



INTRODUCTION TO MPLS

SESSION RST-1601

RST-1601
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

1

Agenda

Cisco.com

- **Background**
- **Technology Basics**
 - What Is MPLS? Where Is it Used?
- **Label Distribution in MPLS Networks**
 - LDP, RSVP, BGP
- **Building MPLS-Based Services**
 - VPNs
 - AToM
 - Traffic Engineering
- **Configurations**
 - Configuring MPLS, LDP, TE
- **Summary**

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

2

BACKGROUND



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

3

Terminology

Cisco.com

- **Acronyms**
 - PE—Provider Edge router
 - P—Provider Core router
 - CE—Customer Edge router (also referred to as CPE)
 - ASBR—Autonomous System Boundary Router
 - RR—Route Reflector
- **TE—Traffic Engineering**
 - TE head end—router that initiates a TE tunnel
 - TE midpoint—router where the TE tunnel transits
- **VPN—Collection of sites that share common policies**
- **AToM—Any Transport over MPLS**
 - Commonly known scheme for building layer 2 circuits over MPLS
 - Attachment circuit—layer 2 circuit between PE and CE
 - Emulated circuit—pseudowire between PEs

RST-1061
9776_05_2004_c1

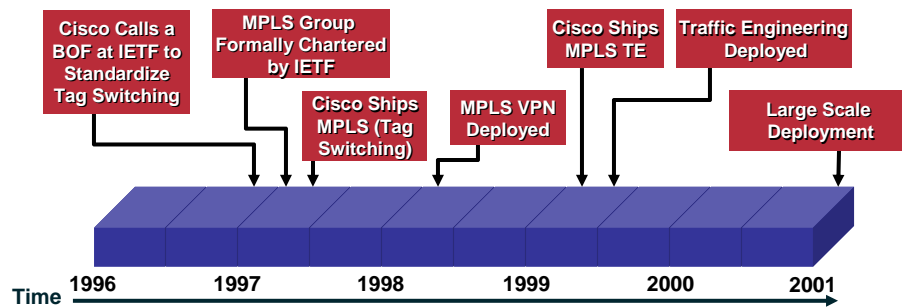
© 2004 Cisco Systems, Inc. All rights reserved.

4

Evolution of MPLS

Cisco.com

- From tag switching
- Proposed in IETF—later combined with other proposals from IBM (ARIS), Toshiba (CSR)



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

5

What Is MPLS?

Cisco.com

- **M**ulti **P**rotocol **L**abel **S**witching
- MPLS is an efficient encapsulation mechanism
- Uses “Labels” appended to packets (IP packets, AAL5 frames) for transport of data
- MPLS packets can run on other Layer 2 technologies such as ATM, FR, PPP, POS, Ethernet
- Other Layer 2 technologies can be run over an MPLS network
- Labels can be used as designators
 - For example—IP prefixes, ATM VC, or a bandwidth guaranteed path
- MPLS is a technology for delivery of IP Services

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

6

Original Motivation of MPLS

Cisco.com

- Allow core routers/networking devices to switch packets based on some simplified header
- Provide a highly scalable mechanism that was topology driven rather than flow driven
- Leverage hardware so that simple forwarding paradigm can be used
- It has evolved a long way from the original goal
 - Hardware became better and looking up longest best match was no longer an issue
 - By associating labels with prefixes, groups of sites or bandwidth paths or light paths new services such as MPLS VPNs and traffic engineering, GMPLS were now possible


RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

7

MPLS as a Foundation for Value-Added Services

Cisco.com



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

8

TECHNOLOGY BASICS




RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

9

Label Header for Packet Media

Cisco.com



Label = 20 Bits

COS/EXP = Class of Service, 3 Bits

S = Bottom of Stack, 1 Bit

TTL = Time to Live, 8 Bits

- Can be used over Ethernet, 802.3, or PPP links
- Uses two new Ethertypes/PPP PIDs
- Contains everything needed at forwarding time
- One word per label

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

10

Encapsulations

Cisco.com

**PPP Header
(Packet over SONET/SDH)**




One or More Labels Appended to the Packet

LAN MAC Label Header



ATM MPLS Cell Header



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

11

Forwarding Equivalence Class

Cisco.com

**Determines How Packets
Are Mapped to LSP**

- IP prefix/host address
- Layer 2 circuits (ATM, FR, PPP, HDLC, Ethernet)
- Groups of addresses/sites—VPN x
- A bridge/switch instance—VSI
- Tunnel interface—traffic engineering

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

12

MPLS Concepts


Cisco.com

At Edge:

- Classify packets
- Label them

Label Imposition

Edge Label
Switch Router
(ATM Switch or
Router)



Label Distribution Protocol

In Core:

- Forward using labels (as opposed to IP addr)
 - Label indicates service class and destination
- Label Swapping or Switching

At Edge:
Remove Labels and
Forward Packets
Label Disposition

- Label Switch Router (LSR)
- Router
 - ATM switch + Label Switch Controller

- Create new services via flexible classification
- Provide the ability to setup bandwidth guaranteed paths
- Enable ATM switches to act as routers

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

13

MPLS Operation

Cisco.com


1a. Existing Routing Protocols (e.g. OSPF, IS-IS)
Establish Reachability to Destination Networks

1b. Label Distribution Protocol (LDP)
Establishes Label to Destination
Network Mappings

4. Edge LSR at
Egress Removes
Label and Delivers
Packet

2. Ingress Edge LSR Receives Packet,
Performs Layer 3 Value-Added
Services, and "Labels" Packets

3. LSR Switches Packets
Using Label Swapping



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

14

LABEL DISTRIBUTION IN MPLS NETWORKS



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

15

Unicast Routing Protocols

Cisco.com

- **OSPF, IS-IS, BGP are needed in the network**
- **They provide reachability**
- **Label distribution protocols distribute labels for prefixes advertised by unicast routing protocols using**
 - Either a dedicated Label Distribution Protocol (LDP)**
 - Extending existing protocols like BGP to distribute labels**

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

16

Label Distribution Protocol

Cisco.com

- Defined in RFC 3035 and 3036
- Used to distribute labels in a MPLS network
 - Uses a TCP session—multiple sessions require multiple TCP sessions
- Forwarding equivalence class
 - How packets are mapped to LSPs (Label Switched Paths)
- Advertise labels per FEC
 - Reach destination a.b.c.d with label x
- Discovery

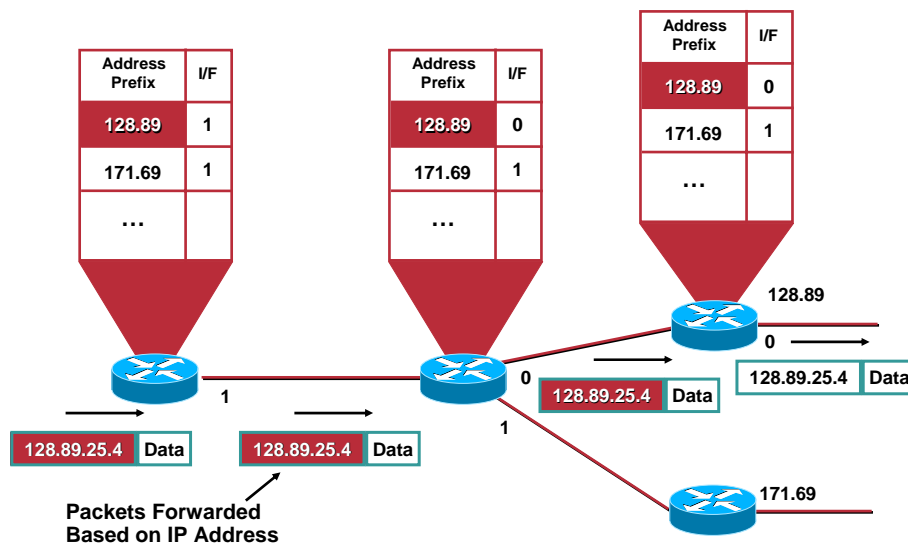
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

