

White Paper

Migration Strategies for IP Service Growth: Cell-switched MPLS or IP-routed MPLS

Chuck Semeria
Marketing Engineer



Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
408 745 2000 or 888 JUNIPER
www.juniper.net

Part Number: 200026-001 03/02

Content

Executive Summary	4
Introduction	4
Initial Assumptions	5
Comparing IP-routed MPLS with IP over ATM	5
Comparing Cell-switched MPLS with IP over ATM	6
Advantages of Cell-switched MPLS Compared with IP over ATM	6
Reducing IGP Stress	6
Simplifying CoS Management	7
Supporting Primarily Cell-Based Services	7
Cell-switched MPLS Is Comparable to IP over ATM	7
Cell-switched MPLS Does Not Provide Better CoS than IP over ATM	7
Cell Tax	8
SAR Bottleneck	8
ATM Support for CBR Service Is Unchanged	8
Comparing Cell-switched MPLS with IP-routed MPLS	8
Two Different Types of Equipment to Manage	8
Cell Tax	8
SAR Bottlenecks	9
Slower Convergence Times	9
Limited Support for Fast Reroute	9
Exhaustion of a Finite Label Space	9
Cell-switched MPLS VC Merge Is Inferior to Label Merge	10
Stability, Scalability, and Interoperability	10
Transition Issues	10
IP over ATM to IP-routed MPLS	10
IP over ATM to Cell-switched MPLS	11
Two Transitions	11
Cell-switched MPLS Is a Transition in the Wrong Direction	11
Conclusion	11
Appendix I: Cell-switched MPLS Operational Model	13
Cell-switched MPLS Terminology	13
Unique Features of ATM LSRs	14
Label-Switching Control Component	14
The Cell-switched MPLS Label Space	16
VPI/VCI Field for Direct Connections	16
VCI Field for Indirect Connections	16
Encapsulation of Labeled Packets on LC ATM Links	17
Label Allocation and Distribution Mechanisms	18
MPLS Label-Request Modes	19
MPLS Label-Retention Modes	19
LDP Operational Modes: IP-routed MPLS vs. Cell-switched MPLS	20
VC Merge and Cell Interleave Issues	20
Label Merging in IP-routed MPLS Networks	20
Label Merging in Cell-switched MPLS Networks	21
Appendix II: References	23
Requests for Comments (RFCs)	23
Textbooks	23

List of Figures

Figure 1: Transition Options	5
Figure 2: Example of a Cell-switched MPLS Network Topology 1.....	4
Figure 3: Ships-in-the-Night Architecture	15
Figure 4: A Direct Connection between Two ATM LSRs	16
Figure 5: An Indirect Connection between Two ATM LSRs	17
Figure 6: MPLS Label-stack Entry	17
Figure 7: Encapsulation of Labeled Packets on LC ATM Links	18
Figure 8: Upstream and Downstream Nodes for a Label Binding	19
Figure 9: Downstream-on-demand vs. Unsolicited-downstream Label Distribution ...	19
Figure 10: How an LSR Performs Label Merging	21
Figure 11: How an ATM LSR Performs VC Merge	22
Figure 12: Assigning Multiple Labels to Eliminate Cell Interleave	22

List of Tables

Table 1: Comparison of IP over ATM, Cell-switched MPLS, and IP-routed MPLS	12
--	----

Executive Summary

If you currently deliver IP services through an IP-over-ATM infrastructure, you will most likely continue to do so until performance and scalability problems begin to impact the ability of your network to deliver these services. The successful delivery of services can be measured in terms of network complexity and resulting operational costs, as well as the performance that is required to deliver a satisfactory experience to your customers.

When the well-known limitations of the IP-over-ATM model start to impact the operation of your network, you will begin to examine new solutions that can overcome the limitations of your existing IP-over-ATM infrastructure. There are two potential transition strategies that you are most likely to consider:

- A transition to a cell-switched MPLS infrastructure
- A transition to an IP-routed MPLS infrastructure

Juniper Networks, Inc. believes that you should stay with your existing IP-over-ATM infrastructure until operational problems, cost structures, or strategic investment decisions force you to make a transition. However, when you make this transition, we recommend that you move directly to IP-routed MPLS, not to cell-switched MPLS. We offer this guidance because a transition from IP over ATM to cell-switched MPLS will not significantly enhance the operation of your network, yet it still requires you to make a second transition to IP-routed MPLS if you want to reach your ultimate goal. We believe that cell-switched MPLS is a slight improvement over IP over ATM, but cell-switched MPLS cannot deliver all of the benefits of IP-routed MPLS if your primary intent is to support the delivery of frame-based services.


Introduction

In this paper, we examine two probable transition strategies that you can use to overcome the well-known performance and scalability problems associated with IP over ATM (see Figure 1):

- IP-routed MPLS
- Cell-switched MPLS

A label-switching router (LSR) is a device that implements both the MPLS control and forwarding planes. IP-routed MPLS refers to the set of procedures that define how an IP router can function as an LSR to switch packets based on the top label carried in the packet's *shim header*. Cell-switched MPLS refers to the set of procedures that define how an ATM switch can function as an LSR and switch cells by inferring the top label, based on the contents of the cell's virtual path identifier (VPI) or VPI/VCI (VPI/virtual channel identifier) fields. In cell-switched MPLS, ATM switches execute network-layer routing protocols (such as OSPF or IS-IS), and all forwarding decisions are based on the results of these routing algorithms. ATM-specific routing protocols and addressing are not required.

For readers who may not be very familiar with cell-switched MPLS, Appendix 1 provides a brief tutorial.

Figure 1: Transition Options

We begin by evaluating the benefits and limitations of each approach compared with the others:

- Comparing IP-routed MPLS with IP over ATM
- Comparing cell-switched MPLS with IP over ATM
- Comparing cell-switched MPLS with IP-routed MPLS

We conclude with a discussion of the transition issues associated with each of these approaches and the role that they play in your decision process.

Initial Assumptions

In this paper we make a number of assumptions:

- We assume that IP-routed MPLS is your ultimate goal because the majority of the traffic that you plan to carry on your network will be IP- or frame-based, not cell-based.
- We assume that if making a transition does not enhance the performance of your network, reduce your operational expenses, or allow you to deploy new revenue-generating services, it is not worth the time, increased investment, and heightened risks of service disruption to make the transition.
- We assume that if you have the choice between making a single-step or a two-step transition to reach your ultimate goal, it is more desirable to make a single-step transition. This is certainly the case if a two-step transition is significantly more expensive than a single-step transition.
- Finally, we assume that when comparing the benefits, limitations, and capabilities of different approaches, you understand that there is a difference between what is *theoretically possible* to implement in a test network and what has been *actually implemented* and proven interoperable in large production networks.

