthernet0/0, addr 10.55.1.30

```

After

dcb-PxTR2#show ip lisp forwarding eid remote 10.25.1.50

```

Prefix      Fwd action  Locator status bits
10.25.1.50/32  encap      0x00000003
packets/bytes  399/33516
path list C184F630, flags 0x49, 4 locks, per-destination
ifnums:
LISP0(15): 1.1.1.35, 1.1.1.50
2 paths
path 69F7464C, path list C184F630, share 10/10, type attached nexthop, for IPv4
nexthop 1.1.1.35 LISP0, adjacency IP midchain out of LISP0, addr 1.1.1.35 7003BB60
path 69F7448C, path list C184F630, share 10/10, type attached nexthop, for IPv4
nexthop 1.1.1.50 LISP0, adjacency IP midchain out of LISP0, addr 1.1.1.50 7003BA00
1 output chain
chain[0]:
loadinfo 6984897C, per-session, 2 choices, flags 0083, 5 locks
flags: Per-session, for-rx-IPv4, 2buckets
2 hash buckets
< 0 > IP midchain out of LISP0, addr 1.1.1.35 7003BB60 IP adj out of
GigabitEthernet0/0, addr 10.55.1.30 7003BCC0
< 1 > IP midchain out of LISP0, addr 1.1.1.50 7003BA00 IP adj out of
GigabitEthernet0/0, addr 10.55.1.30 7003BCC0
Subblocks:
None

```

```
dcb-PxTR2#show adjacency lisp 0 detail
Protocol Interface      Address
IP          LISP0      1.1.1.35(5)
                                0 packets, 0 bytes
                                epoch 0
                                sourced in sev-epoch 9
                                Encap length 36
                                4500000000000400000116D760A37011D
                                01010123000010F5000000008002D30E
                                00000000
                                LISP
                                Next chain element:
                                IP adj out of GigabitEthernet0/0, addr 10.55.1.30
IP          LISP0      1.1.1.50(5)
                                399 packets, 47880 bytes
                                epoch 0
                                sourced in sev-epoch 9
                                Encap length 36
                                4500000000000400000116D670A37011D
                                01010132000010F5000000008003A653
                                00000000
                                LISP
                                Next chain element:
                                IP adj out of GigabitEthernet0/0, addr 10.55.1.30
```

Durga Nagulpally**DCI Systems Engineer, Systems Development Unit, Cisco Systems**

Durga Nagulpally is a Systems validation Engineer in SDU. Her background in the networking industry spans over 15 years in protocol design, development test, Beta test, Pre-Qual and Sustaining engineering roles at Cisco and other Silicon Valley companies. Durga has experience in a wide range of technologies including Routing, Switching, IP Multicast, WAAS-WCCP WAN optimization, ATM, FR, PPP/oE, L2TP, AAA/Radius, Subscriber Aggregation and DCI related A-VPLS, OTV, ACE, GSS and LISP. In her current role, Durga focuses on DCI Systems releases for the enterprise customer. Her most recent and successful collaborative Cisco solution buildout and validation efforts include VMware vSphere 5.0 - Metro vMotion solution Certification, and Cisco LISP VM-Mobility & Path Optimization Solution with NetApp FlexCache. Durga has a Master's degree in electrical and computer engineering from the University of Arizona.



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, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn is a service mark; and Access Registrar, Aironet, 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 Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. 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. (0805R)