

Cisco 7600 Series Ethernet over MPLS

1. Objective & Target Audience

The objective of this document is to provide a detailed description of the features, benefits and market opportunities associated with the Ethernet-over-MPLS (EoMPLS) feature set supported on the Cisco 7600 Series router, which is based on the IETF Martini Draft for Ethernet transport service over MPLS[2]. This technology brief is targeted towards Service Providers that are planning to offer Metropolitan Area Network services based on MPLS, or to the customers of these Service Providers for such services.

Because EoMPLS technology is closely associated with MPLS technology, this document assumes that the reader has at least some basic knowledge of MPLS technology. As such, this document discusses MPLS technology only to the extent of describing how EoMPLS technology integrates into an MPLS backbone network, and does not provide a detailed explanation of the configuration and operation of the MPLS backbone network.

Finally, this document contains various acronyms, and also a number of references to technical documents. Each of the acronyms is defined at the end of this document in the section titled, "DEFINITIONS". The technical references (notated throughout this document as [#], where # is the reference number) are found in the section titled, "REFERENCE DOCUMENTS".

2. Market Opportunity

One of the most important emerging market opportunities for Service Providers focused on corporate customers is the requirement for optical speed data networks to connect multiple corporate sites within a specific metropolitan geographical area. This market requirement is frequently referred to as a Metropolitan Area Network (or "Metro" for short). The corporate customers that require these Metro services range in size from the Fortune 500 Enterprise customers all the way down to medium and small commercial businesses located in commercial business parks or building complexes. The common characteristic of virtually all of these customers is their need to establish private networks between two or more locations in a specific Metro service region.

Those corporate customers that are looking for Metro services have two basic requirements:

- Transparent LAN services (TLS) between various corporate facilities within the Metro service region.
- IP services for general purpose, high-speed access to the Internet or to other corporate facilities located beyond the Metro service region.



As Service Providers attempt to meet these emerging requirements, they must evaluate new networking technologies and architectures at both the data link and physical layers, as well as at the network layer. Obviously, these new technologies must accommodate the service requirements of the Metro customer.

One of the most promising areas of networking technology to address these requirements is MPLS, in general, and EoMPLS in particular: MPLS provides an excellent foundation for providing private network services across a wide range of data link and physical network technologies throughout the service providers network domain, whether that domain is in a single Metro service area, or is an international network spanning many Metro service areas.

EoMPLS technology leverages an existing MPLS backbone network to deliver Transparent LAN Services (i.e. TLS) based on Ethernet connectivity to the customer site. The concept of Transparent LAN Services is straightforward: it is the ability to connect two Ethernet networks, which are geographically separate, and have the two networks appear as a single logical Ethernet or VLAN domain. The introduction of such a VLAN transport capability will allow service providers to deliver a service that allows VLAN networks in different locations within a metro service area to be cost-effectively connected at transmission speeds equivalent to Fast Ethernet or Gig Ethernet.


When EoMPLS is deployed in conjunction with MPLS VPN, the Service Provider will be able to provide tremendous flexibility in the variety of both L2 and L3 network services that can be provisioned for their Metro customers, and will do so over a single, simplified, integrated MPLS backbone network.

3. Functional Characteristics

3.1 Overview

For the purpose of this discussion, the network reference model shown in Figure 1, below, will be used through the remainder of this document.

Figure 1 EoMPLS Logical Topology



Although EoMPLS is capable of delivering TLS functionality, it should not be confused with traditional LAN bridging. Unlike traditional LAN bridging, EoMPLS does not perform any layer-2 look up to determine if the destination MAC address resides on the local or remote segment, and does not perform any layer-2 address learning, as traditional LAN bridging does. Instead, EoMPLS is more analogous to the transport of Layer 2 Frame Relay packets through an ATM backbone (which, in the future, could also be migrated to an MPLS backbone). The discussion of support for AAL5 and Frame relay has been presented in other documents ([2] & [3]), so the remainder of this document is limited to the support for transporting Ethernet VLAN packets over an MPLS network as per [1].