Comparing IP-routed MPLS with IP over ATM

There are a number of well-known limitations associated with the traditional IP-over-ATM approach when supporting the delivery of IP- and frame-based services:

- IP over ATM has the potential to create bottlenecks leading into the core resulting from the lack of segmentation and reassembly (SAR) functionality on OC-48 and faster interfaces.
- IP over ATM increases network complexity by requiring providers to manage two separate control planes and two fundamentally different types of networking equipment.
- IP over ATM results in an inefficient use of network bandwidth due to the traditional ATM cell tax.

- The IP Differentiated Services (DiffServ) approach to class of service (CoS) does not map well to existing ATM quality of service (QoS) mechanisms.
- IP over ATM requires the deployment of n-squared routing adjacencies that can lead to Interior Gateway Protocol (IGP) stress during failure scenarios. While the stress may be ameliorated using techniques such as mesh groups, the use of mesh groups adds complexity at the routing layer to overcome the characteristics of the ATM layer.

In the world's largest IP networks, the deployment of IP-routed MPLS has demonstrated its ability to overcome these well-known limitations of IP over ATM:

- IP-routed MPLS eliminates the potential for SAR bottlenecks, because it does not rely on ATM as a transport.
- IP-routed MPLS eliminates the complexity of managing two separate control planes and two fundamentally different types of networking equipment, because it requires only a single IP/MPLS control plane and runs over frame-based switches.
- IP-routed MPLS provides more efficient use of network bandwidth, because it eliminates the cell tax by not using ATM as a transport.
- IP-routed MPLS can support IP DiffServ CoS, because frame-based LSRs have the ability to read and write the EXP bits in the MPLS label stack that carry the IP DiffServ information. Also, frame-based LSRs can natively apply DiffServ-defined PHBs, based on the value of these EXP bit values. These capabilities support the delivery of multiple service levels within a single LSP.
- The peer-to-peer nature of IP-routed MPLS eliminates IGP stress during failure scenarios, because it does not require the establishment of n-squared adjacencies between all of the edge routers.
- The peer-to-peer nature of IP-routed MPLS eliminates the need to manage a complex logical topology (n-squared PVCs) within a logical circuit-based underlying layer.

Comparing Cell-switched MPLS with IP over ATM

There are certain aspects in which a cell-switched MPLS infrastructure has some advantages when compared with IP over ATM. However, there are a number of issues where a transition to cell-switched MPLS is no better than staying with your existing IP-over-ATM network.

Advantages of Cell-switched MPLS Compared with IP over ATM

There are three cases where a transition from IP over ATM to cell-switched MPLS can remove complexity and simplify network operations.

Reducing IGP Stress

One of the major problems when trying to scale IP over ATM is that the deployment of n-squared permanent virtual connections (PVCs) leads to IGP stress during failure situations. The peer-to-peer nature of cell-switched MPLS eliminates the need to establish an overlay network consisting of a full mesh of n-squared routing adjacencies. By eliminating the need to establish a full mesh of n-squared routing adjacencies, the IGP stress that occurs during failure scenarios in IP over ATM can be eliminated.

Simplifying CoS Management

Cell-switched MPLS can simplify CoS management in your network, because the dual IP/MPLS and ATM control planes in an IP-over-ATM network can be consolidated into a single IP/MPLS control plane. Cell-switched MPLS supports IP CoS by establishing multiple LSPs, with each LSP dedicated to a specific traffic class. The essential difference between IP over ATM and cell-switched MPLS is that, instead of using ATM signaling to establish the PVCs, MPLS signaling establishes the LSPs. So, while it may be simpler to manage CoS in cell-switched MPLS because the task is accomplished using a single control plane, the CoS delivered by cell-switched MPLS is similar, if not identical, to the CoS delivered by IP over ATM.

However, while a single control plane has benefits when compared with the challenges of coordinating two independent control planes, the deployment of cell-switched MPLS does not eliminate the management complexity associated with the deployment of two fundamentally different types of equipment in the core of your network—IP routers and ATM switches.

Finally, if you plan to run ATM to support cell-based services, plus cell-switched MPLS to support frame-based services over a common infrastructure, you will still have two control planes to manage—an ATM control plane and an IP/MPLS control plane. While these two control planes can operate as ships in the night, you still need to coordinate the partitioning of the ATM VPI/VCI space between the two control planes and the amount of link bandwidth allocated to each of the control planes. Thus, cell-switched MPLS can simplify CoS management in your network only if you plan to support frame-based services but not cell-based services. However, if you are planning to offer both cell-based services and frame-based services on the same infrastructure, then CoS management becomes even more complex.

Supporting Primarily Cell-Based Services

Cell-switched MPLS offers an enhancement to IP over ATM for providers that

- want to keep their existing ATM infrastructures,
- do not have the budget to purchase new frame-based LSRs,
- believe in the future of cell-based networks, or
- transmit mostly cell-based traffic, not frame-based.

For these cases, cell-switched MPLS has the potential to simplify CoS management and to improve routing scalability by eliminating the n-squared PVC problem associated with IP over ATM.

Cell-switched MPLS Is Comparable to IP over ATM

While there are cases where cell-switched MPLS is an improvement over your existing IP-over-ATM network, there are a number of cases where cell-switched MPLS is simply comparable to IP over ATM, thus offers no significant improvement.

Cell-switched MPLS Does Not Provide Better CoS than IP over ATM

As discussed above, cell-switched MPLS does not deliver better CoS for frame-based services than IP over ATM. In IP over ATM, CoS is supported by using ATM signaling to establish multiple PVCs, with each PVC carrying a specific class of traffic. In a cell-switched MPLS network, CoS is supported by using MPLS signaling to establish multiple LSPs, with each LSP

carrying a specific class of traffic. If your goal is to support CoS for IP packets, a transition from IP over ATM to cell-switched MPLS does not provide better CoS for your frame-based IP traffic.

Cell Tax

Neither IP over ATM nor cell-switched MPLS make the most efficient use of network bandwidth, due to the ATM cell tax. Assuming an Internet traffic mix, applying ATM Adaptation Layer type 5 (AAL5) encapsulation, and including nulls in the last unfilled cell, this means that both IP over ATM and cell-switched MPLS can consume up to 25 percent more bandwidth than a frame-based solution, because both of these approaches run over a cell-based transport. Undertaking a transition from IP over ATM to cell-switched MPLS does not eliminate the ATM cell tax; therefore, moving to cell-switched MPLS does not improve bandwidth use on your network.

SAR Bottleneck

IP routers at the edges of both IP-over-ATM and cell-switched MPLS networks require SAR interfaces to communicate with ATM switches in the core of the network. While OC-48 SARs may exist with very limited feature sets, *theory vs. practice* limits the performance of full-feature SAR interfaces to OC-12 and below. The inability of SARs to support OC-48 and faster interfaces places a hard limit on the maximum size of feeder links from the routing layer to the cell core. A transition from IP over ATM to cell-switched MPLS does not eliminate the potential for the creation of SAR bottlenecks.

