



2 NETWORKERS



3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

1



Introduction to Real-Time ATM Services

Session 3001



3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

2

“

Cisco's IP+ATM switches help Service Provider to build a **single multiservice** infrastructure optimized for **current services*** as well as new **IP-enabled services.**

”

* ATM, FR, private line

3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

3

Agenda

- **Multiservice Requirements and Benefits**
- **Technical Motivation for MPLS with ATM**
- **Applications and Transports**
- **Summary**

3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

4

Agenda

- **Multiservice Requirements and Benefits**
- Technical Motivation for MPLS with ATM
- Applications and Transports
- Summary

Myths

If bandwidth is free, why does my Internet access cost so much?

Which bandwidth?
Free, available

If ATM is dead... ,
why are Service Providers buying it?

If there is so much bandwidth everywhere,
then why can't I get any at my home



Even the LAN switches are shipping QoS!

The transition from scarcity to abundance is going to be long enough!

IP and ATM are competing technologies!



Reality



Multiservice Concept and Requirements

- Data requirements are driving network design, not voice
- Economics encourage network/service integration (“convergence”)
- Flexible bandwidth management (traffic engineering)
- Multiple service performance and predictability (availability, serviceability)

Market Monitor

“

The SPs are moving from managing multilayer networks to managing multiservice network ...

”

Joe Skorupa, Director at RHK

Benefits

- **Operations cost**
- **New/integrated services**
- **Best of technologies**
- **Resource management**

Agenda

- **Multiservice Requirements and Benefits**
- **Technical Motivation for MPLS with ATM**
- **Applications and Transports**
- **Summary**

Technical Motivation for MPLS with ATM

“

**No single technology today
provides all the performance
and flexibility needed today
by the WAN clients**

”

Technical Motivation for MPLS with ATM

	IP/Connectionless	ATM/Connection
Advantage	?	?
Challenge	?	?

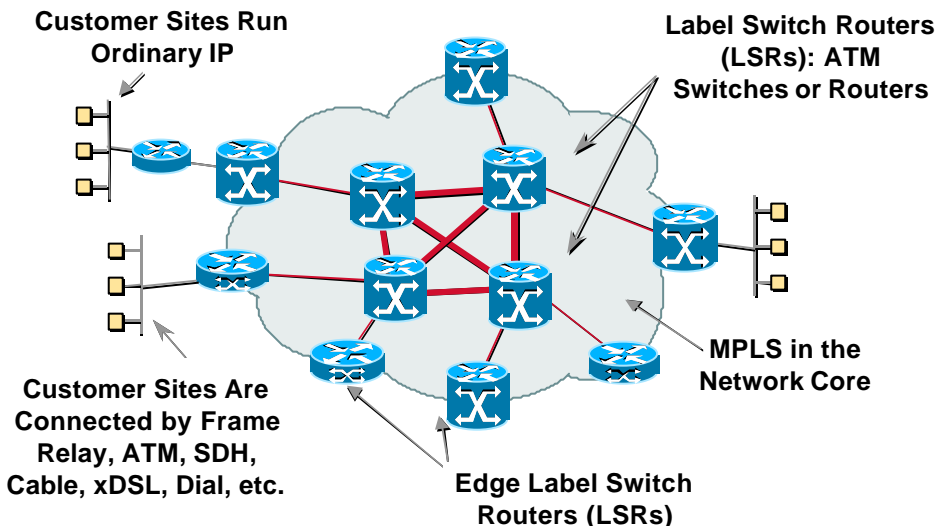
Technical Motivation for MPLS with ATM

- **Leverage** the existent
- **Migrating** the traditional infrastructure/services to IP-optimized networks

Using MPLS as the control plane

Multiprotocol Label Switching


MPLS Concepts



MPLS: Mechanisms

1a. Existing routing protocols (e.g. OSPF, IS-IS) establish reachability to destination networks.

1b. Label Distribution Protocol (LDP) summarizes destination information with labels



2. Ingress Edge LSR receives packet, performs Layer 3 value-added services, and "label" packets

3. Label Switch Routers (ATM switches or gigabit routers) switch labeled packets


4. Edge LSR at egress removes label and delivers packet

Non-ATM MPLS Example: Routing Information

In Lbl	Address Prefix	Out l'face	Out Lbl
	128.89	1	
	171.69	1	
	

In Lbl	Address Prefix	Out l'face	Out Lbl
	128.89	0	
	171.69	1	
	

In Lbl	Address Prefix	Out l'face	Out Lbl
	128.89	0	
	




(Example shows RIP updates for simplicity: OSPF or IS-IS are more likely. Routing tables for this network will be identical in any case.)

Non-ATM MPLS Example: Assigning Labels

In Lbl	Address Prefix	Out l'face	Out Lbl
-	128.89	1	4
-	171.69	1	5
...

In Lbl	Address Prefix	Out l'face	Out Lbl
4	128.89	0	9
5	171.69	1	7
...

In Lbl	Address Prefix	Out l'face	Out Lbl
9	128.89	0	-
...