17

Router Example: Forwarding Packets

Cisco.com



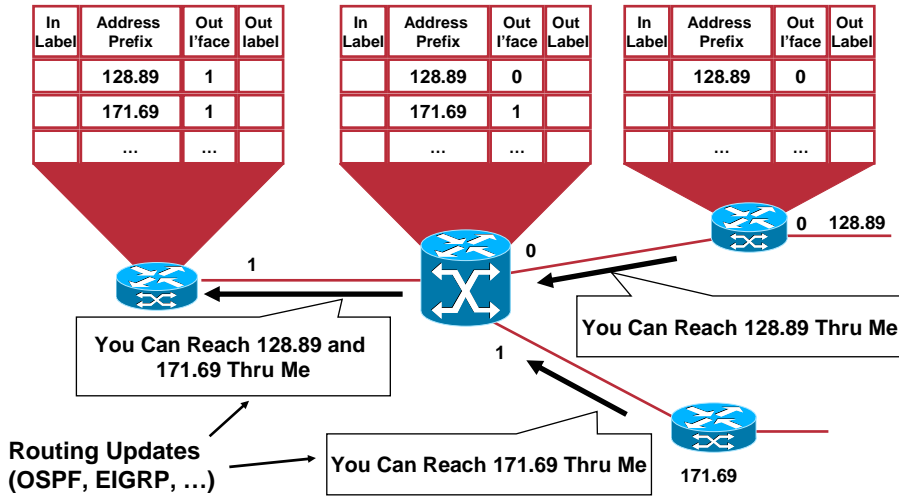
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

18

MPLS Example: Routing Information

Cisco.com



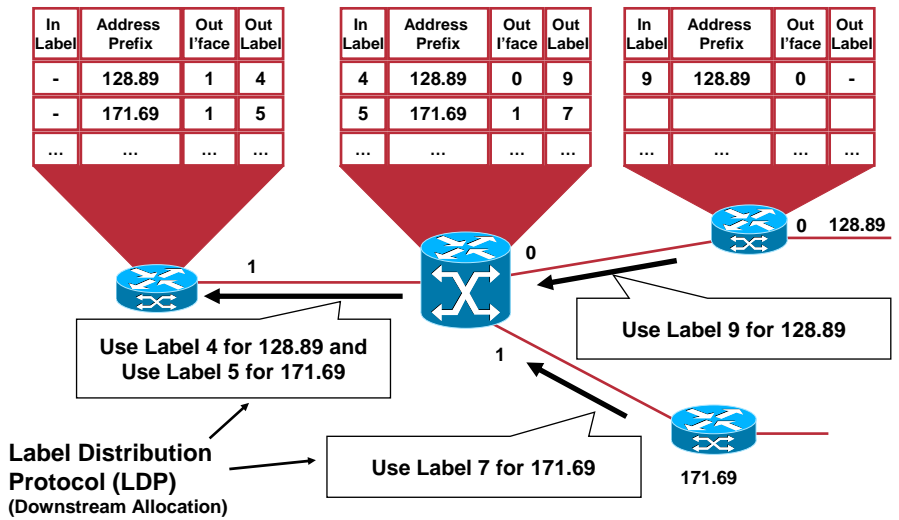
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

19

MPLS Example: Assigning Labels

Cisco.com



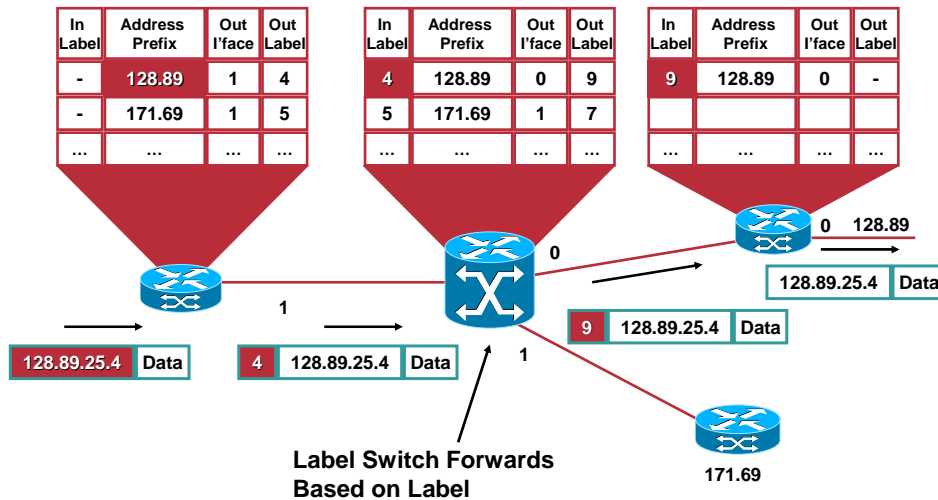
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

20

MPLS Example: Forwarding Packets

Cisco.com



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

21

Label Distribution Modes

Cisco.com

- **Downstream unsolicited**
Downstream node just advertises labels for prefixes/FEC reachable via that device
Previous example
- **Downstream on-demand**
Upstream node requests a label for a learnt prefix via the downstream node
Next example—ATM MPLS

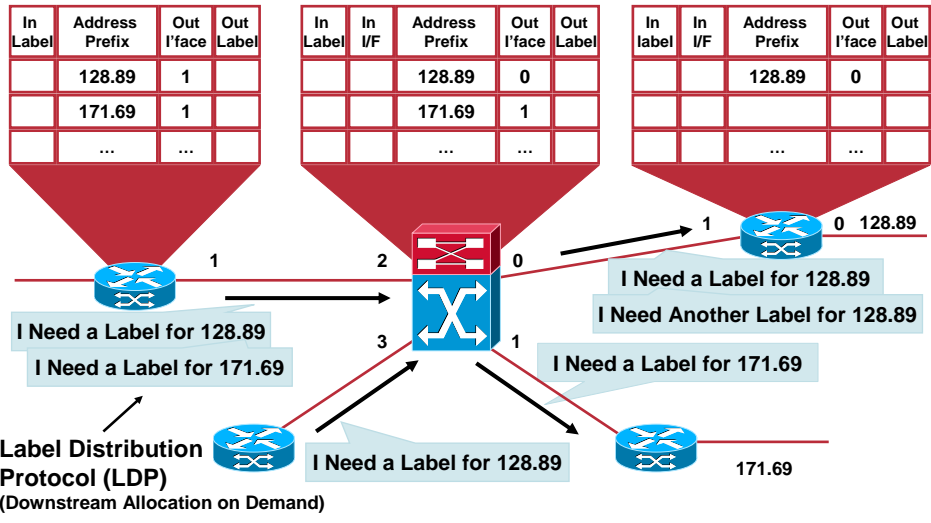
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