ATM Support for CBR Service Is Unchanged

ATM provides a constant bit rate (CBR) service to support circuit emulation. The same ATM infrastructure that supports the transport of IP packets using IP over ATM or cell-switched MPLS can also support CBR cell-based services. In this sense, a transition from IP over ATM to cell-switched MPLS has no impact on the ability of the network to continue supporting CBR services.

Comparing Cell-switched MPLS with IP-routed MPLS

While cell-switched MPLS offers advantages when compared with IP over ATM, cell-switched MPLS cannot deliver all of the benefits of IP-routed MPLS if your primary objective is to support the delivery of frame-based services.

Two Different Types of Equipment to Manage

Cell-switched MPLS requires you to manage two fundamentally different types of equipment in your network—IP routers and ATM switches. In contrast, IP-routed MPLS requires you to manage only one type of equipment—frame-based LSRs. A frame-based LSR is an IP router configured to perform a specialized function.

Cell Tax

Cell-switched MPLS does not support the most efficient use of network bandwidth, due to the ATM cell tax. This means that cell-switched MPLS can consume up to 25 percent more bandwidth than IP-routed MPLS. In fact, TCP/IP is one of the worst protocols to run over an ATM transport, because the transmission of a single ACK requires not one but two ATM cells.

SAR Bottlenecks

Cell-switched MPLS requires SAR interfaces to communicate with cell-based switches in the core of the network. In contrast, IP-routed MPLS does not require SAR interfaces, because it does not run over cell switching. As a result, IP-routed MPLS eliminates SAR bottlenecks and allows edge systems to access the core at OC-48 and greater speeds.

Slower Convergence Times

ATM LSRs operate in *downstream-on-demand* and *conservative label-retention* modes. As a result, ATM LSRs maintain only a single label mapping for each prefix and are required to request a new label from a downstream LSR if the path fails. In contrast, frame-based LSRs operate in *unsolicited-downstream* and *liberal label-retention* modes. Therefore, frame-based LSRs can maintain multiple label mappings for each prefix and have a backup path that is ready for immediate use if the next hop to a destination changes. Also, for ATM LSRs that do not support virtual circuit (VC) merge capability, when the next hop changes, the LSR may need to request not one but multiple labels from the new next hop. The difference in the operational modes of LDP in cell-switched MPLS and IP-routed MPLS allows IP-routed MPLS to respond more rapidly to path failures.

Limited Support for Fast Reroute

MPLS fast reroute can significantly reduce packet losses during periods of network instability. However, the fast-reroute mechanisms supported by IP-routed MPLS are generally not supported by cell-switched MPLS.

- The limited support for label-stacking operations in cell-switched MPLS using the VCI space makes it very difficult to support fast-reroute mechanisms that are based on the use of label stacking. While support for a two-level hierarchy using VPI and VCI label stacking is possible, this is another theory vs. reality issue.
- Other fast-reroute mechanisms require end-to-end signaling from the ingress point to the egress point of an LSP. This requires an end-to-end signaling protocol that is similar to the model supported by RSVP-TE or CR-LDP. While RSVP-TE or CR-LDP could conceivably be deployed by cell-switched MPLS, this is yet another theory vs. reality issue.

Exhaustion of a Finite Label Space

Both the label space used by IP-routed MPLS and the VPI/VCI label space used by cell-switched MPLS are finite resources. If the finite label space becomes exhausted, cell-switched MPLS degrades less gracefully for certain applications (such as traffic engineering) than for IP-routed MPLS. The failure mode for cell-switched MPLS is an inability to establish new LSPs. Given the same circumstances, IP-routed MPLS can continue to provide connectivity using the underlying IGP route to forward traffic to its destination. However, in the case of an application such as VPNs, both cell-switched MPLS and IP-routed MPLS have the same failure mode, because they are not able to provide connectivity by establishing new LSPs.

The use of arbitrary label stacking allows you to create a hierarchy to conserve a finite label space and eliminate these potential problems. While it is very easy for IP-routed MPLS to support arbitrary label stacks, this is another of the theory vs. reality issues for cell-switched MPLS.

Cell-switched MPLS VC Merge Is Inferior to Label Merge

When cell-switched MPLS performs VC merge in a multipoint-to-point PVC, cells from two different frames arriving from two different tributaries in the merge cannot be interleaved on the output interface. Because cells cannot be interleaved, all of the cells from the first frame that presents a cell to the interface past the merge point must be transmitted in sequence. This means that even if a cell from another frame arrives at the merge point and the cell is ready to be transmitted, this cell must be buffered and cannot be sent until all cells comprising the first frame are transmitted. This creates a situation in an ATM LSR that is identical to a small packet waiting behind a larger packet in a frame-based LSR. However, the situation is even worse in cell-switched MPLS, because the cell overhead requires the consumption of more bandwidth.

If ATM LSRs do not support VC merge, then providers have to be careful about their label use. The lack of VC merge capabilities requires an increase in the number of LSPs that need to be established to support connectivity requirements, yet also consumes significantly more labels than when VC merge is supported. The limitations of cell-switched MPLS to support label-stacking operations means that IP-routed MPLS can provide better scalability with respect to the forwarding (label) space.

Stability, Scalability, and Interoperability

It is important to determine the stability, scalability, and interoperability of each vendor's implementation of cell-switched MPLS compared with IP-routed MPLS. While there are many things that can be done in theory, a critical issue when choosing between alternative transition strategies is to determine what protocols (and features of those protocols) have actually been implemented and are currently available in products being shipped by vendors. It is also important to consider the deployment history of each implementation's signaling (RSVP-TE and CR-LDP) and routing (OSPF and IS-IS) protocols. You need to ask several questions:

- Has the implementation been deployed in large production networks?
- Has the implementation demonstrated its stability in these deployments?
- Has the implementation proven its scalability in rapidly growing production environments?
- Has the implementation verified its interoperability in third-party multivendor testing labs?
- Will the implementation deliver the feature velocity required to support mission-critical applications in hostile environments?

Transition Issues

After comparing and contrasting each of these approaches, we conclude with a discussion of the benefits and limitations associated with making a direct transition to IP-routed MPLS or an indirect transition to IP-routed MPLS via cell-switched MPLS.

IP over ATM to IP-routed MPLS

In terms of capital expenses, a direct transition from IP over ATM to IP-routed MPLS is more expensive than a transition from IP over ATM to cell-switched MPLS, because you are required to purchase new frame-based LSRs. However, after completing this transition, your network will immediately begin to receive the performance and scalability benefits of IP-routed MPLS,

and you are not faced with a second transition. Also, if you are still required to support ATM cell traffic over your new IP-routed MPLS infrastructure, then features such as Layer 2 VPNs and cell relay may provide useful tools to achieve this goal. Finally, the service provider community has a significant amount of experience making the transition from IP over ATM to IP-routed MPLS in large production networks.

IP over ATM to Cell-switched MPLS

There are two issues related to making a transition from IP over ATM to cell-switched MPLS:

- A second transition is required to achieve your goal of IP-routed MPLS.
- A transition to cell-switched MPLS is a transition in the wrong direction.