3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com


17

Non-ATM MPLS Example: Forwarding Packets

In Lbl	Address Prefix	Out l'face	Out Lbl
-	128.89	1	4
-	171.69	1	5
...

In Lbl	Address Prefix	Out l'face	Out Lbl
4	128.89	0	9
5	171.69	1	7
...

In Lbl	Address Prefix	Out l'face	Out Lbl
9	128.89	0	-
...




3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

18

MPLS: Another View



The MPLS Protocol Stack Consists of IP Routing and MPLS Signaling, Together Setting up Connectivity at Layer 2. LDP Signaling Is Fairly Straightforward. Much of the Complexity in MPLS Signaling Is in IP Routing.

Carrying Labels


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



LAN MAC Tag Header




ATM Cell Header




Carrying Labels on ATM Links

ATM Cell Header



- Label information is carried in part of the VPI and VCI header on cells
- A each Label is a different Virtual Circuit Identifier (VCI)
- One or more VPIs must reserved for use by MPLS. Traditional ATM services may use VCs with other VPIs

MPLS Is an ATM Control Stack

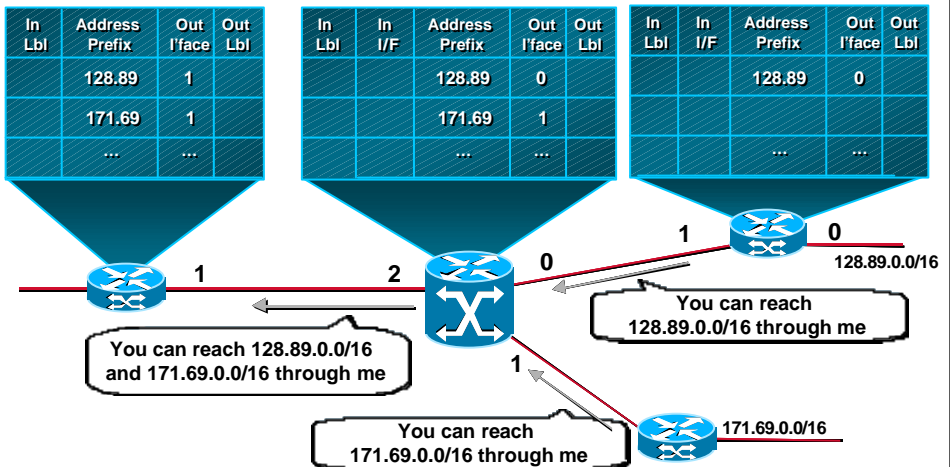


- MPLS and ATM Forum protocols can co-exist on the same links and switches
- ATM MPLS sets up Label VCs (LVCs) according to IP Routing information
- Label VCs (LVCs) are not PVCs or SVCs
- LVCs, SVCs and PVCs can coexist on same links in different VPIs

Label Distribution for ATM

- Uses LDP in “Downstream-on-Demand” mode
- To route a label ATM link, Virtual Channel (VC) labels are requested when topology changes
- Precedence associated with LVC
- Some LDP extensions for negotiation of ATM specific parameters

ATM MPLS Example: Routing Information




(Example shows RIP updates for simplicity: OSPF or IS-IS are more likely)

ATM MPLS Example: Requesting Labels

In Lbl	Address Prefix	Out I'face	Out Lbl
	128.89	1	
	171.69	1	
...

In Lbl	In I/F	Address Prefix	Out I'face	Out Lbl
		128.89	0	
		171.69	1	
...

In Lbl	In I/F	Address Prefix	Out I'face	Out Lbl
		128.89	0	
...



3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com


25

ATM MPLS Example: Assigning Labels

In Lbl	Address Prefix	Out I'face	Out Lbl
-	128.89	1	4
-	171.69	1	5
...

In Lbl	In I/F	Address Prefix	Out I'face	Out Lbl
4	2	128.89	0	9
8	3	128.89	0	10
5	2	171.69	1	7

In Lbl	In I/F	Address Prefix	Out I'face	Out Lbl
9	1	128.89	0	-
10	1	128.89	0	-
...



3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

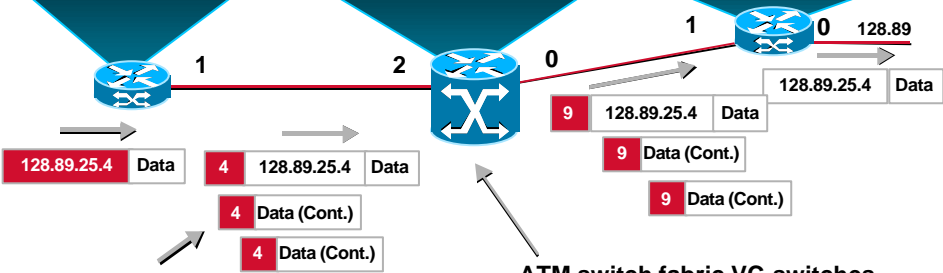
26

ATM MPLS Example: Packet Forwarding

In Lbl	Address Prefix	Out I'face	Out Lbl
-	128.89	1	4
-	171.69	1	5
...