22

ATM MPLS Example: Requesting Labels

Cisco.com



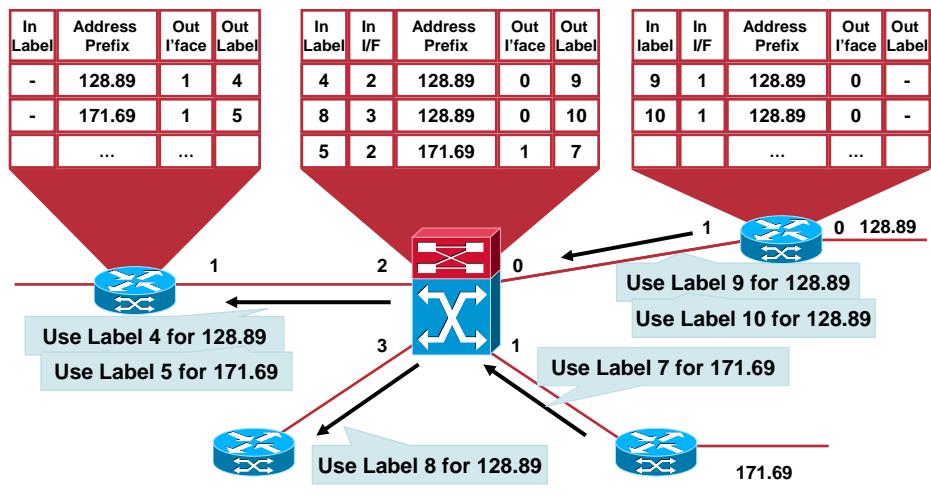
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

23

ATM MPLS Example: Assigning Labels

Cisco.com



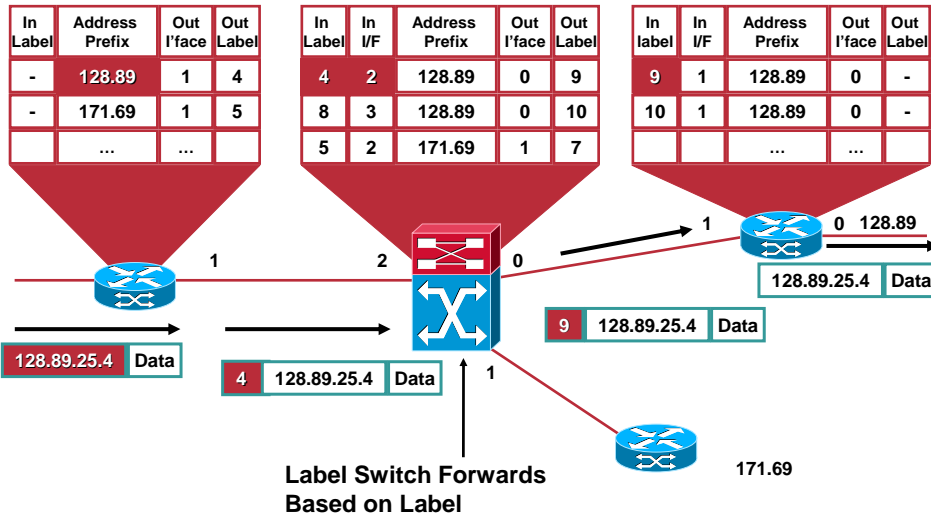
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

24

ATM MPLS Example: Packet Forwarding

Cisco.com



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

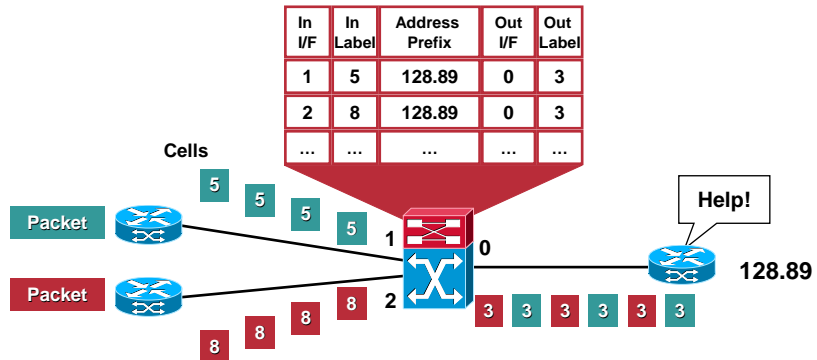
25

Why Multiple Labels with ATM?

Cisco.com

If Didn't Allocate Multiple Labels:

- Cells of different packets would have same label (VPI/VCI)
- Egress router can't reassemble packets



RST-1061
9776_05_2004_c1

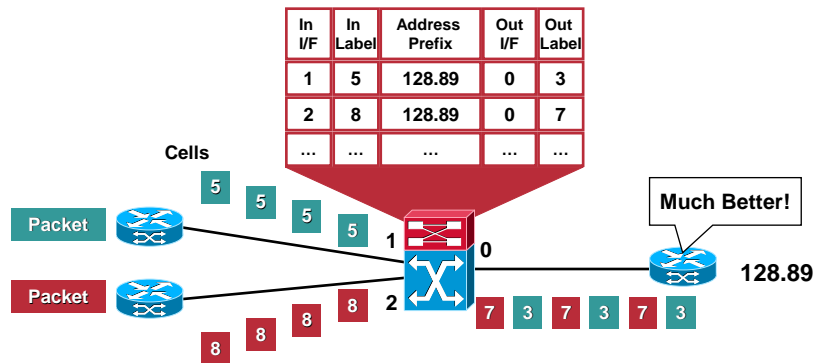
© 2004 Cisco Systems, Inc. All rights reserved.

26

Multiple Labels

Cisco.com

Multiple Labels Enables Edge Router to Reassemble Packets Correctly



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

27

Label Distribution Protocol

Cisco.com

Label Merge

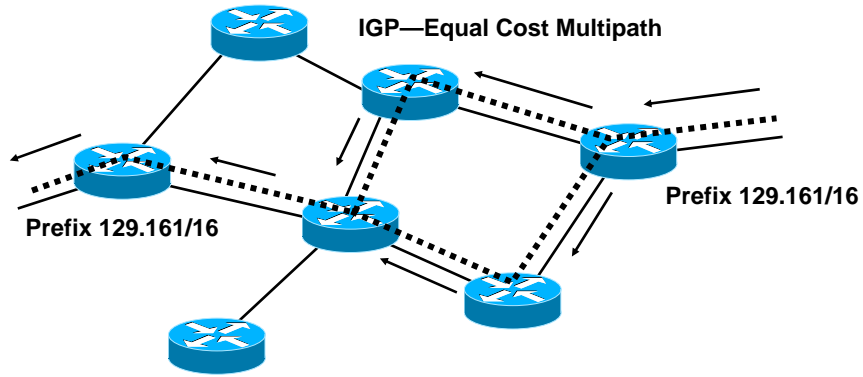
- Done by default for packet networks—unique label advertised per FEC
- Requires VC merge for ATM networks

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

28

LDP: Label Merge



Labels for Prefix 129.161 Are Advertised Along Both Paths

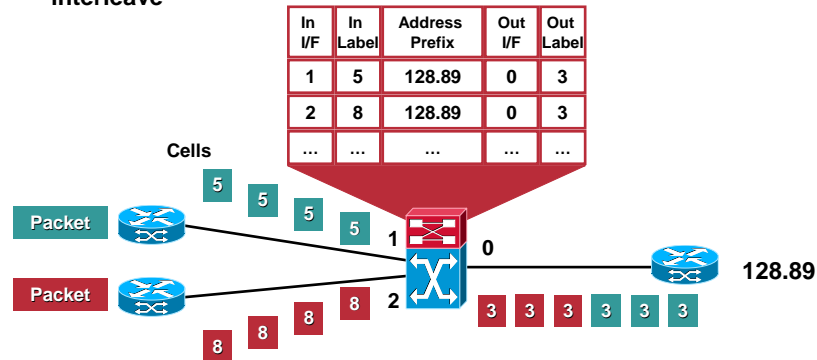
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