Two Transitions

An initial transition from IP over ATM to cell-switched MPLS is less expensive than a direct transition to IP-routed MPLS, because you can cross over by completing a software upgrade and adding a routing processor module to your ATM switches. The challenge with this approach is that, at some point, you will eventually need to make a second transition from cell-switched MPLS to IP-routed MPLS when the performance and scalability limitations of an ATM infrastructure start to impact the operation of your network. The combined expense of undertaking a two-step transition strategy will be significantly greater than if you make a direct, single-step transition from IP over ATM to IP-routed MPLS.

Cell-switched MPLS Is a Transition in the Wrong Direction

The natural way for an IP router to operate is to maintain a primary route for each prefix, as well as one or more backup routes that are ready for immediate use if the primary route fails. When deploying IP-routed MPLS and using LDP unsolicited downstream mode with liberal label-retention modes, the same model applies. However, when deploying cell-switched MPLS and using LDP downstream-on-demand mode with conservative label-retention modes, when the next hop for a specific destination changes, an ATM LSR cannot forward traffic to this destination. The ATM LSR is required to request and receive label-binding information for that destination from the new next hop before it can forward traffic across the backup route.

Operating some of your LSRs in one mode for IP-routed MPLS links and some in another mode for cell-switched MPLS links adds complexity to the second transition from cell-switched MPLS to IP-routed MPLS. Finally, when compared with the operational experience of making an IP over ATM to IP-routed MPLS transition, the service provider community has very little, if any, experience making a transition from cell-switched MPLS to IP-routed MPLS in large production environments.

Conclusion

At the beginning of this paper, we assumed that the bulk of the traffic traversing your network supports the delivery of IP-based services to your customers. Table 1 shows how successful cell-switched MPLS and IP-routed MPLS are at surmounting the well-known limitations of IP over ATM.

Table 1: Comparison of IP over ATM, Cell-switched MPLS, and IP-routed MPLS

	IP over ATM	Cell-switched MPLS	IP-routed MPLS
Elimination of SAR bottleneck	No	No	Yes
A single control plane to manage	No	Yes ^a	Yes
A single type of equipment to manage	No	No	Yes
Elimination of ATM cell tax	No	No	Yes
Native support for IP DiffServ CoS	No	No	Yes
Elimination of IGP stress	No	Yes	Yes

a. Only if frame-based services, not native ATM services, are supported.

From an examination of this table, you can see that there are a couple of areas where cell-switched MPLS provides some advantages over IP over ATM, but there are also a number of cases where cell-switched MPLS offers no significant improvement to IP-over-ATM networks. In contrast, IP-routed MPLS is superior to IP over ATM in all instances, and there are a number of important cases when cell-switched MPLS is inferior to IP-routed MPLS. If your ultimate goal is to support the delivery of IP-based services, then IP-routed MPLS is clearly a superior solution to cell-switched MPLS.

Transitions involving a changeover from IP over ATM to IP-routed MPLS have taken place in a number of the world's largest provider networks. In contrast, changeovers from IP over ATM to cell-switched MPLS have yet to take place in large production networks. Also, if you assume that the bulk of the traffic in your network will be frame-based, you will eventually have to make a second transition to IP-routed MPLS when the limitations of an ATM transport begin to impact the performance and scalability of your network. While your initial transition from IP over ATM to cell-switched MPLS may be less expensive than a direct transition to IP-routed MPLS, the combined expense of undertaking a two-step transition strategy will be significantly greater than taking a single-step transition directly to IP-routed MPLS.

Juniper Networks believes that IP-routed MPLS is the ideal solution for providers that believe the future will focus on frame-based applications. If you do not believe in the future of frames, then our guidance is to stay with your existing ATM infrastructure and continue to use IP over ATM to support your frame-based services. However, if you believe in the future of frames and need to preserve your investment in ATM, our advice is to stay with your existing IP over ATM infrastructure, cap your investment in ATM, and start investing in IP-routed MPLS. A transition to cell-switched MPLS offers only a slight improvement to IP over ATM, and that approach falls short of delivering the performance and scalability of IP-routed MPLS.

Appendix I: Cell-switched MPLS Operational Model

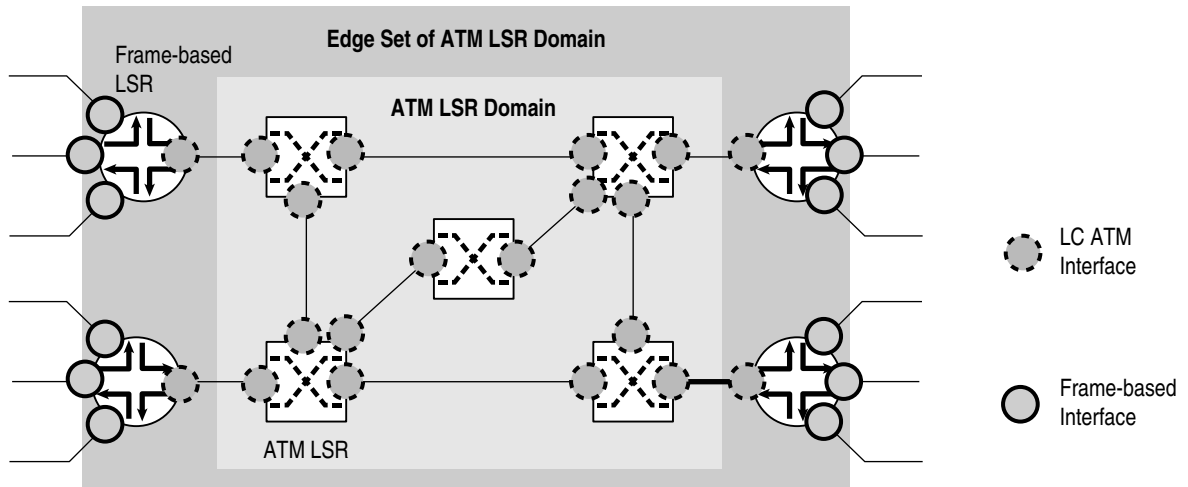
Multiprotocol Label Switching (MPLS) enhances the capabilities of IP networks by providing a mechanism that supports traffic engineering, virtual private networks (VPNs), and the integration of IP routing with Layer 2 technologies or optical switching. The initial deployment of MPLS has occurred in packet-based networks where labels are carried in a *shim header* that is inserted between each packet's link-layer and network-layer headers. In a packet-based network, the label-switching routers (LSRs) exchange native IP packets to support control functions and forward labeled IP packets to transport user data.

RFCs 3031 and 3035 define procedures that allow an Asynchronous Transfer Mode (ATM) switch to function as an LSR. When operating as an LSR, an ATM switch uses IP addressing, runs network-layer routing protocols to determine routes, executes MPLS-specific label distribution and signalling protocols, and forwards cells, based on the results of these algorithms. Conventional ATM addressing, routing, and signaling protocols are not used for determining how to forward cell-switched MPLS traffic.