In Lbl	In I/F	Address Prefix	Out I'face	Out Lbl
4	2	128.89	0	9
8	3	128.89	0	10
5	2	171.69	1	7

In Lbl	In I/F	Address Prefix	Out I'face	Out Lbl
9	1	128.89	0	-
10	1	128.89	0	-
...



Each distinct label is a distinct VCI. Label 4 = VCI 4

ATM switch fabric VC-switches the cells of the packet

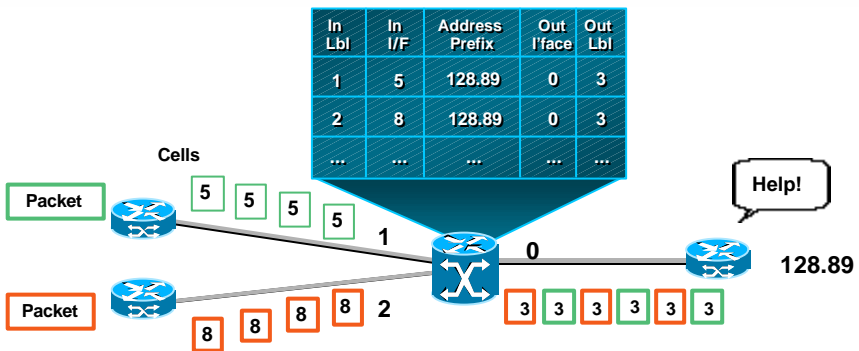
3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

27

Why Multiple Labels with ATM?



- If didn't allocate multiple labels:

Cells of different packets would have same label (VPI/VCI)

Egress router can't reassemble packets

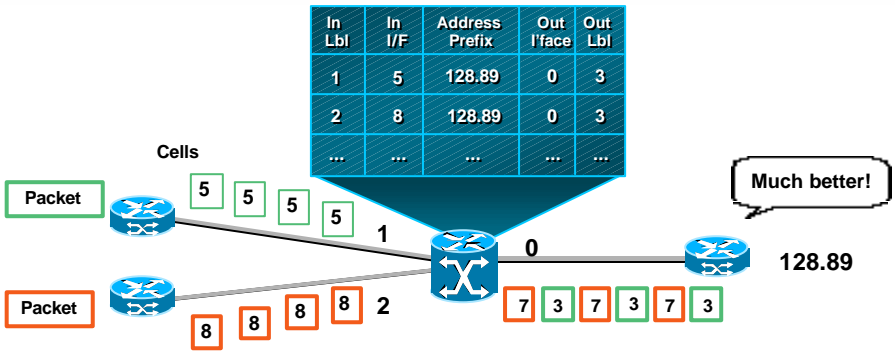
3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

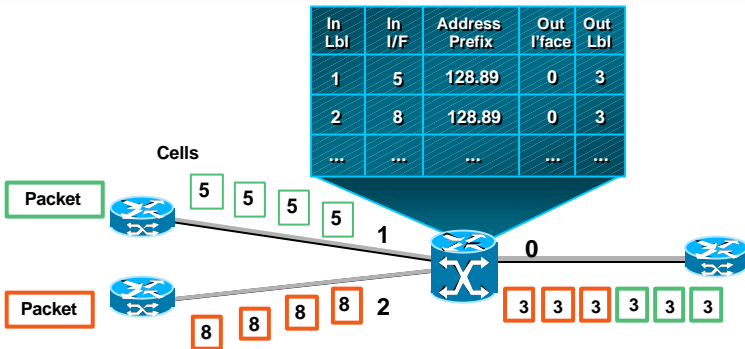
28

Multiple Labels



- Multiple labels enables edge router to reassemble packets correctly

VC Merge



- With ATM switch that can merge VCs:
 - Can reuse outgoing label
 - Hardware prevents cell interleave
 - Fewer labels required



Delivering IP+ATM Platforms

3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

31

Delivering IP+ATM Platforms



- **Cisco BPX 8650 with Label Switch Controller adds IP and MPLS support**
- **Installed Cisco BPX 8620 base can be field-upgraded to support MPLS**



- **Cisco MGX 8850 contains Route Processor Modules (RPMs) for IP**
- **Full range of L2/L3 services**
- **Universal chassis for flexibility**



3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

32

Cisco BPX 8650



Cisco BPX 8620



Label Switch
Controller



Cisco BPX
8650 ATM
LSR

- The Cisco BPX 8650 combines a Cisco BPX 8620 ATM switch with a Label Switch Controller based on a Cisco 7204 router
- The Cisco BPX 8650 has been shipping with MPLS since August 1998

Cisco MGX 8800



Cisco MGX 8800



RPM (LSC/LSR)



Cisco MGX 8800
IP+ATM Switch

- The Cisco MGX 8800 IP+ATM switch incorporates:
 - A multiservice access switch with many different interfaces: Frame Relay, ATM, serial, voice, etc
 - Full-featured routers with Edge LSR function
 - Also, full LSR function
 - Full shipment of MPLS function since 1999