VC Merge

With ATM Switch that Can Merge VC's:

- Can reuse outgoing label
- Fewer labels required
- Hardware prevents cell interleave
- For very large networks



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

Label Retention Modes

Cisco.com

- **In downstream unsolicited mode—label mapping advertisements are received for all routes from all peers**
- **Liberal label retention**
 - These mappings are retained regardless of whether the LSR is the next hop for the advertised mapping
 - Once Labels are allocated to a prefix these labels are retained
 - Reaction to routing changes is fast
- **Conservative label retention**
 - Used with DOD mode
 - Label mappings are retained only if they are used to forward packets
 - Can save some label space—however, reacts slower to changes

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

31

Label Allocation Modes

Cisco.com

- **Independent mode**
 - Labels are allocated independently of neighbors' bindings
 - As long as the router has routes—it allocates a label irrespective of the neighbor
- **Ordered mode**
 - Labels are allocated only after the bindings from neighbors are received
 - Takes care of propagation delays in routing changes

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

32

LDP

Cisco.com

- **Neighbor discovery**
 - Discover directly attached neighbors—pt-to-pt links (including Ethernet)
 - Establish a session
 - Exchange prefix/FEC and label information
- **Extended neighbor discovery**
 - Establish peer relationship with another router that is not a neighbor
 - Exchange FEC and label information
 - May be needed to exchange service labels

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

33

TDP and LDP

Cisco.com

- **Tag distribution protocol—Cisco proprietary**
 - Pre-cursor to LDP
 - Used for Cisco tag switching
- **TDP and LDP supported on the same device**
 - Per neighbor/link basis
 - Per target basis
- **LDP is a superset of TDP**
- **Uses the same label/TAG**
- **Has different message formats**

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

34

Other Label Distribution Protocols: RSVP

Cisco.com

- Used in MPLS traffic engineering
- Additions to RSVP signaling protocol
- Leverage the admission control mechanism of RSVP to create an LSP with bandwidth
- Label requests are sent in PATH messages and binding is done with RESV messages
- EXPLICIT-ROUTE object defines the path over which setup messages should be routed
- Using RSVP has several advantages

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

35

Other Label Distribution Protocols: BGP

Cisco.com

- Used in the context of MPLS VPNs
- Need multiprotocol extensions to BGP
- Routers need to be BGP peers
- Label mapping info carried as part of NLRI (Network Layer Reacheability Information)

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

36

MPLS Control and Forwarding Planes

Cisco.com

- Control plane used to distribute labels—BGP, LDP, RSVP
- Forwarding plane consists of label imposition, swapping and disposition—no matter what the control plane
- Key: There is a separation of control plane and forwarding plane

Basic MPLS: destination-based unicast
 Labels divorce forwarding from IP address
 Many additional options for assigning labels
 Labels define destination and service



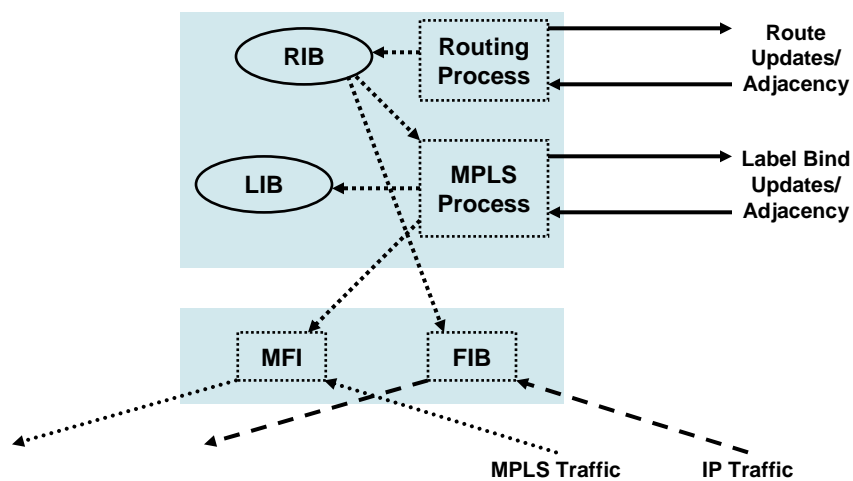
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

37

Control and Forward Plane Separation

Cisco.com



RST-1061
9776_05_2004_c1


© 2004 Cisco Systems, Inc. All rights reserved.

38

Label Stacking

Cisco.com

- There may be more than one label in an MPLS packet
- As we know labels correspond to forwarding equivalence classes
 - Example—there can be one label for routing the packet to an egress point and another that separates a customer A packet from customer B
 - Inner labels can be used to designate services/FECs, etc.
 - E.g. VPNs, fast reroute
- Outer label used to route/switch the MPLS packets in the network
- Last label in the stack is marked with EOS bit
- Allows building services such as
 - MPLS VPNs
 - Traffic engineering and fast re-route
 - VPNs over traffic engineered core
 - Any transport over MPLS



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

39

MPLS VPNS LAYER 2 AND LAYER 3



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

40

What Is a VPN?

Cisco.com

- VPN is a **set of sites** which are allowed to **communicate** with each other
- VPN is defined by a **set of administrative policies**
 - Policies determine both connectivity and QoS among sites
 - Policies **established by VPN customers**
 - Policies could be **implemented** completely by **VPN service providers**
 - Using BGP/MPLS VPN mechanisms

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

41

What Is a VPN? (Cont.)

Cisco.com

- **Flexible inter-site connectivity**
 - Ranging from complete to partial mesh
- **Sites may be either within the same or in different organizations**
 - VPN can be either intranet or extranet
- **Site may be in more than one VPN**
 - VPNs may overlap
- **Not all sites have to be connected to the same service provider**
 - VPN can span multiple providers

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

42

VPNs

Cisco.com

- **Layer 2 VPNs**

- Customer endpoints (CPE) connected via Layer 2 such as Frame Relay DLCI, ATM VC or point-to-point connection
 - If it connects IP routers then peering or routing relationship is between the endpoints
 - Multiple logical connections (one with each endpoint)

- **Layer 3 VPNs**

- Customer end points peer with provider routers
 - Single peering relationship
 - No mesh of connections
 - Provider network responsible for
 - Distributing routing information to VPN sites
 - Separation of routing tables from one VPN to another

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

43

LAYER 3 VPNs



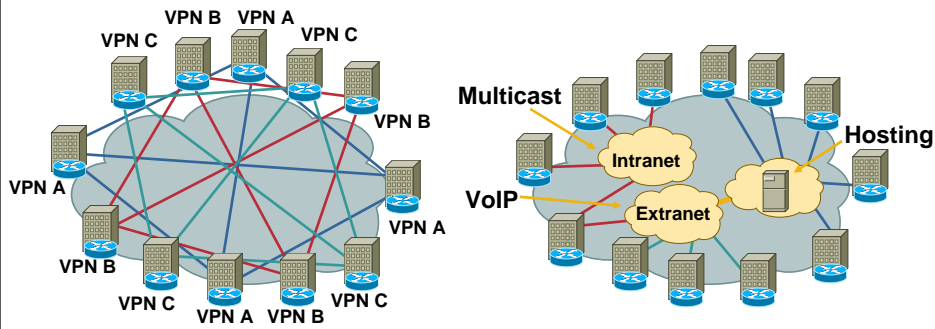
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

44

Service Provider Benefits of MPLS-Based VPNs

Cisco.com



Overlay VPN

- Pushes content outside the network
- Costs scale exponentially
- Transport dependent
- Groups endpoints, not groups
- Complex overlay with QoS, tunnels, IP

MPLS-Based VPNs

- Enables content hosting **inside** the network
- "Flat" cost curve
- Transport independent
- Easy grouping of users and services
- Enables QoS inside the VPNs


RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

45

Using Labels to Build an IP VPN

Cisco.com



- **The network distributes labels to each VPN**
 - Only labels for other VPN members are distributed
 - Each VPN is provisioned automatically by IP routing
- **Privacy and QoS of ATM without tunnels or encryption**
 - Each **network** is as secure as a Frame Relay **connection**
- **One mechanism (labels) for QoS and VPNs—no tradeoffs**

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

46

How Does It Work?