Cell-switched MPLS Terminology

This section defines a number of terms that are specific to the deployment of a cell-switched MPLS network where ATM switches function as LSRs (see Figure 2):

- Label-switching-controlled ATM (LC ATM) interface—A switch or router ATM interface that is controlled by the device's label-switching control component. When a cell is received on an LC ATM interface, the cell is treated as a labeled packet, and the top label is implied from the VCI or VPI/VCI fields in the cell header.
- ATM label-switching router (ATM LSR)—An LSR with several LC ATM interfaces. ATM LSRs execute traditional network-layer routing algorithms (such as OSPF or IS-IS), and cell forwarding is based on the results of these algorithms. An ATM LSR forwards cells between its LC ATM interfaces using the label carried in the cell's VCI or VPI/VCI fields, without reassembling the cells into frames before forwarding the cells across the switching fabric.
- Frame-based LSR—An LSR that forwards complete frames between its interfaces. An example of a frame-based LSR is a traditional IP router that forwards frames between pairs of Ethernet interfaces or frames between an Ethernet interface and an LC ATM interface. Frame-based LSRs may have zero, one, or more LC ATM interfaces.
- ATM LSR domain—A set of ATM LSRs that are interconnected by LC ATM interfaces.
- The edge set of an ATM LSR domain—The set of frame-based LSRs that are connected to the ATM LSR domain by LC ATM interfaces.

Figure 2: Example of a Cell-switched MPLS Network Topology

Unique Features of ATM LSRs

When adapting conventional MPLS mechanisms to run over a cell-based technology, there are a number of procedures that need to be developed to overcome the constraints imposed by the pre-existing ATM cell format and possibly pre-existing limitations of ATM switching hardware:

- ATM LSRs must exchange control-plane IP packets (for label distribution and routing updates) over an ATM-control virtual circuit (VC), because ATM switches cannot exchange native IP packets directly across an ATM interface.
- The use of the VCI or VPI/VCI fields in the ATM cell header to carry the label dictates the size and placement of labels in the cell header.
- The pushing or popping of labels within an ATM LSR domain is not supported, because the individual cells of an ATM Adaptation Layer type 5 (AAL5) PDU must be reassembled in order to access the label stack. These operations can be performed only by frame-based LSRs.
- New mechanisms must be designed to support the detection and prevention of loops, because ATM cell headers do not carry a time-to-live (TTL) field.
- The network topology must be carefully designed so that ATM LSRs are not asked to perform Layer 3 lookups or route summarization, because, as Layer 2 devices, they are incapable of supporting these functions.

This appendix describes how these constraints have been overcome and explains the various mechanisms that are defined to allow an ATM switch to function as an ATM LSR.

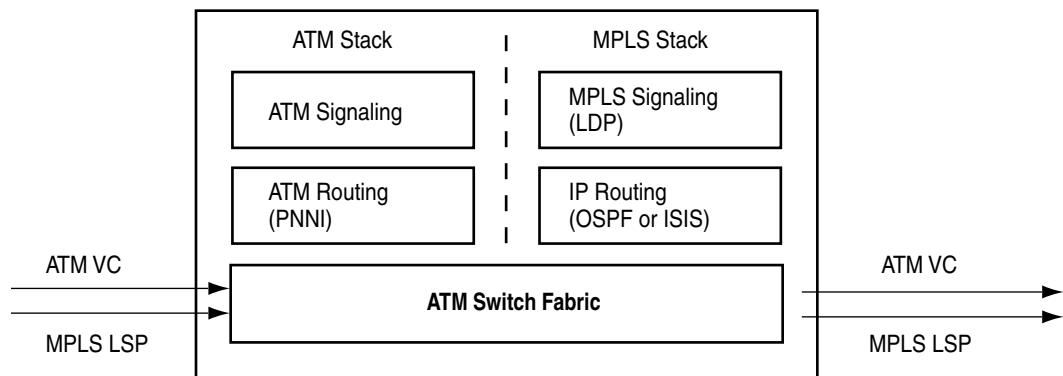
Label-Switching Control Component

To support cell-switched MPLS, an ATM LSR is required to implement a label-switching control component. The cell-switched MPLS label-switching control component consists of a number of mechanisms:

- A label allocation, distribution, and maintenance mechanism.—The Label Distribution Protocol (LDP) is required to support this function, although other protocols (such as RSVP-TE, CR-LDP, PIM, and BGP) may also be used to distribute label bindings. If these other protocols are deployed in the network, *each* ATM LSR is required to participate in these protocols to support label distribution.
- A Layer 3 routing protocol, such as OSPF or IS-IS.—An ATM LSR participates as a peer in these IETF-defined routing protocols to provide network topology information to the cell-switched MPLS control component. This means that all cell-forwarding decisions are based on information learned from the IETF protocols, not the protocols defined by the ITU or the ATM Forum (such as UNI, NNI, or PNNI). An ATM LSR is *not* required to support the traditional ATM control component.

The existence of a cell-switched MPLS control component on an ATM LSR does not prevent the coexistence of the traditional ATM control component on the same switch and the same interfaces. In fact, it is likely that the initial deployment of cell-switched MPLS will occur on existing hardware that supports dual control planes. Each of the two control components would be co-resident on the ATM switch, but would operate independently as *ships in the night*. (see Figure 3.)

Figure 3: Ships-in-the-Night Architecture



When deploying the ships-in-the-night model of operation, the physical ATM network is divided into two disjointed topologies—one topology maintained by the traditional ATM control component and the other topology supported by the cell-switched MPLS control component. Among the major issues that ATM LSR switch vendors and network operators need to resolve are how to allocate network resources between these two co-resident but mutually exclusive control components. For example:

- How should the native ATM VPI/VCI label space be distributed between the ATM and the cell-switched MPLS control components?
- How should the bandwidth on network interfaces be allocated between the ATM and cell-switched MPLS control components?
- How are CoS and QoS supported by each of the control components without negatively impacting the ability of the other control component to deliver these services for its applications?

On some ATM switches, support for these two distinct control planes can be provided by the existing control processor. Other ATM switches will require the insertion of an additional control processor—one control processor for the ATM control plane and a second control processor for the cell-switched MPLS control plane. Still other ATM switches will require the deployment of a dedicated external controller to support the cell-switched MPLS control plane.

The Cell-switched MPLS Label Space

Assuming the use of AAL5 encapsulation by the edge set of an ATM LSR domain, there are two mechanisms that allow a label to be encoded in an ATM cell header:


- The VPI/VCI field
- The VCI field

VPI/VCI Field for Direct Connections

When two ATM LSRs are directly connected over an LC ATM interface (see Figure 4), they may agree to use the combined VPI/VCI field to encode the top label in the MPLS label stack. When using this encoding technique, an ATM LSR cannot perform *push* or *pop* operations on the label stack (see the discussion that follows describing the encapsulation of labeled packets).