MPLS on Cisco BPX 8650— Theory of Operation

3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

35

Cisco BPX 8600 Switching

- **16x16 switch fabric, 9.6 Gbps–19.2 Gbps for multicast**
- **A pair of BCC cards in a hot standby:**
 - Close to zero data loss occurs in case of failure
- **An online switch card permanently connects the 14 cards to each other. It has $14 \times 14 = 196$ connections. The switch fabric **is not programmed for each VC connection****

3001
1237_05_2000_c2

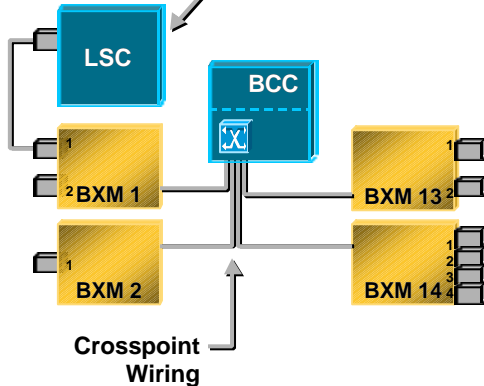
© 2000, Cisco Systems, Inc.

cisco.com

36

Setting Up an LVC Cross-Connect: 0

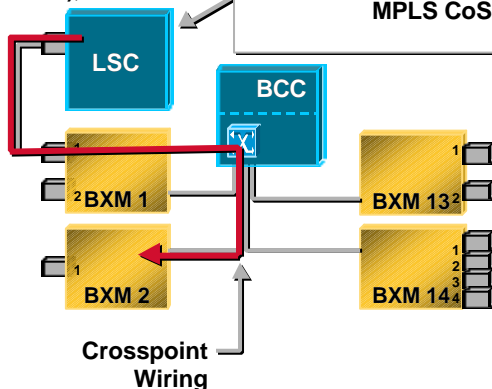
LDP tells me I want a connection with MPLS CoS 3 between (Card 2, Port 1, VPI 22, VCI 1332) and (C 13, P 2, VPI 56, VCI 3456)



Setting Up an LVC Cross-Connect: 1

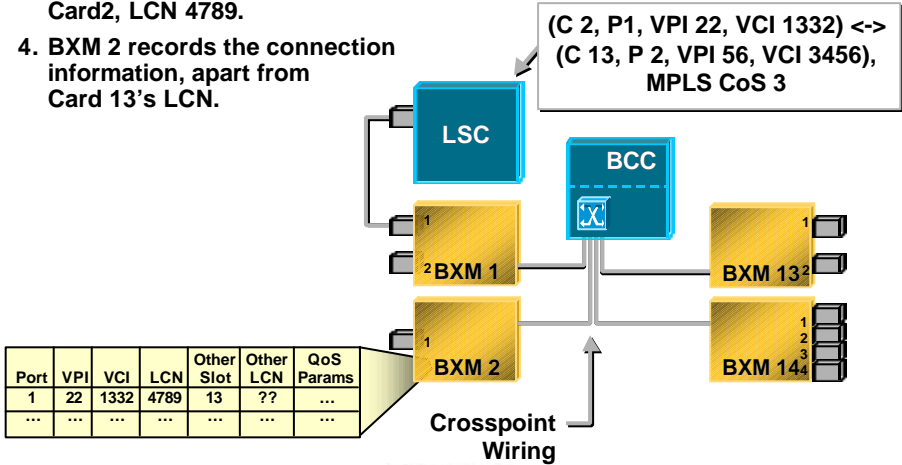
1. LSC can send the connection message to either Card 2 or Card 13. It will pick one at random, say Card 2.
2. LSC sends a VSI message to Card 2, requesting a connection: (C 2, P1, VPI 22, VCI 1332) <-> (C 13, P 2, VPI 56, VCI 3456), MPLS CoS 3

