Designing Multiprotocol Label Switching Networks

Jeremy Lawrence, Cisco Systems

ABSTRACT

Multiprotocol label switching adds to the capabilities of IP networks in several ways. Despite new capabilities, MPLS technology has much in common with ordinary IP networks. In turn, the design process for MPLS networks has much in common with the design of any IP network. This article examines MPLS and IP technology with particular emphasis on what is common between them. The common design steps of MPLS networks and other IP networks are outlined briefly, and those issues specific to MPLS networks are covered in more detail. This article emphasizes MPLS point of presence design, routing design issues for MPLS, and provisioning of sufficient label space.

INTRODUCTION

Multiprotocol label switching (MPLS) is an extension to the existing Internet Protocol (IP) architecture [1]. By adding new capabilities to the IP architecture, MPLS enables support of new features and applications, which have been described previously [2]. The new applications include traffic engineering [3, 4], IP virtual private networks [5, 6], integration of IP routing and layer 2 or optical switching [7], and other applications. The previous descriptions of MPLS have concentrated, quite rightly, on what is new in MPLS (i.e., the differences between MPLS and traditional IP routing and forwarding). However, an important consideration is that there is much in common between the technology of MPLS networks and IP routed networks. The overlap in technology means that there is much in common between the process of design of MPLS networks and the design of other IP routed networks. This article starts with a description of MPLS technology that concentrates as much on the similarities between MPLS and other IP routed networks as on the differences. The design process of MPLS networks is then outlined, and the new design issues specific to all MPLS networks are described in detail.

MPLS NETWORK STRUCTURES

An MPLS network such as that in Fig. 1a consists of edge label switch routers (edge LSRs)[1] around a core of label switch routers (LSRs). Customer sites are connected to the carrier's MPLS network, or equivalently a large organization's network backbone. Figure 1a shows nine customer sites and four edge LSRs, but more typically there will be hundreds or more customer sites per edge LSR. The customer premises equipment connected to an MPLS network typically runs ordinary IP forwarding rather than MPLS, and is typically a router or a LAN switch. Since the customer equipment typically does not run MPLS, the edge LSRs are part of the carrier's network and under the carrier's administration.

A carrier's MPLS network will often be connected to one or more other IP networks as part of the Internet. An IP connection to another carrier might be an MPLS link, although use of MPLS on intercarrier links is usually not required. As with any interconnection of carrier networks in the Internet, the Border Gateway Protocol [8] would typically operate over links to other carriers, to exchange routing information with them. The neighboring IP networks may use MPLS internally, but not necessarily.

The links between customer equipment, edge LSRs, and/or LSRs may be of virtually any type. Traditional non-MPLS switches may be used to carry connections to MPLS equipment, typically using some sort of permanent virtual circuits or optical lightpaths. Figure 1b illustrates several ways of doing this.

PROTOCOLS IN IP AND MPLS NETWORKS

IP PROTOCOLS

MPLS networks use a superset of the protocols used in ordinary IP routers, and most MPLS networks are fully interoperable with ordinary IP equipment. The IP protocol architecture, or IP stack, consists of a large number of protocols, a subset of which is used on IP routers, as opposed to hosts or servers. Figure 2a shows the most

[1] Edge LSRs are also known as label edge routers. Most commercially available LSRs have at least limited edge LSR capability, and vice versa. The term "edge LSR" conveys this overlap more effectively.