The two directly connected LDP peers initially establish a non-MPLS connection that supports the transmission of unlabeled packets between each ATM LSR. This non-MPLS connection is used to carry LDP packets between the two peers and may also be used for other unlabeled packets, such as routing protocol updates. The default VPI/VCI value for the non-MPLS control connection between directly connected ATM LSRs is 0/32. (see Figure 4.)

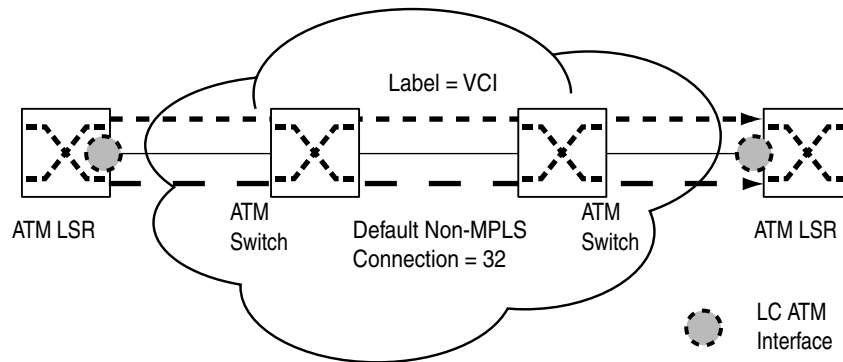
Figure 4: A Direct Connection between Two ATM LSRs



VCI Field for Indirect Connections

Sometimes it is desirable to consider two ATM LSRs as adjacent in a given LSP across an LC ATM interface, even if the connection between the two ATM LSRs is made through an ATM cloud via an ATM VP (see Figure 5). For this case, the bits in the VPI field are not available to MPLS, so the top label in the label stack must be encoded exclusively within the VCI field. The default VCI value for the non-MPLS control connection between two indirectly connected ATM LSRs is 32.

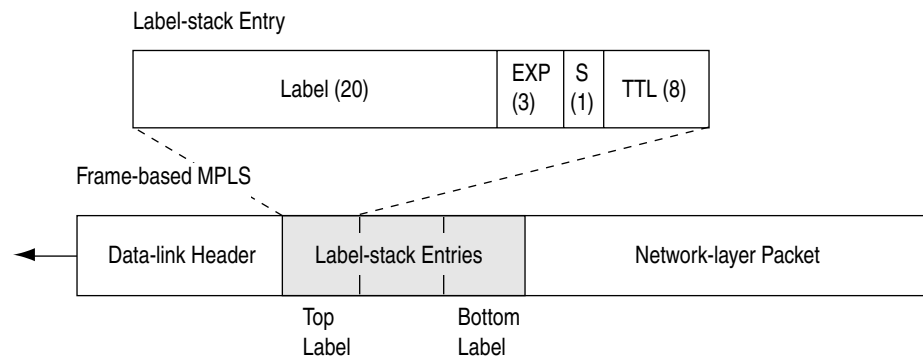
Figure 5: An Indirect Connection between Two ATM LSRs



Encapsulation of Labeled Packets on LC ATM Links

IP-routed MPLS supports a label stack of an arbitrary depth. An IP-routed MPLS label stack is composed of a sequence of label-stack entries. Each label-stack entry contains 4 octets that include the label value (20 bits), the EXP field (3 bits), the bottom of stack marker (1 bit), and the TTL sub-fields (8 bits). Label-stack entries are encoded after each packet’s data-link header, but before the network-layer header and payload (see Figure 6).

Figure 6: MPLS Label-stack Entry



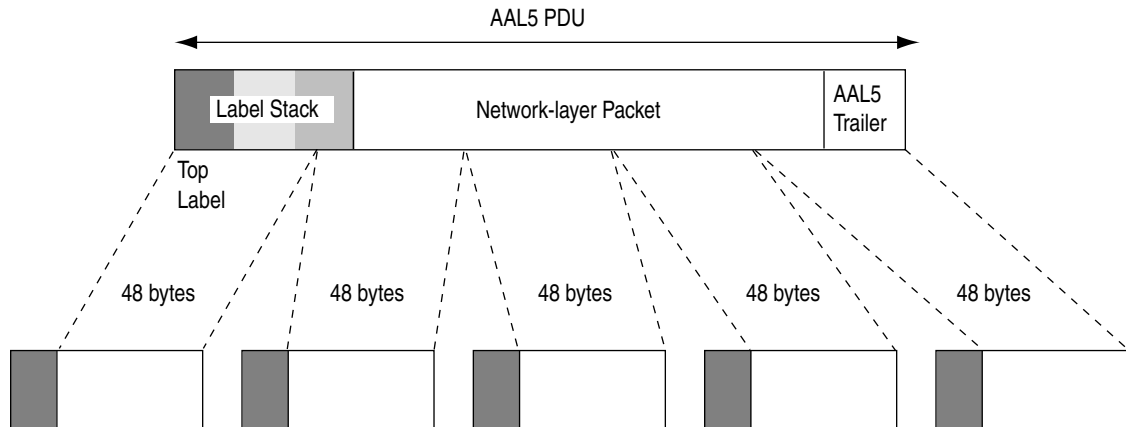
The encapsulation functions described in this section apply only to the edge set of an ATM LSR domain. In an ATM LSR domain, the ingress frame-based edge set LSR encapsulates labeled or unlabeled packets by using AAL5, then segments the AAL5 PDU into 53-byte ATM cells, and then forwards the cells into the core of the ATM LSR domain. The core ATM LSRs:

- Receive each cell on an incoming interface with a known VPI/VCI value (input label),
- Look up this value in their cross-connect tables to determine the outbound interface and corresponding VPI/VCI value (outbound label),
- Rewrite the VPI/VCI value (outbound label) into the outbound cell header, and
- Forward the cell to the appropriate outbound interface for transmission to the next-hop ATM LSR.

When all of the cells that constitute the original packet arrive at the egress frame-based edge set LSR, they are reassembled and then forwarded to their ultimate destinations as an unlabeled IP packet or as a labeled MPLS frame.

The ingress frame-based edge set LSR encapsulates the network-layer packet as an AAL5 PDU before the packet is segmented into cells (see Figure 7). To support label stacks of arbitrary depth on LC ATM links, the AAL5 PDU is encoded by using the null encapsulation defined in Section 6.1 of RFC 2684. If a packet has a label stack that contains four entries, then the AAL5 PDU contains a label stack that consists of four label-stack entries. The value of the top label in the label stack is encoded in the ATM VPI/VCI field of each cell that is segmented from a given AAL5 PDU.


Figure 7: Encapsulation of Labeled Packets on LC ATM Links



This encoding process does not permit core ATM LSRs to execute push or pop functions on the label stack. The reason for this is that these operations require not only modifications to the VPI/VCI field in the ATM cell header, but also modifications to the label stack. Because the label stack is not present in each of the cells segmented from a given AAL5 PDU, a push or pop operation requires that the entire packet be reassembled before the encoded label stack can be accessed and then modified. Packet reassembly can be performed only by frame-based LSRs, (not by cell-based ATM LSRs), so push and pop functions cannot be performed by cell-based LSRs.