(C 2, P1, VPI 22, VCI 1332) <-> (C 13, P 2, VPI 56, VCI 3456), MPLS CoS 3



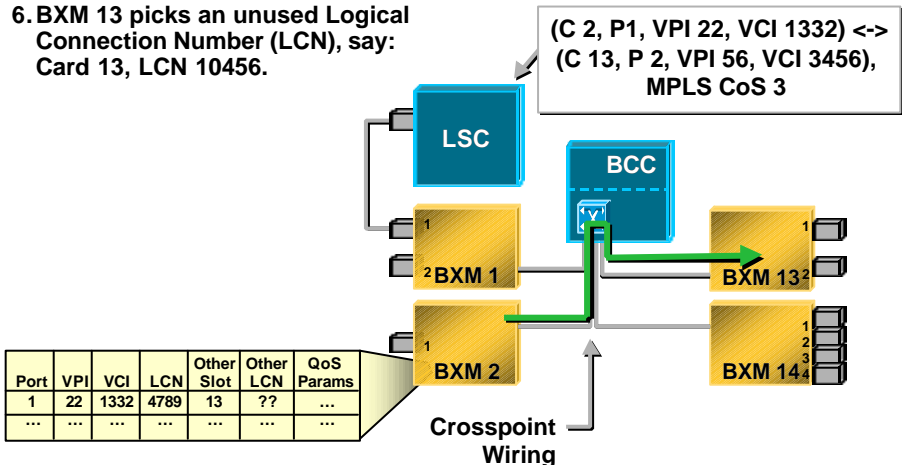
Setting Up an LVC Cross-Connect: 2

- BXM 2 picks an unused Logical Connection Number (LCN), say: Card2, LCN 4789.
- BXM 2 records the connection information, apart from Card 13's LCN.



Setting Up an LVC Cross-Connect: 3

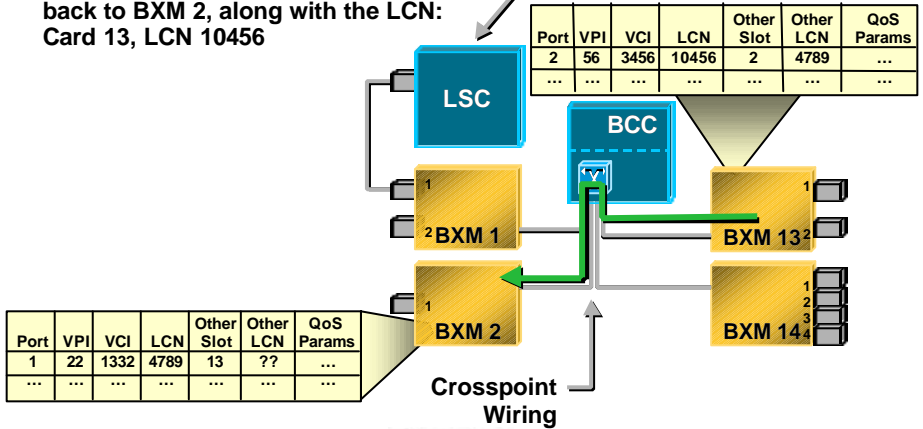
- BXM 2 sends the connection request on to BXM 13, along with the LCN: Card 2, LCN 4789.
- BXM 13 picks an unused Logical Connection Number (LCN), say: Card 13, LCN 10456.



Setting Up an LVC Cross-Connect: 4

7. BXM 13 programs the connection information in hardware.
8. BXM 13 sends a VSI confirmation back to BXM 2, along with the LCN: Card 13, LCN 10456

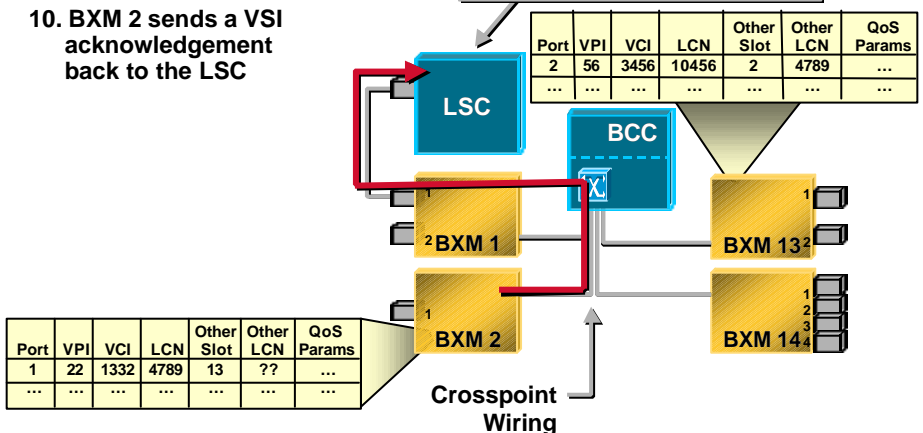
(C 2, P1, VPI 22, VCI 1332) <->
(C 13, P 2, VPI 56, VCI 3456),
MPLS CoS 3