3.2 Characteristics of EoMPLS Services

Although EoMPLS technology offers several important benefits to service providers that intend to offer Transparent LAN Services to their Metro customers, it also has several characteristics, which the service provider and user need to understand to make effective use of this technology:

1. Establishing an EoMPLS circuit requires that the Service Provider customer be assigned a specific physical port on an LER device, such as a Cisco 7600. The identification of that physical port is a critical element in the binding of the MPLS Label assigned to the customers EoMPLS VC.
2. A customer may have more than one EoMPLS VC per physical port as long as the Ethernet traffic transmitted from the customer site to the PE device can be distinguished as having specific 802.1Q headers for each EoMPLS VC by the PE device. This requires coordination between the Service Provider and customer in the provisioning of the EoMPLS service.
3. EoMPLS VC's are point-to-point transmissions only, as explicitly specified in the IETF Martinin draft specifications.
4. Traffic sent between the imposition/disposition routers (between LERs) over an EoMPLS VC will take the same path across the IP/MPLS backbone. The LSP may change due to routing changes inside the provider network.
5. Adding/removing a point-to-point layer 2 VC requires configuration of the two VC endpoints (at the two LERs). Provisioning a VC will involve defining an endpoint of a Layer 2 VC at each of the VLAN interfaces at the PE router on the interface that connects to the CE.
6. The two LERs at the ingress/egress points of the IP/MPLS backbone (the PE routers) are the only routers with knowledge of the layer 2 transport VCs. All other LSRs will have no table entries for the layer 2 transport VCs. This means that only the PEs require software with EoMPLS functionality.

3.3 Technical Requirements for EoMPLS LER's

In order for the Cisco 7600 to perform as an LER to support transport of layer2 VLAN packets over MPLS, it must be able to forward 802.1QVLAN layer 2 VCs across an IP/MPLS backbone. The features required to provide this capability are:

- Targeted LDP session between the peers
- Virtual Circuit between peers
- Two level labeling
- Label imposition/disposition
- CoS mapping

3.3.1 LDP session

The ingress PE needs a tunnel label from its upstream IGP neighbor to route the packets it receives from the ingress interface of the ingress PE to the egress PE. The ingress PE and its IGP neighbor are local label distribution peers and maintain a LDP session over which tunnel labels are distributed. They must all have the FEC that contains the host route for the egress PE, and the tunnel label is bound to that FEC.

In addition, one targeted LDP session is required between the ingress PE and egress PE, which are remote label distribution peers, to exchange VC labels. All VC label bindings exchanged over this LDP session use the Virtual Circuit FEC element type 128 as defined in [2] via the LDP “downstream unsolicited mode” described in [4]. Since only a single LDP session is required, one is only created if not already present. A session may already be active because another application has already established a targeted LDP session between the two PE routers. LDP, which is only a transport protocol, shall determine that a LDP message is for the AToM application by checking for the existence of the Virtual Circuit FEC element type 128.

Support for the Virtual Circuit FEC element type 128 is required in the following LDP messages:

- Label Mapping Request



- Label Mapping
- Label Mapping Withdraw

If using conservative label retention, the ingress PE also needs to send Label-Request messages for all locally configured VCs. If using liberal label retention, the ingress PE does not need to send Label-Request messages for configured VCs, but must prepare to respond to Label-Request sent by other PEs which use conservative label retention. In AToM, we are going to migrate to using liberal label retention mode, but initially the implementation will employ conservative label retention mode.

3.3.2 Two level labeling

The layer 2 transport service over MPLS is implemented through the use of two level label switching between the edge routers. The label used to route the packet over the MPLS backbone to the destination PE is called the “tunnel label” (see [1]). The label used to determine the egress interface is referred to as the VC label. The egress PE allocates a VC label and binds the Layer 2 egress interface to the VC in question, then it signals this label to the ingress PE via the targeted LDP session.