Cisco.com

- **Simple idea**

Use a label to designate VPN prefix

Route that VPN packet to egress PE advertising that prefix

Use the IGP label to the VPN packet to the egress node

- **How is it done?**

Routers need to maintain separate VPN routing tables called VRFs (Virtual Routing and Forwarding Tables)

Routers then export and import routes using BGP extensions to identify and separate one VPNs routes from another

Routers then exchange labels for VPN routes in addition to IGP routes

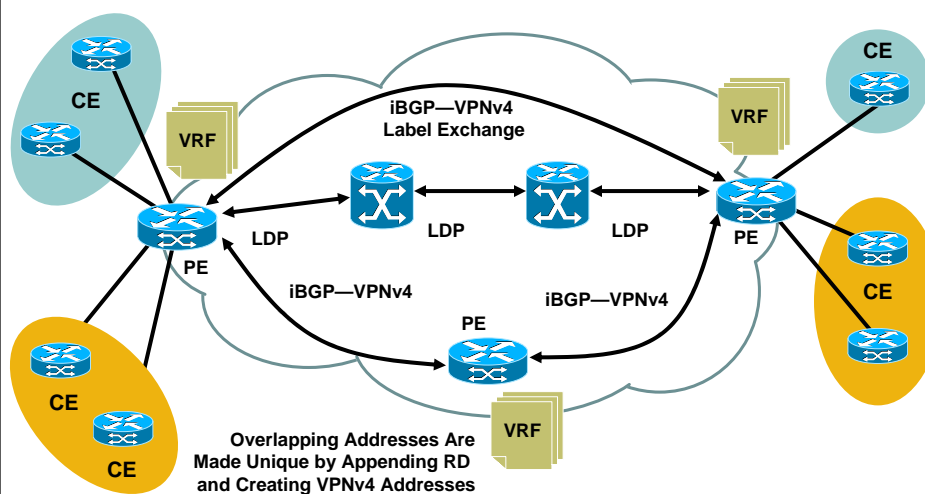
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

47

RFC 2547: MPLS VPNs

Cisco.com



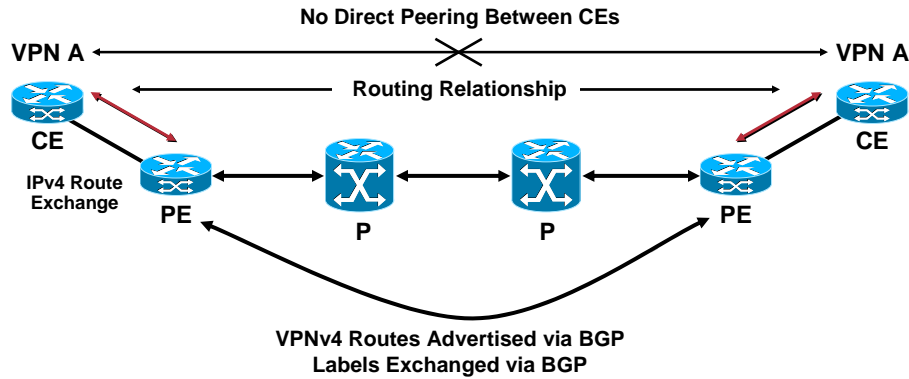
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

48

Control Plane Path

Cisco.com



- RD—8 Byte field—assigned by provider—significant to the provider network only
- VPNv4 Address: RD+VPN Prefix
- Unique RD per VPN makes the VPNv4 address unique

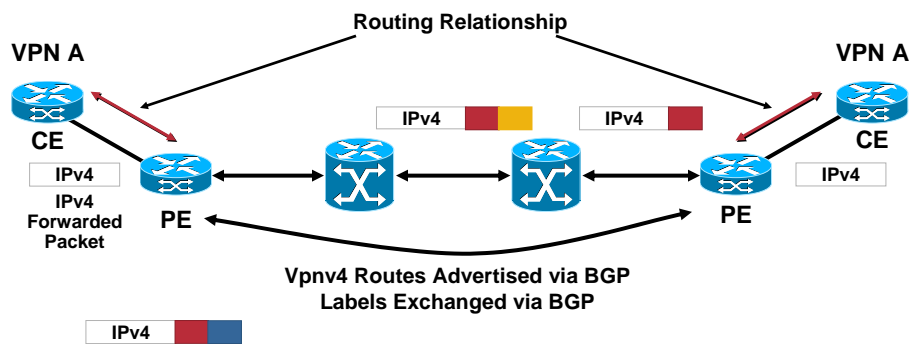
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

49

Data Plane Path

Cisco.com



- Ingress PE is imposing 2 labels

RST-1061
9776_05_2004_c1

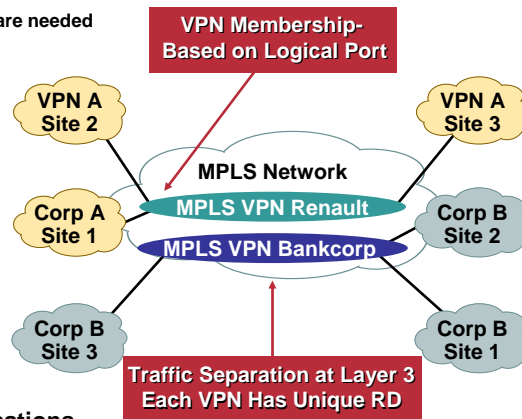
© 2004 Cisco Systems, Inc. All rights reserved.

50

MPLS-Based IP-VPN Architecture

Cisco.com

- **Scalable VPNs**
 - Add more PEs if more VPNs are needed
 - No N² mesh
 - VPNs are built in to the cloud
- **IP QoS and traffic engineering**
- **Easy to manage and no VC mesh provisioning required**
- **Provides a level of security/separation equivalent to Frame Relay and ATM**
- **Supports the deployment of new value-added applications**
- **Customer IP address freedom**



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

51

Key Features

Cisco.com

- **No constraints on addressing plans used by VPNs—a VPN customer may:**
 - Use globally unique and routable/non-routable addresses,
 - Use private addresses (RFC1918)
- **Security:**
 - Basic security is comparable to that provided by FR/ATM-based VPNs without providing data encryption
 - VPN customer may still use IPSec-based mechanisms
 - e.g., CE-CE IPSec-based encryption

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

52

Key Features (Cont.)