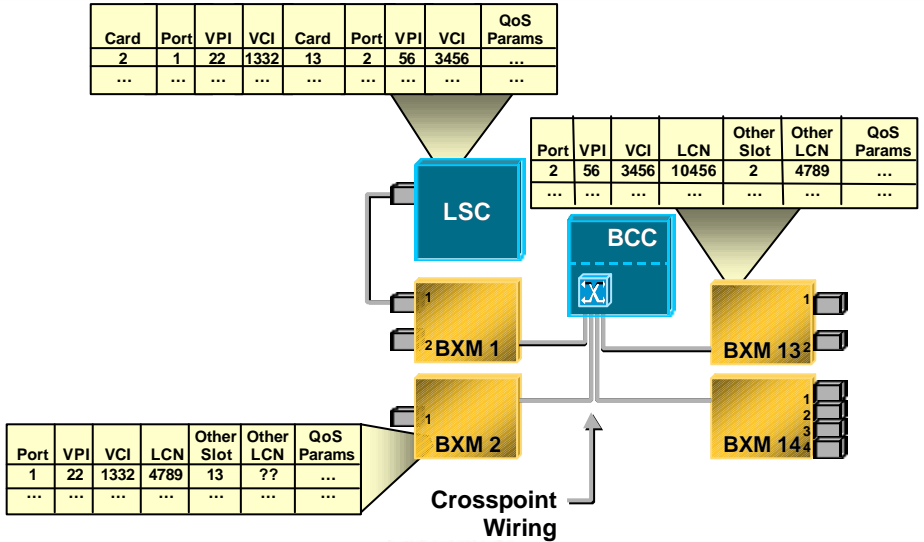
Setting Up an LVC Cross-Connect: 5

9. BXM 2 programs the full connection information in hardware.
10. BXM 2 sends a VSI acknowledgement back to the LSC

(C 2, P1, VPI 22, VCI 1332) <->
(C 13, P 2, VPI 56, VCI 3456),
MPLS CoS 3




Setting Up an LVC Cross-Connect: Final State



Setting Up an LVC Cross-Connect: Notes

- All hardware configuration of connection information occurs on the BXMs. The BCC hardware is not configured with any information about connections
- VSI is used by the LSC to program connections. It will also be used for PNNI
- AutoRoute programs the hardware similarly for PVC connections, but uses different software: the AutoRoute software runs on the BCC, and uses the BPX switch's communication bus to communicate with the BXMs. VSI is not used for PVCs

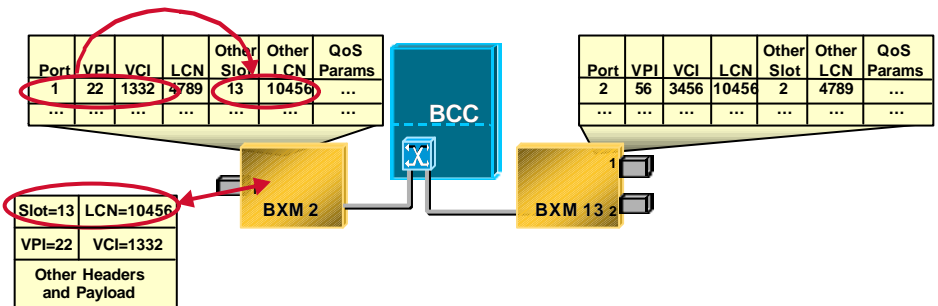
Forwarding a Cell on Connection: 1



0. Any software information in the LSC (or BCC) is not used in cell forwarding. The LSC plays no part in cell forwarding and is not shown.

1. Cell (53 bytes) arrives on Card 2, Port 1, VPI 22, VCI 1332


Forwarding a Cell on Connection: 2



2(a). Card 2 hardware looks up Port 1, VPI 22, VCI 1332 in its connection table, to find that the destination Slot is 13, and destination LCN is 10456.


2(b). Card 2 hardware encapsulates the cell in a 'Bframe header' with entries: Slot = 13, LCN = 10456

Forwarding a Cell on Connection: 3




- 3(a). Card 2 hardware forwards the cell, in the Bframe header, on its link to the switch fabric.
- 3(b). The switch fabric hardware reads that the destination Slot is 13 from the Bframe header, and forwards the cell out it's link to Card 13.

Forwarding a Cell on Connection: 4



- 4(a). Card 13 hardware reads the LCN=10456 from the Bframe header, and looks up the details in its connection table.
 - 4(b). Card 13 hardware finds that the cell is destined for Port 2, VPI 56, VCI 3456.
- (Note that at the moment the cell still has VPI=22, VCI=1332, the values it arrived with.)

Forwarding a Cell on Connection: 5



5(a). Card 13 hardware writes VPI=56, VCI=3456 into the cell header and discards the Bframe header.

5(b). Card 13 outputs the cell (after queueing and bandwidth control) on port 2.

Summary: Cell arrived on Card 2, Port 1, VPI 22, VC 1332.
Cell departs on Card 13, Port 2, VPI 56, VCI 3456.

Agenda

- Multiservice Requirements and Benefits
- Technical Motivation for MPLS with ATM
- **Applications and Transports**
- Summary

Applications and Transports

- IP and ATM integration
- IP-VPN
- MPLS CoS
- Traffic engineering

IP and ATM Integration

- A single IP+ATM network
- Versus**
- Dual-backbone network


IP+ATM Networks

- In an IP+ATM network, both edge and core switches support:

ATM and FR PVC and SVC services

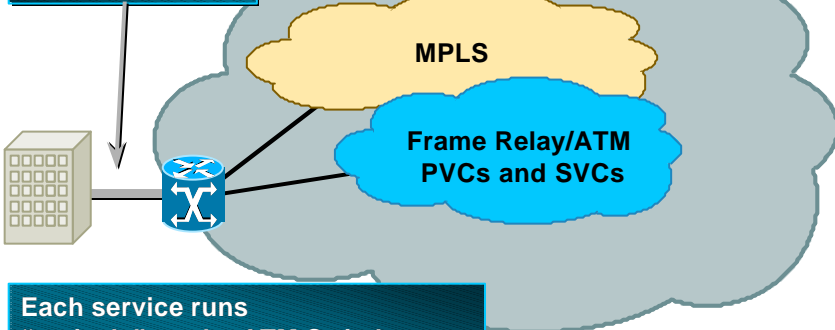
MPLS services and VPNs

- All services are fully integrated on one set of switches and one set of links



IP+ATM Networks

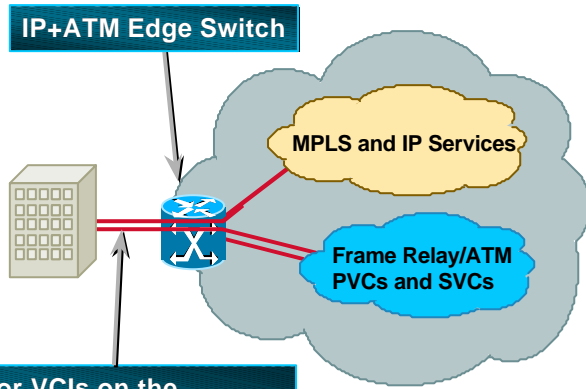
Unified Access to Multiple Services



Each service runs "natively" on the ATM Switches

- No overlays
- Resources partitioned by service

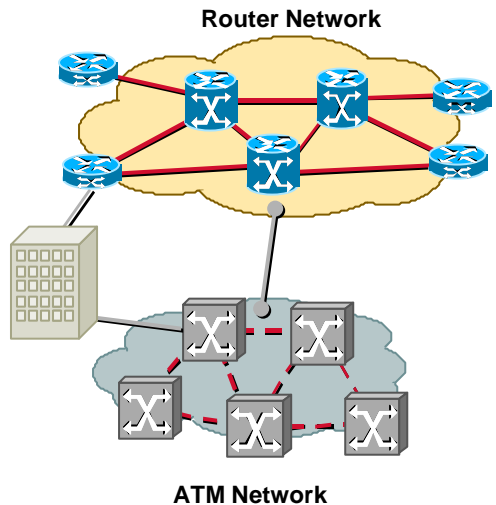
IP+ATM Edge Switches



Different DLCIs or VCIs on the same customer link can go to different services in the IP+ATM network.

Common Situation in Carriers: Separate Networks

- Carry IP services on router network and ATM/FR services on ATM network
- Management cost



IP-VPN


3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

57

Benefits of Internet-Scale VPNs



**VPN-Aware Network:
VPNs Are “Built-in”
Rather Than “Overlaid”**

3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com


58

MPLS VPN Goals


- Provide an IP VPN service, using a L3 backbone network
- Improve provisioning scalability for Provider
- Simplify routing configuration for customer
- Support IP CoS mechanisms

Layer 3 Architecture

- Two labels:
Inner VPN Label
Outer label implicit inside the ATM Cell header



MPLS VPN Routing Architecture




- **P** router = Provider Router (Core LSR)
- **PE** router = Provider Edge router (Edge LSR)
knows which VPN each CE belongs to (by sub-interface)
- **CE** router = Customer Edge router
- RD (Route Distinguisher) = uniquely identify a VPN (AS#,VPN_ID)
- IPv4 Addresses are unique within VPN
- IPv4 Addresses might overlap across VPNs

MPLS Class of Service

CoS on ATM-LSRs


- **Three modes:**
 - ATM Forum PVC**
 - Multi-‘VC’ (LSP)**
 - Single ‘VC’ (LSP)**

ATM Forum PVC Mode




- **Usually used in non-MPLS enabled ATM core**
- **Looks like a packet interface to MPLS CoS**
- **BW and other parameters configured on the PVC**
- **Requires significant amount of configuration**

Multi-VC Mode




- **MPLS ATM core provides MPLS CoS at each link by establishing parallel LVCs (up to 4)**
- **Classification** at MPLS ATM edge based on IP precedence
- **Scheduling policy** can assign weight to each class on per-link basis
- **Drop policy** can use switch pre-existing discard thresholds mechanisms on switches

Multi-VC Mode Example



- **Queuing on all links is per-class WFQ**
- **E.g., premium gets 80% of link, standard gets 20%**
Effect is to make premium see unloaded network
- **Unused bandwidth available to other classes**
- **No per-router-pair configuration required**

Multiple LVC Model




Operation

- Multiple labels per destination (one LVC for each class)
- Edge:
 - CAR classifies packets
 - WRED per interface
 - Class-based weighted fair queuing
- Core: class-based WFQ
- No end-to-end VCs required

■ Gold
 ■ Silver
 ■ Bronze
 Class-Based Queuing

Prefix	Class	Label
128.89	1	1217
128.89	2	1218
128.89	3	1219

Single ABR 'VC' Mode



- ABR control algorithms are enabled on LSRs
- ATM-LSRs push back on edge LSRs to eliminate congestion in ATM-LSR cloud
- Non-ATM-LSRs respond to ABR feedback, manage per-VC queues using WRED, WFQ
- Drop decisions made where MPLS CoS is visible—outside the ATM-LSR cloud


Single VC Mode (Cont.)