3.3.3 Label Imposition/Disposition

When the layer 2 PDU arrives at the ingress interface of the ingress PE, the router must perform label imposition, and switch the packet to the appropriate outgoing MPLS interface which will route the packet to the egress LER for the VC in question. When the egress PE receives the packet it receives the packet with only the VC label because its neighbor (known as the penultimate router) pops the tunnel label prior to forwarding the packet. The egress PE uses the VC label to perform disposition and switch the packet to the appropriate egress interface.

3.3.4 Class of Service (Quality of Service) Mapping

MPLS provides quality of service using the 3 experimental bits in a label to determine the queue of packets. In order to support QoS from PE to PE, the experimental bits in both the VC and Tunnel labels will have to be set. The experimental bits need to be set in the VC label because the tunnel label is popped at the penultimate router. In the case of EoMPLS, two methods of setting experimental bits are provided.

3.3.4.1 Static setting of experimental bits

The service provider will be able to configure the PE to set the EXP bits in the labels to a given value based in ingress interface.

3.3.4.2 Using VLAN user priority bits to determine Experimental bit settings

The 3 user priority bits will be used to index into a table of 8 values. The value for a given index will be used to set the experimental bits. This method may cause out of order packets when packets have different user priorities. (NOTE: issue when sequence numbers supported). See [2] for more details on this requirement.

3.4 MPLS VC Circuit Set up

A Virtual Circuit is an LSP tunnel between the ingress and egress PE, which consists of two LSPs since a uni-directional LSP is required to transport Layer 2 PDUs in each direction. A two-level label stack as defined in [3], where the Level 1 label is the VC label and the Level 2 label is the VLAN tunnel label, is used to switch packets back and forth between the ingress and egress PE.

The VC label is provided to the ingress PE by the egress PE of a particular LSP to direct traffic to a particular egress interface on the egress PE. A VC label is assigned by the egress PE during the VC setup and represents the binding between the egress interface and a given VC ID. A VC is identified by a unique and configurable VC ID that is recognized by both the ingress and egress PE. During a VC setup, the ingress and egress PE exchange VC label bindings for the specified VC ID. The VC setup procedures are transport independent. The detailed VC setup procedure occurs as follows:

- An MPLS L2 transport route is entered on the ingress interface on PE1.



- PE1 starts a remote LDP session with PE2 if there is none already existing. Both PEs receive LDP KeepAlive messages from each other, reach OPERATIONAL state, and are ready to exchange label bindings.
- The physical layer of the ingress interface on PE1 comes up. PE1 realizes there is a VC configured for the ingress interface over which Layer 2 PDUs received from CE1 are forwarded, so it allocates a local VC label and binds it to VC ID configured under the ingress interface.
- PE1 encodes this binding with the VC label TLV and VC FEC TLV and sends it to PE2 in a Label-Mapping message.
- PE1 receives a Label-Mapping message from PE2 with a VC FEC TLV and VC label TLV. In the VC FEC TLV, the VC ID has a match with a locally configured VC ID. The VC label encoded in the VC label TLV is the outgoing VC label that PE1 is going to use when forwarding Layer 2 PDUs to PE2 for that particular VC.
- PE1 might receive a Label-Request message from some other PE with a specific VC FEC TLV at any time during the OPERATIONAL state. PE1 examines the VC ID encoded in the FEC element, and responds to the peer a Label-Mapping with the local VC label corresponding to the VC ID.
- PE2 performs the same steps 1-6 as PE1. Once both exchange the VC labels for a particular VC ID, we say the VC with that VC ID is fully established.
- When one LSP of a VC is taken down for some reason, for example, the CE-PE link goes down, or the VC configuration is removed from one PE router, the PE router must send a Label-Withdraw message to its remote LDP peer to withdraw the VC label it previously advertised.