Cisco.com

- **Quality of Service:**

Flexible and scalable support for a CoS-based networks

- **Scalability:**

Total capacity of the system isn't bounded by the capacity of an individual component

Scale to virtually unlimited number of VPNs per VPN Service Provider and scale to thousands of sites per VPN

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

53

Key Features (Cont.)

Cisco.com

- **Connectivity to the Internet:**

VPN Service Provider may also provide connectivity to the Internet to its VPN customers

Common infrastructure is used for both VPN and the Internet connectivity services

- **Simplifies operations and management for VPN Service Providers:**

No need for VPN Service Providers to set up and manage a separate backbone or "virtual backbone" for each VPN

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

54

BGP/MPLS VPN: Summary

Cisco.com

- **Supports large scale VPN service**
- **Increases value add by the VPN Service Provider**
- **Decreases Service Provider cost of providing VPN services**
- **Mechanisms are general enough to enable VPN Service Provider to support a wide range of VPN customers**

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

55

LAYER 2 VPNS



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

56

Layer 2 VPNs

Cisco.com

Similar to L3VPN

- Designate a label for the circuit
- Exchange that label information with the egress PE
- Encapsulate the incoming traffic (Layer 2 frames)
- Apply label (learnt through the exchange)
- Forward the MPLS packet (I2 encapsulated to destination on an LSP)
- At the egress
 - Lookup the L2 label
 - Forward the packet onto the L2 attachment circuit

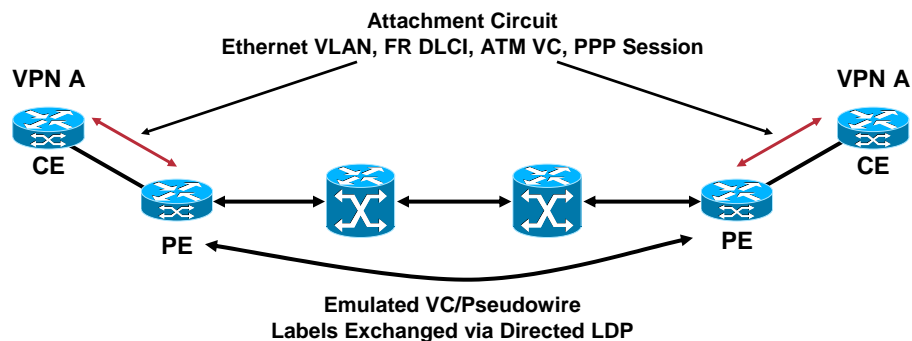
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

57

Architecture

Cisco.com



RST-1061
9776_05_2004_c1

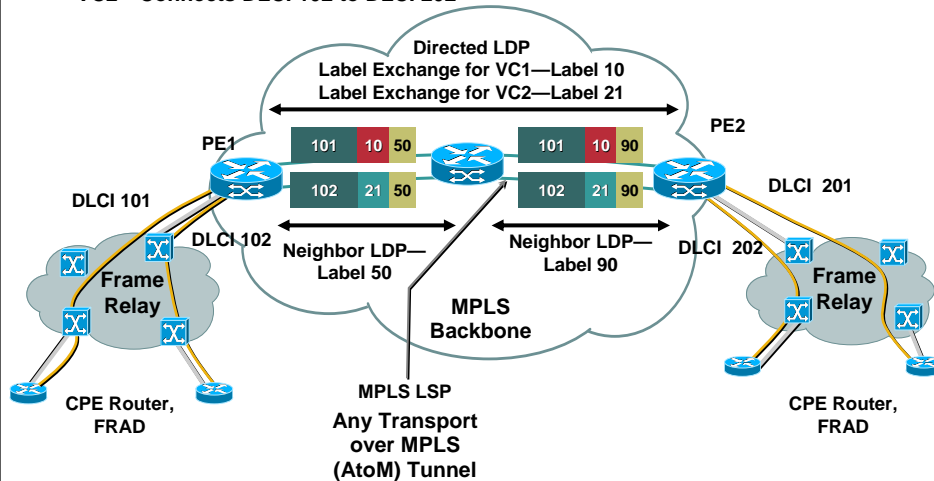
© 2004 Cisco Systems, Inc. All rights reserved.

58

Frame Relay over MPLS: Example

Cisco.com

VC1—Connects DLCI 101 to DLCI 201
VC2—Connects DLCI 102 to DLCI 202



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

59

Summary

Cisco.com

- Easy way of transporting Layer 2 frames
- Can be used to transport ATM AAL5 frames, cells, FR DLCI, PPP sessions, Ethernet VLANs
- Point-to-point transport with QoS guarantees
- Combine with TE and QoS to emulate Layer 2 service over a packet infrastructure
- Easy migration towards network convergence

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

60

MPLS TRAFFIC ENGINEERING



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

61

What Is MPLS Traffic Engineering?

Cisco.com

- **Process of routing data traffic in order to balance the traffic load on the various links, routers, and switches in the network**
- **Key in most networks where multiple parallel or alternate paths are available**

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

62

Why Traffic Engineering?

Cisco.com

- **Congestion in the network due to changing traffic patterns**
 - Election news, online trading, major sports events
- **Better utilization of available bandwidth**
 - Route on the non-shortest path
- **Route around failed links/nodes**
 - Fast rerouting around failures, transparently to users
 - Like SONET APS (Automatic Protection Switching)
- **Build new services—virtual leased line services**
 - VoIP toll-bypass applications, point-to-point bandwidth guarantees
- **Capacity planning**
 - TE improves aggregate availability of the network


RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

63

IP Routing and the Fish

Cisco.com



IP (Mostly) Uses Destination-Based Least-Cost Routing
Flows from R8 and R1 Merge at R2 and Become Indistinguishable
From R2, Traffic to R3, R4, R5 Use Upper Route

Alternate Path Under-Utilized

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

64


The Problem with Shortest-Path

Cisco.com

Node	Next-Hop	Cost
B	B	10
C	C	10
D	C	20
E	B	20
F	B	30
G	B	30

- Some links are **DS3**, some are **OC-3**
- Router A has 40Mb of traffic for Router F, 40Mb of traffic for Router G
- Massive (44%) packet loss at Router B->Router E!

Changing to A->C->D->E won't help



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.


65

How MPLS TE Solves the Problem

Cisco.com

Node	Next-Hop	Cost
B	B	10
C	C	10
D	C	20
E	B	20
F	Tunnel 0	30
G	Tunnel 1	30

- Router A sees all links
- Router A computes paths on properties other than just shortest cost
- **No link oversubscribed!**



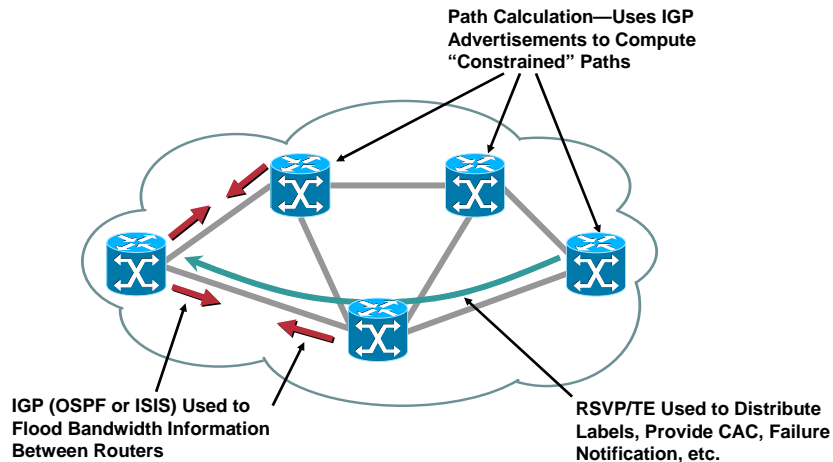
RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

66

TE Fundamentals: “Building Blocks”

Cisco.com



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

67

Information Distribution

Cisco.com

- You need a link-state protocol as your IGP
IS-IS or OSPF
- Link-state requirement is **only** for MPLS-TE!
Not a requirement for VPNs, etc!
- Why do I need a link-state protocol?
To make sure info gets flooded
To build a picture of the entire network
- Information flooded includes link, bandwidth, attributes, etc.


RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

68

Example

Cisco.com



- **PATH** messages are sent with requested bandwidth
- **RESV** messages are sent with label bindings for the TE tunnel
- Tunnels can be explicitly routes
- Admission control at each hop to see if the bandwidth requirement can be met
- Packets are mapped to the tunnel via
 - Static routed
 - Autoroute
 - Policy route
- Packets follow the tunnel—LSP

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

69

Benefits of TE over Policy Routing

Cisco.com

- **Policy Routing**
 - Hop-by-hop decision making
 - No accounting of bandwidth
- **Traffic Engineering**
 - Headend-based
 - Accounts for available link bandwidth
 - Admission control


RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

70

Applications of MPLS TE: MPLS Fast Reroute

Cisco.com



Mimic SONET APS
Reroute in 50ms or Less

- Multiple hops can be by-passed; R2 swaps the label which R4 expects before pushing the label for R6
- R2 locally patches traffic onto the link with R6


RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

71

Link Protection

Cisco.com



- Primary tunnel: A -> B -> D -> E
- Backup tunnel: B -> C -> D (preprovisioned)
- Recovery = ~50ms

*Actual Time Varies—Well Below 50ms in Lab Tests, Can Also Be Higher


RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

72

Node Protection

Cisco.com



- Primary tunnel: A -> B -> D -> E -> F
- Backup tunnel: B -> C -> E (pre-provisioned)
- Recovery = ~100ms



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

73

TE DEPLOYMENT SCENARIOS



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

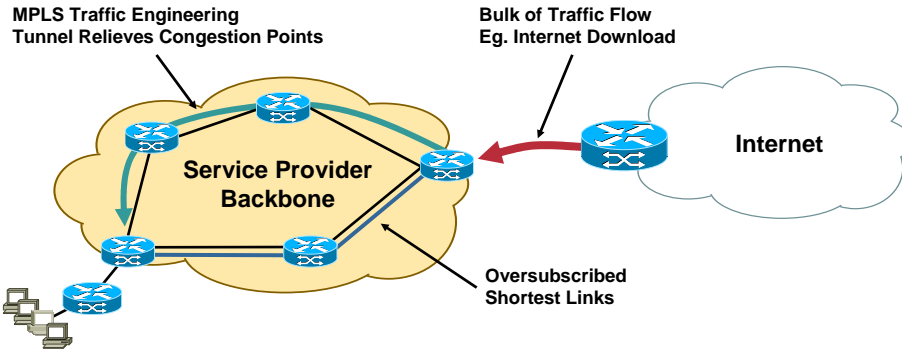
74

Tactical TE Deployment

Cisco.com

Requirement: Need to Handle Scattered Congestion Points in the Network

Solution: Deploy MPLS TE on Only Those Nodes That Face Congestion



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

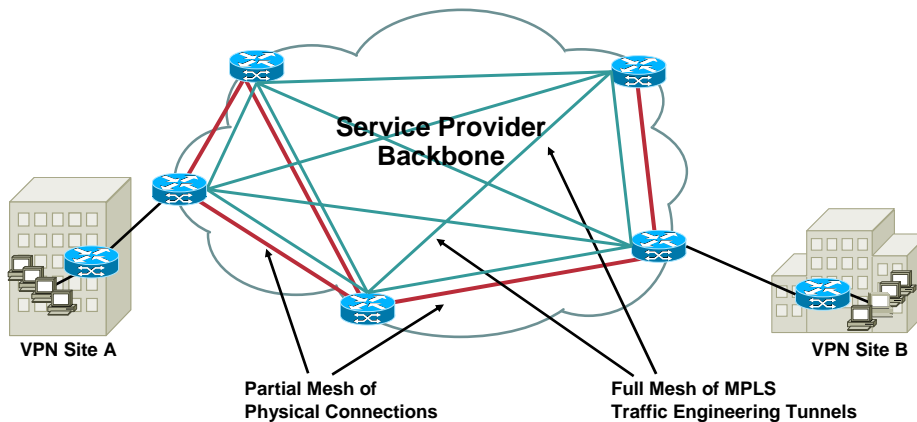
75

Full Mesh TE Deployment

Cisco.com

Requirement: Need to Increase "Bandwidth Inventory" Across the Network

Solution: Deploy MPLS TE with a Full Logical Mesh over a Partial Physical Mesh and Use Offline Capacity Planning Tool



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

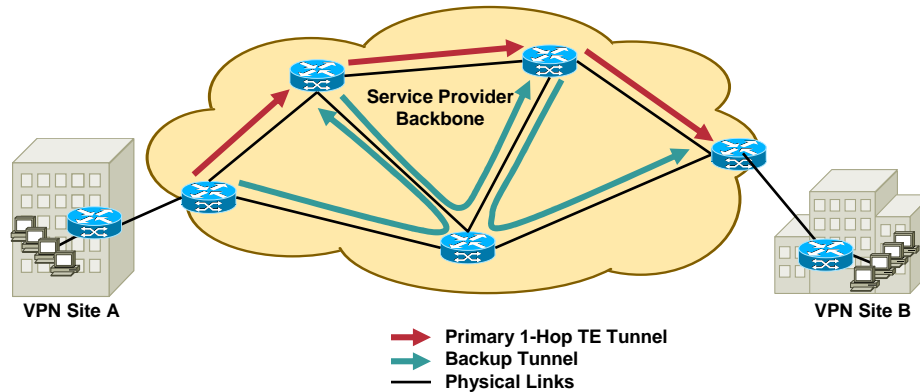
76

1-Hop TE Deployment

Cisco.com

**Requirement: Need Protection Only—Minimize Packet Loss
Lots of Bandwidth in the Core**

**Solution: Deploy MPLS Fast Reroute for Less than 50ms Failover Time with
1-Hop Primary TE Tunnels and Backup Tunnel for Each**



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

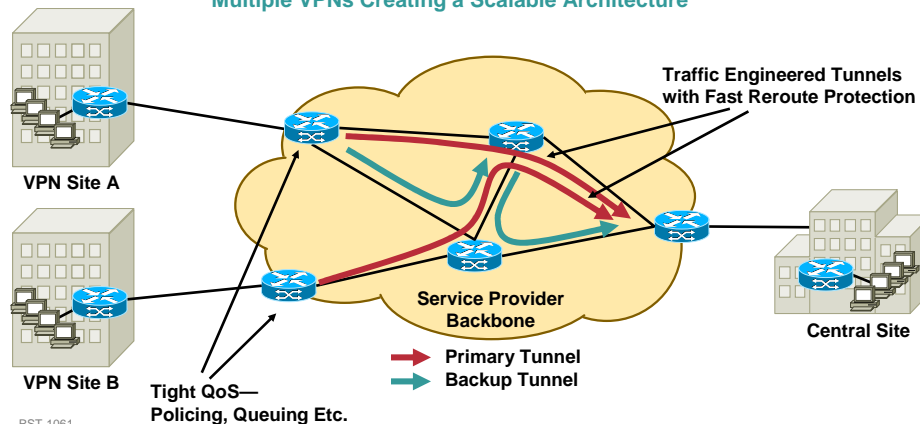
77

Virtual Leased Line Deployment

Cisco.com

**Requirement: Need to Create Dedicated Point-to-Point Circuits with Bandwidth
Guarantees—Virtual Leased Line (VLL)**

**Solution: Deploy MPLS TE (or DS-TE) with QoS; Forward Traffic from L3 VPN
or L2 VPN into a TE Tunnel; Unlike ATM PVCs, Use 1 TE Tunnel for
Multiple VPNs Creating a Scalable Architecture**