Label Allocation and Distribution Mechanisms

When two LSRs agree to exchange a label binding, one LSR is referred to as the *upstream node* and the other LSR is referred to as the *downstream node* with respect to the label binding. In this relationship, it is assumed that labeled data packets or cells will flow from the upstream node to the downstream node (see Figure 8).

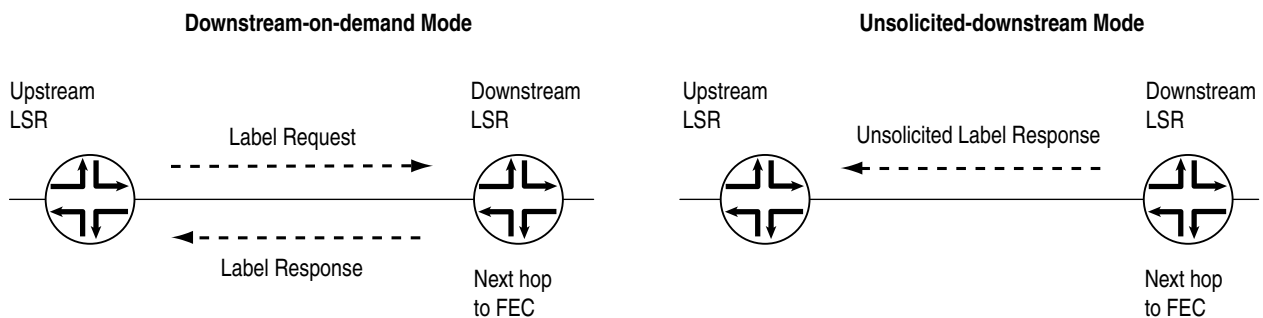
Figure 8: Upstream and Downstream Nodes for a Label Binding

MPLS Label-Request Modes

The distribution of label bindings is typically initiated using one of two mechanisms:

- Downstream-on-demand label distribution
- Unsolicited-downstream label distribution

Downstream-on-demand label distribution occurs when an upstream LSR explicitly requests that the next-hop LSR toward a FEC provide a label binding for that FEC. Unsolicited downstream label distribution occurs when a downstream LSR advertises a label binding for an FEC to an upstream neighbor without the upstream neighbor specifically requesting the label binding (see Figure 9).

Figure 9: Downstream-on-demand vs. Unsolicited-downstream Label Distribution

MPLS Label-Retention Modes

When an upstream LSR receives a label binding for a given FEC from a downstream LSR, it may choose to retain or to discard the label binding, based on its support for one of two mechanisms:

- Conservative label retention
- Liberal label retention

If conservative label-retention mode is supported, an upstream LSR maintains the received label binding for an FEC only if the label binding is received from the downstream LSR that the upstream LSR has selected as the next hop for that FEC. If liberal label retention-mode is supported, an upstream LSR maintains the received label binding for an FEC even if the label binding is received from a downstream LSR that the upstream LSR has not selected as the next hop for that FEC.

Conservative label-retention mode has the benefit of requiring an LSR to maintain fewer label bindings, but this provides slower adaptation to routing changes, because label bindings for alternative paths are not established and must be defined before an alternative path can be

used. On the other hand, liberal label-retention mode has the disadvantage of requiring an LSR to maintain more label bindings, but this allows quicker adaptation to routing changes, because label bindings for alternative paths are already established.

LDP Operational Modes: IP-routed MPLS vs. Cell-switched MPLS

The LDP functions by using different operational modes when deployed in an IP-routed MPLS network than when deployed in a cell-switched MPLS network.

- In an IP-routed MPLS network, LDP executes in unsolicited downstream with liberal label-retention mode. This requires LSRs to maintain more label bindings but allows them to respond very quickly to routing failures. This is the more natural mode of operation for an IP routing node.
- In a cell-switched MPLS network, LDP executes in downstream-on-demand with conservative label-retention mode. This requires ATM LSRs to maintain fewer label bindings but causes them to respond less quickly to routing failures. This is a much less natural mode of operation for an IP routing node at the edge of a cell-switched MPLS core.

The impact of the different operational modes of LDP in packet-based and cell-based networks should not be overlooked when considering a transition from cell-switched MPLS to IP-routed MPLS.

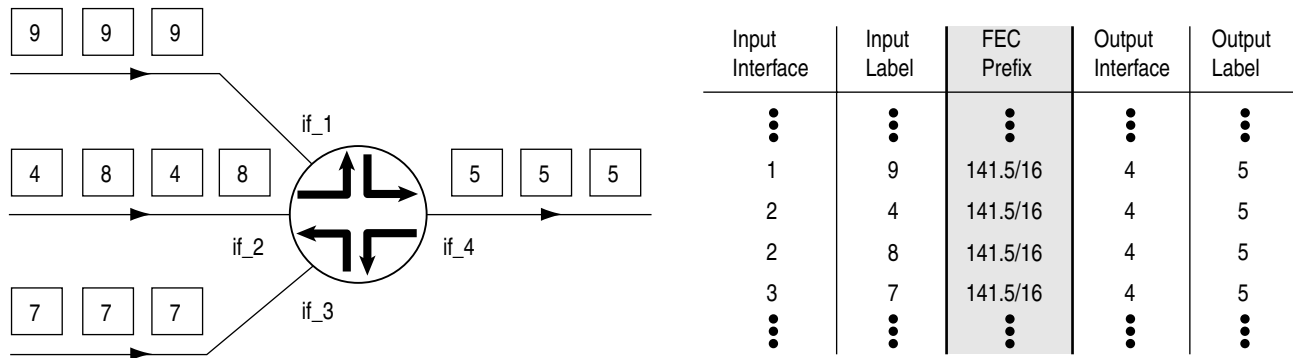
VC Merge and Cell Interleave Issues

When using frame-based technologies, two or more incoming labels from different traffic flows can be mapped to a single outgoing label to support the merging of multiple traffic flows into a single traffic flow. Packets from the different traffic flows can be interleaved as desired on links that are downstream of the merge point. However, when merging traffic flows over cell-based technologies, it is vital that cells associated with a given packet remain contiguous and are not interleaved with cells from other packets.

Label Merging in IP-routed MPLS Networks

An LSR can receive frames that are members of the same FEC on different incoming interfaces or with different labels on the same interface. If the LSR can forward all of the frames in the FEC over a single interface with the same outgoing label, then the LSR is capable of performing *label merging* (see Figure 10).

Figure 10: How an LSR Performs Label Merging



The ability of LSRs to support label merging can significantly reduce the total number of labels that need to be allocated across an LSR domain. An LSR capable of label merging needs only one outgoing label for each FEC, even if multiple requests for label bindings to that FEC are received from upstream neighbors. Without label merging, the total number of outgoing labels required for each FEC could be as large as the number of nodes in the network. On the other hand, once labels are merged and packets have been transmitted, the information indicating that the packets arrived on different interfaces or with different labels is forever lost.