4. End User Interface

4.1 EoMPLS CLI commands

The following commands have been added to IOS to support the transport of VLAN packets over MPLS.

4.1.1 Command for routing a layer 2 transport VC

```
[no] mpls l2transport route <destination> <vcid>
```

<destination> is used to select the LSP as well as specify the destination of the targeted LDP session

<vc-id> is the common identifier that is used by each end of the VC to identify the VC.

This command is used at the ingress/egress VLAN interface to route a layer 2 transport over MPLS point-to-point VC between two PEs. The destination cannot be an IP address on the router that the command is defined from (i.e.: only PE to another PE VCs are supported). The <vc-id> is used by the two PEs as a common reference point. This command must be invoked at both ends to establish a bi-directional VC which is comprised of two uni-directional LSPs. A VC will not be established if not properly defined from both ends.

Example: layer two transport VC between PE1's VLAN 3 interface and PE2's VLAN 4 interface.

We assume that PE1 has an IP address 10.0.0.1/32 that is discovered by PE2 via routing. Similarly it is assumed PE2 has IP address 20.0.0.1/32 that is discovered by PE1 via routing.

A <vc-id> equal to 333 is arbitrarily chosen for this example. Service providers may choose to manage assignment of <vc-id>s to avoid the possibility for misconfiguring and for simplifying troubleshooting. <vc-id>s need to be unique per pair of routers. In order to simplify the implementation, <vc-id> will be required to be unique on a per router basis. This



simplifying requirement will not cause problems in terms of interoperability since there is 32 bits available for VC IDs. Specifying this restriction eliminates many possibilities for mis-configuration which might otherwise be difficult to troubleshoot.

It is assumed that the PE is properly configured with an MPLS connection. Details on how to configure an MPLS network is beyond the scope of this document.

At PE 1 the following config is required at the VLAN interface being routed using layer two transport:

```
(config)#interface vlan 3
(config-if)# mpls l2transport route 20.0.0.1 333
```

At PE 2 the following config is required at the VLAN interface being routed using layer two transport:

```
(config)#interface vlan 4
(config-if)# mpls l2transport route 10.0.0.1 333
```

4.1.2 Line card show commands

The only new show commands required in the linecard used to show the layer 2 transport imposition and disposition information that has been programmed into the card. This command is provided primarily to troubleshoot any synchronization issues between the control plane and data plane.

```
show mpls l2transport [imposition | disposition] [detailed]
[[<vc-id> | <vc-id-min>] <vc-id-max>]
```

4.1.3 Control Plane show commands

The following command is used to vc info:

```
show mpls l2transport vc [detailed] [[<vc-id> | <vc-id-min>] <vc-id-max>]
[<vc-id>] – specify a single VC to be displayed.
[[<vc-id-min>] <vc-id-max>] – allows a user to specify a range of vc-ids to be displayed.
```

The following command will be provided to summarize the number of vc's routed to a given destination address and also to summarize all the active VCs on the PE's MPLS interfaces

```
show mpls l2transport summary
```

4.1.4 Control Plane and Linecard (data plane) debug commands

The following debug command will be supported on both the Route Processor(RP) and Linecard:

```
debug mpls l2transport {vlan}{control | switching | distributed }
control – control plane label setup/teardown
switching – no distributed data plane debug output
distributed – distributed data plane debug (ie Linecard debugging).
```

4.1.5 Commands for CoS

Either the Modular QoS CLI will be used or a new command will be introduced for setting the experimental bits. The command being proposed is as follows:

```
[no] mpls l2transport exp-map <val0><val1><val2>.....<val7>
```



If the experimental bit mapping is not explicitly specified, the default will be a direct mapping of the user priority bits to the experimental bits.