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

78

MPLS TE Summary

Cisco.com

- **Useful for rerouting traffic in congested environments**
- **Build innovative services like virtual leased line**
- **Build protection solutions using MPLS FRR**

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

79

MPLS and QoS

Cisco.com

- **IP QoS mechanisms are leveraged for MPLS**
- **IP Precedence bits copied into MPLS EXP field**
- **IP DSCP can be mapped to MPLS EXP**
- **WRED, queuing can be done on MPLS EXP**
- **DiffServ mechanisms provide similar per hop behavior with MPLS**

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

80

Configuring MPLS

Cisco.com

Step 1	Router# configure terminal	Enables Configuration Mode
Step 2	Router(config)# ip cef [distributed]	Configures Cisco Express Forwarding
Step 3	Router(config)# interface interface	Specifies the Interface to Configure
Step 4	Router(config-if)# mpls ip	Configures MPLS Hop-by-Hop Forwarding for a Specified Interface
Step 5	Router(config-if)# mpls label protocol ldp	Configures the Use of LDP for a Specific Interface; Sets the Default Label Distribution Protocol for the Specified Interface To Be LDP, Overriding Any Default Set by the Global MPLS Label Protocol Command
Step 6	Router# configure terminal Router(config)# mpls label protocol ldp	Configures the Use of LDP on All Interfaces; Sets the Default Label Distribution Protocol for All Interfaces To Be LDP

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

81

Show Commands

Cisco.com

```
Router# show mpls interfaces
Interface IP Tunnel Operational
Ethernet1/1/1 Yes (tdp) No No
Ethernet1/1/2 Yes (tdp) Yes No
Ethernet1/1/3 Yes (tdp) Yes Yes
POS2/0/0 Yes (tdp) No No
ATM0/0.1 Yes (tdp) No No (ATM labels)
ATM3/0.1 Yes (ldp) No Yes (ATM labels)
ATM0/0.2 Yes (tdp) No Yes
```

```
Router# show mpls ldp discovery
Local LDP Identifier:
118.1.1.1:0
Discovery Sources:
Interfaces:
POS2/0 (ldp): xmit/recv
LDP Id: 155.0.0.55:0
Tunnel1 (ldp): Targeted -> 133.0.0.33
Targeted Hellos:
118.1.1.1 -> 133.0.0.33 (ldp): active, xmit/recv
LDP Id: 133.0.0.33:0
118.1.1.1 -> 168.7.0.16 (tdp): passive, xmit/recv
TDP Id: 168.7.0.16:0
```

```
show mpls ip binding [vrf vpn-name] [network
{mask | length} [longer-prefixes]]
[local-label {atm vpi vci | label [- label]]}
[remote-label {atm vpi vci | label [- label]]}
[neighbor address] [local]
[interface interface] [generic | atm]
show mpls ip binding summary
```

```
Router# show mpls ip binding 194.44.44.0 24
194.44.44.0/24
in label: 24
in vc label: 1/37 lsr: 203.0.7.7:2 ATM1/0.8
Active egress (vcd 56)
out label: imp-null lsr: 155.0.0.55:0 inuse
Router#
```

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

82

TE: Configuration

Cisco.com

Step 1	Router(config-if)# mpls traffic-eng tunnels	Enables MPLS Traffic Engineering Tunnels on an Interface
Step 2	Router(config-if)# ip rsvp bandwidth bandwidth	Enables RSVP for IP on an Interface and Specifies the Amount of Bandwidth That Will Be Reserved; For a Description of the ip rsvp Command Syntax, See the <i>Quality of Service Solutions Command Reference</i>
Step 1	Router(config)# router ospf process-id	Configures an OSPF Routing Process for IP; You Are Placed in Router Configuration Mode; The <i>process-id</i> is an Internally Used Identification Parameter for an OSPF routing process; It Is Locally Assigned and Can Be Any Positive Integer; Assign a Unique Value for Each OSPF Routing Process
Step 2	Router(config- router)# mpls traffic-eng area 0	Turns on MPLS Traffic Engineering for OSPF Area 0
Step 3	Router(config- router)# mpls traffic-eng router-id loopback0	Specifies that the Traffic Engineering Router Identifier for The Node Is the IP Address Associated with Interface loopback0

Show mpls traffic-eng <->

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

83

SUMMARY




RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

84


MPLS: The Key Technology for L2 and L3 Service Delivery

Cisco.com



IP+ATM: MPLS Brings IP and ATM Together

- Eliminates IP "over" ATM overhead and complexity
- One network for Internet, Business IP VPNs, and transport



Network-Based VPNs with MPLS: a Foundation for Value-Added Service Delivery

- Flexible user and service grouping (biz-to-biz)
- Flexibility of IP and the QoS and privacy of ATM
- Enables application and content hosting inside each VPN
- Transport independent
- Low provisioning costs enable affordable managed services


RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

85


MPLS: The Key Technology for L2 and L3 Service Delivery

Cisco.com



MPLS Traffic Engineering

- Provides Routing on diverse paths to avoid congestion
- Better utilization of the network
- Better availability using Protection Solution (FRR)



Guaranteed Bandwidth Services

- Combine MPLS Traffic Engineering and QoS
- Deliver Point-to-point bandwidth guaranteed pipes
- Leverage the capability of Traffic Engineering
- Build Solution like Virtual leased line and Toll Trunking


RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

86


MPLS: The Key Technology for IP Service Delivery

Cisco.com



IP+Optical Integration

- Eliminates IP “over” Optical Complexity
- Uses MPLS as a control Plane for setting up lightpaths (wavelengths)
- One control plane for Internet, Business IP VPNs, and optical transport



Any Transport over MPLS

- Transport ATM, FR, Ethernet, PPP over MPLS
- Provide Services to existing installed base
- Protect Investment in the installed gear
- Leverage capabilities of the packet core
- Combine with other packet based services such as MPLS VPNs

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

87

Further Reading

Cisco.com

- <http://www.cisco.com/go/mpls>
- **MPLS and VPN Architectures—Jim Guichard, Ivan Papelnjak—Cisco Press®**
- **Traffic Engineering with MPLS—Eric Osborne, Ajay Simha—Cisco Press**

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

88

Q AND A



RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

89

Complete Your Online Session Evaluation!

Cisco.com

- WHAT:** Complete an online session evaluation and your name will be entered into a daily drawing
- WHY:** Win fabulous prizes! Give us your feedback!
- WHERE:** Go to the Internet stations located throughout the Convention Center
- HOW:** Winners will be posted on the onsite Networkers Website; four winners per day

RST-1061
9776_05_2004_c1

© 2004 Cisco Systems, Inc. All rights reserved.

90

CISCO SYSTEMS