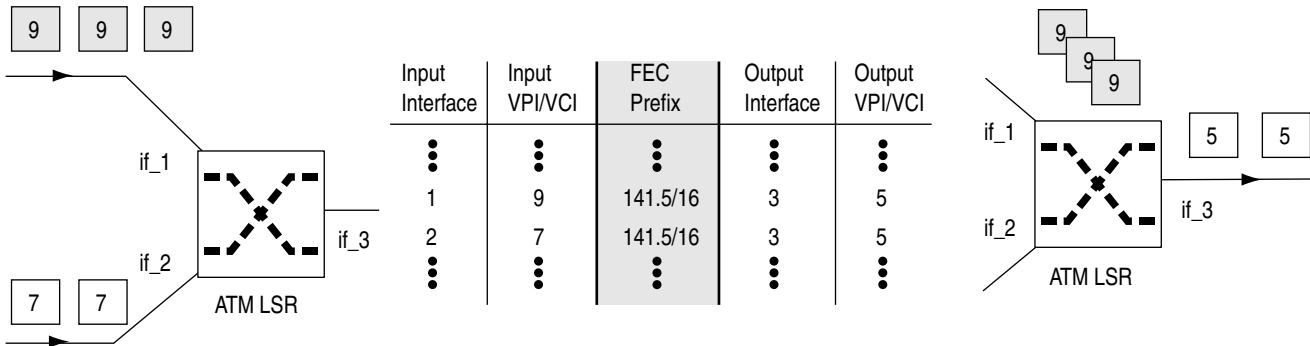
Label Merging in Cell-switched MPLS Networks

In a cell-switched MPLS network, the ingress frame-based edge set LSR encapsulates each MPLS frame in an AAL5 PDU that is segmented into ATM cells for transmission across an LC ATM interface. AAL5 requires that all of the cells segmented from a given AAL5 PDU be transmitted on the given VC in sequence; otherwise, they cannot be reassembled by the egress frame-based edge set LSR. Because cells that arrive simultaneously at an ATM LSR from different sources have the potential to be interleaved, each ATM LSR must ensure that cells do not interleave if they are merged on a common VC (the concatenation of the VPI and VCI) toward the same FEC. If an ATM LSR cannot guarantee that the cells will not be interleaved, then it is required to maintain more than one label binding for each FEC.

VC Merge

VC merge is the mechanism that allows an ATM LSR to receive cells on different incoming VCs and to transmit them over a single, outgoing VC without causing the cells of different AAL5 PDUs to become interleaved. To eliminate the cell interleave problem, the ATM LSR buffers all of the cells from each packet until the last cell of each packet is received. At this point, all of the cells from a particular AAL5 PDU are transmitted in sequence on the outbound link (see Figure 11).

Figure 11: How an ATM LSR Performs VC Merge

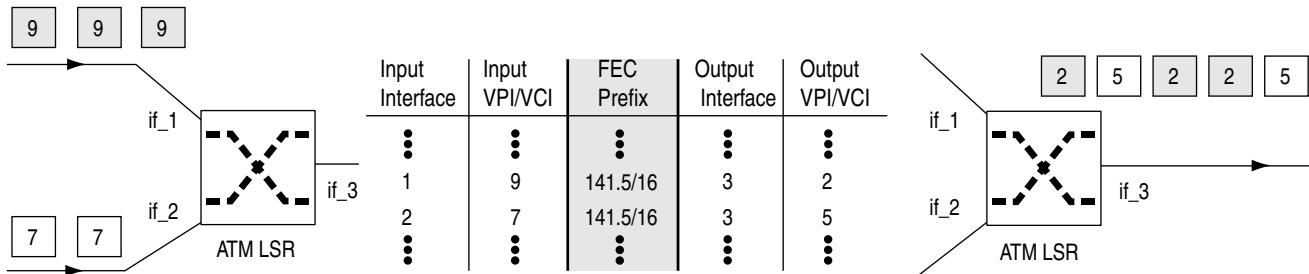


To support VC merge, an ATM switch may need specialized hardware or per-VC buffering capabilities. Thus, some ATM LSRs will be able to support this feature, while others will not.

Multiple Labels per FEC

If an ATM LSR does not support VC merge, then it is required to maintain more than one label binding for each FEC. In this solution, the ATM LSR that does not support VC merge uses downstream-on-demand mode to request a new label binding each time that one of its upstream neighbors makes a requests for a label binding. The non-merging ATM LSR requests a new label binding even if it already has a label binding for the specific FEC. The assignment of multiple labels to an FEC allows the ATM LSR to distinguish between flows it receives on different interfaces and to ensure that cells from different AAL5 PDUs are not interleaved (see Figure 12).

Figure 12: Assigning Multiple Labels to Eliminate Cell Interleave



Appendix II: References

Requests for Comments (RFCs)

- RFC 2684. Grossman, D. and Heinanen, J. "Multiprotocol Encapsulation over ATM Adaptation Layer 5." September 1999.
- RFC 2702. Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and McManus, J. "Requirements for Traffic Engineering over MPLS." September 1999.
- RFC 3031. Rosen, E., Viswanathan, A., and Callon, R. "Multiprotocol Label Switching Architecture." January 2001.
- RFC 3032. Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li T., and Conta, A. "MPLS Label Stack Encoding." January 2001.
- RFC 3034. Conta, A., Doolan, P., and Malis, A. "Use of Label Switching on Frame Relay Networks." January 2001.
- RFC 3035. Davie, B., Lawrence, J., McClooghrie, K., Rosen, E., Swallow, G., Rekhter, Y., and Doolan, P. "MPLS Using LDP and ATM VC Switching." January 2001.
- RFC 3036. Andersson, L., Doolan, P., Feldman, N., Fredette, A., and Thomas, B. "LDP Specification." January 2001.
- RFC 3038. Nagami, K., Katsubi, Y., Demizu N., Esaki H., and Doolan, P. "VCID Notification over ATM Link for LDP." January 2001.

Textbooks

- Davie, Bruce S. and Rekhter, Yakov. *MPLS: Technology and Applications*. Morgan Kaufmann Publishers, 2000 (ISBN 1558606564).
- Metz, Christopher. *IP Switching: Protocols and Architectures*. McGraw-Hill, 1998, (ISBN 0070419531).
- Pepelnjak, Ivan and Guichard, Jim. *MPLS and VPN Architectures: A Practical Guide to Understanding, Designing and Deploying MPLS and MPLS-enabled VPNs*. Cisco Press, 2000 (ISBN 1587050021).

Copyright © 2002, Juniper Networks, Inc. All rights reserved. Juniper Networks is registered in the U.S. Patent and Trademark Office and in other countries as a trademark of Juniper Networks, Inc. Broadband Cable Processor, G10, Internet Processor, Internet Processor II, JUNOS, JUNOScript, M5, M10, M20, M40, M40e, and M160 are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. All specifications are subject to change without notice.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.