Example:

```
(config)#interface vlan 3
(config-if)# mpls l2transport route 20.0.0.1 333
(config-if)# mpls l2transport exp-map 0 1 2 3 4 4 4 4
```

The above mapping causes packets with user priority 0, 1, 2, 3 to be directly mapped to the experimental bits. User priority settings 4, 5, 6, 7 all cause the experimental bits to be set to 4.

5. EoMPLS Restrictions

5.1 Supported VC types

Only the transport of VLAN packets will be supported in the initial implementation for [1].

5.2 Dynamic IP labeling

In order to support this feature, Dynamic IP labeling (“mpls ip”) must be enabled on all paths between the two imposition/disposition LERs. Failure to do so will result in the packet being discarded before it reaches the disposition PE.

5.3 Summarized Routes

Routes from a PE discovered by its peers must be unsummarized (i.e.: address/32). This is required in order to insure that there is an LSP from PE to PE.

5.4 Layer 2 packet size

Fragmentation is not supported for layer 2 packets transmitted across the MPLS backbone. Therefore, in order to deploy the layer transport service, the network operator must make sure that the MTU of all intermediate links between the endpoints is sufficient to carry the largest layer 2 transported packet that will be received. The MTU setting of the ingress PE will need to match with the setting at the egress PE. A VC will not properly establish if MTU sizes mismatch. Checking of the MTU on MPLS interfaces will not be done by the EoMPLS application.

5.5 Explicit routes between endpoints


In the case where there are multiple LSPs between the two PEs, it may be desirable to be able to select which LSP is used to route a VC. For example, multiple traffic engineering paths may be created in order to carry various grades of traffic. The first implementation will allow the service provider to select more than one LSP to a destination router, but the implementation may not fully interoperate with other vendors implementation. The reason for this initially is the lack of clear guidelines and procedures in [1] which explains how to set up the LDP session or how many should be setup. The Cisco implementation will employ multiple LDP targeted discoveries for to each destination address specified. A Cisco router will respond by re-using an existing targeted session if one exists. Other vendors routers may not be capable of this.

6. Configuration Example

As an example of how IOS can be configured to support an EoMPLS virtual circuit, the IOS configuration example shown in Figure 2 is useful. Based on the following diagram, the IOS commands needed to configure the Cisco 7600 #1 to support two EoMPLS connections from Gig Ethernet port with IP address 10.1.1.1 through the MPLS core to Cisco 7600 #2 GE port with IP address 10.1.1.2 are shown below:



Figure 2 OS Configuration Example for EoMPLS



```
mpls label protocol ldp
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.255
!
interface GigabitEthernet4/1
 ip address 10.4.1.2 255.255.0.0
 no negotiation auto
 tag-switching ip
!
interface Vlan25
 no ip address
 mpls l2transport route 10.1.1.2 50
!
interface Vlan60
 no ip address
 mpls l2transport route 10.1.1.2 1021
!
```

Once Cisco 7600 #1 is configured with these commands, then symmetrical configuration of Cisco 7600 #2 must be performed to complete the EoMPLS circuit. Once both Cisco 7600 #1 and #2 are configured, and assuming that MPLS Label Distribution Protocol (LDP) is operating throughout the MPLS core, then LDP sessions connecting Cisco 7600 #1 and #2 to the core MPLS network will automatically signal the MPLS VC forwarding information throughout the MPLS network.

7. Reference Documents

- {1} draft-martini-l2circuit-trans-mpls-05.txt
- [2] draft-martini-l2circuit-encap-mpls-01.txt
- [3] RFC 3031, Multiprotocol Label Switching Architecture
- [4] RFC3036, LDP Specification



8. Definitions

This section defines words, acronyms, and actions which may not be readily understood.