- **ABR parameters:**

**MCR is effectively zero
(to avoid loss/blocking)**

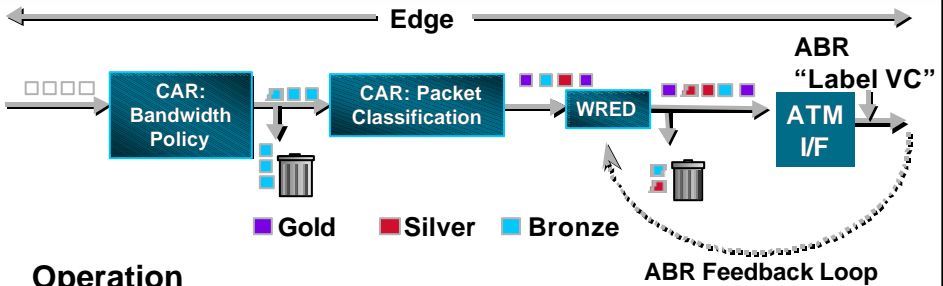
**Division of bandwidth
among competing VCs
may be uniform or configured**

Single VC Mode Example



- **Fair sharing of link A-B is not always desirable**
- **Configure relative bandwidth on router-pair basis, e.g. Tinytown-Nowheresville = 1; Chicago-New York = 100**
- **Relative BW carried in TDP, used by ABR algorithms**

Utilizing ATM ABR in an IP Service



Operation

- One label per destination (in SP network)
- One ABR LVC for all service classes
- Core: ABR "pushes" congestion to the edge
- Edge: WRED creates distinct classes
- ABR provides congestion notification directly to IP queues without an end-to-end PVC

Prefix	Label
128.89	10
129.45	354

Example of IP+ATM QoS Implementation: Cisco 8650 QBIN and SCT

COS in BPX Switches

- **Classification is at MPLS-ATM edge:**

Up to four LVCs created per dest. prefix, based on ToS

Scheduling policy on BXMs based on weights given to each ToS

Drop policy based on cell discard features of QBINs


Scheduling Policy: Cisco 8650 QBINs

- **16 CoS QBINs per interface**
- **AutoRoute created VCs are placed into QBIN based on service type**
CBR in QBIN 1, VBR in QBIN 2, etc.
- **MPLS uses QBINs 10-14**

IP Awareness in an IP+ATM Switch

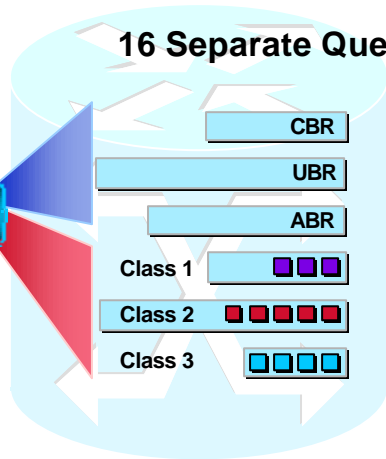
ATM Partition:

- Bandwidth allocated by VC



MPLS Partition:


- Bandwidth allocated by class
- Separate labels per class



ATM Queues

IP Queues

Class-Based WFQ: QoS Guarantees and Bandwidth Efficiency



- Weights guarantee minimum bandwidth
- Buffering controls latency
- Unused capacity is shared amongst the other classes
Proportional to their weights
- Each queue can be separately configured for QoS (e.g. ABR)
- Benefits:
 - Maximize transport of paying traffic
 - No loss of service class guarantees
 - No wasted bandwidth as with PVCs

IP Traffic Engineering

- IP Traffic Engineering is used to deal with issues in the operation of IP networks. These indirectly affect customer quality of service

IP+ATM Wins



Agenda

- Multiservice Requirements and Benefits
- Technical Motivation for MPLS with ATM
- Applications and Transports
- **Summary**



Cisco's IP+ATM switches help Service Provider to build a **single multiservice infrastructure optimized for **current services*** as well as new **IP-enabled services.****



* ATM, FR, private line

More Reading

- **IP+ATM technical info:**

<http://wwwin.cisco.com/mssbu/Tech/ipatm/index.html>

- **Sales Tools**

http://wwwin.cisco.com/mssbu/SalesTools/resource_list.html

<http://wwwin.cisco.com/mssbu/SalesTools/programs.html>

- **MSSBU MarketWatch**

<http://wwwin.cisco.com/mssbu/SalesTools/markets/watch.html>



Introduction to Real-Time ATM Services

Session 3001



Please Complete Your Evaluation Form

Session 3001

3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

83

CISCO SYSTEMS



EMPOWERING THE INTERNET GENERATIONSM

3001
1237_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

84