<i>AToM</i>	Any Transport over MPLS.
<i>VLAN over MPLS</i>	Encapsulation of VLAN layer 2 PDUs across a MPLS backbone.
<i>Customer Edge (CE)</i>	The router in the customer's network that is used to connect to the PE Router.
<i>Dynamic MPLS tunnel</i>	An IOS tunnel interface that defines the endpoints for encapsulation of traffic across a MPLS backbone.
<i>L2transport over MPLS</i>	General term used to refer to the encapsulation Layer 2 PDUs and transport of them over an MPLS backbone. VLAN over MPLS is an example of L2transport over MPLS. Also sometimes referred to by the acronym AToM.
<i>Label imposition</i>	The act of putting label(s) on a packet. Label imposition is done by the "label imposition router" or "Label Edge Router (LER)". In the case of L2transport over MPLS, this is the router that receives a layer2 packet and encapsulates it for the MPLS backbone.
<i>Label disposition</i>	The act of removing label(s) from a packet. This is done by the "label disposition router" or "Label Edge Router (LER)". In the case of L2transport over MPLS, this is the router that receives a packet, removes the last label and transmits the layer2 PDU out the appropriate interface.
<i>Label Edge Router (LER)</i>	The edge router that is used to perform label imposition and disposition. L2 transport over MPLS application runs on the Label Edge Router.
<i>Label Switched Path (LSP)</i>	A route identified by a label used to switch packets to a destination. The label used to forward a packet using a LSP is referred to as the tunnel label.
<i>LDP</i>	Label Distribution Protocol.
<i>Provider Edge (PE)</i>	The LER in the provider network that is used to connect to the Customer Router.
<i>TDP</i>	Tag Distribution Protocol.
<i>Tunnel label</i>	Refers to the label provided to the ingress LSR by the directly connected LSR which is used by the ingress LER to route a packet to a specific destination address. The imposition router uses this label to switch the packet over a LSP to the destination router. This term is used in
<i>VC</i>	Virtual Circuit. In case of L2 transport over MPLS, a VC is configured between two different interfaces on a/may LERs. The VC is used to define a point to point circuit over which Layer 2 PDUs are transported. Since a uni-directional LSP is used to transport PDUs in each direction, a VC requires 2 LSPs to carry layer 2 traffic back and forth across the Virtual Circuit.

VC label

Refers to the label provided from the disposition LER. The imposition LER prepends this label so that the disposition router knows which output interface and VC to route a packet to. When the packet travels over a MPLS backbone, the tunnel label is popped by the neighboring LSR prior to passing it to the egress LER. The egress LER uses the VC label to identify the VC and output interface that the packet belongs to.

Virtual Circuit Identifier (VCID) A 32 bit identifier used to uniquely identify a VC. The VCID is used by the LER(s) to identify endpoints of a VCs. The VCID must be unique per tunnel for each transport type.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy Les Moulineaux
Cedex 9
France
www.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 17, 99 Walker Street
North Sydney
NSW 2059 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco.com Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE
Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico
The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia
Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2001 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement. Printed in the USA. AccessPath, AtmDirector, Browse with Me, CCDA, CCDE, CCDP, CCIE, CCNA, CCNP, CCSI, CD-PAC, *CiscoLink*, the Cisco NetWorks logo, the Cisco Powered Network logo, Cisco Systems Networking Academy, the Cisco Systems Networking Academy logo, Discover All That's Possible, Fast Step, Follow Me Browsing, FormShare, FrameShare, GigaStack, IGX, Internet Quotient, IP/VC, IQ Breakthrough, IQ Expertise, IQ FastTrack, the IQ Logo, IQ Net Readiness Scorecard, MGX, the Networkers logo, *Packet*, PIX, RateMUX, ScriptBuilder, ScriptShare, SlideCast, SMARTnet, TransPath, Voice LAN, Wavelength Router, WebViewer are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Empowering the Internet Generation, are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastSwitch, IOS, IP/TV, LightStream, MICA, Network Registrar, Post-Routing, Pre-Routing, Registrar, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. or its affiliates in the U.S. and certain other countries.

All other brands, names, or trademarks mentioned in this document or Web site 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. (0101R)