 0163-6804/01/$10.00

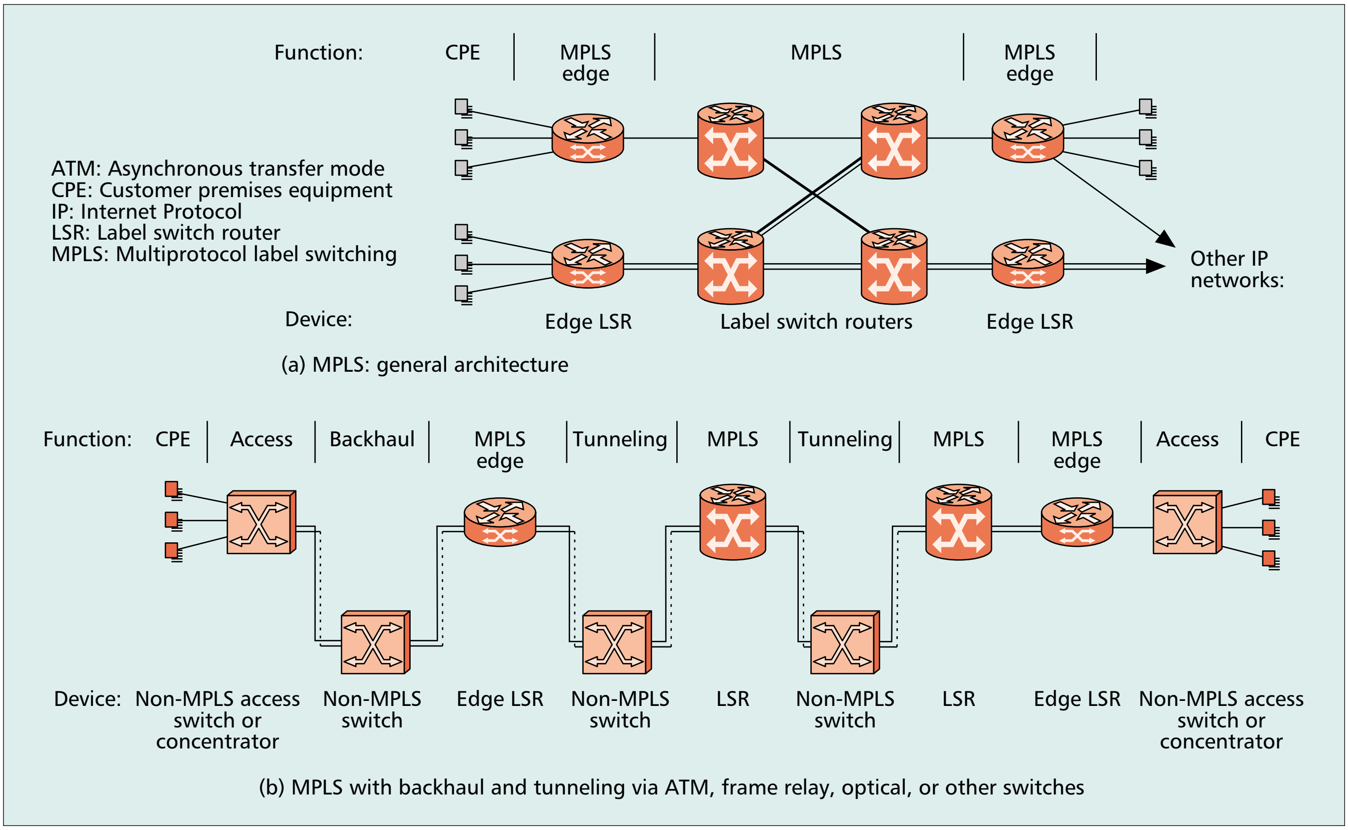

Figure 1. *MPLS network structures.*

important of these protocols as a basis of comparison with the MPLS stack.

The protocols that run on an IP router consist of routing and forwarding protocols, and others. Routing protocols inform a router about the routes to reach the known IP destinations. Each router will, in general, run an interior gateway (routing) protocol, which deals with routing inside a network. There is a choice of standard interior gateway protocols:

- The Routing Information Protocol has many limitations, and is becoming obsolete.
- Open Shortest Path First is widely used.
- Interior System to Interior System is similar to Open Shortest Path First, and is also common among large providers and carriers.

Typically, all routers under one distinct administration will run the same interior gateway protocol. In general, a router used to communicate between different administrations will run another routing protocol called the Border Gateway Protocol, which deals with exchange of information between administrative domains.

Routing protocols automatically build a table of forwarding information in each router. The forwarding table lists the IP destination prefixes that are known, and which data link on the router is the "next hop" when moving IP packets to each known destination prefix. In routers, the forwarding table is used by IP when moving packets toward their destination. IP may be referred to as a *forwarding protocol*, since it moves packets rather than determines routes for them. Version 4 of IP (IPv4) is the most widely used for IP applications, although use of IPv6 is also beginning.

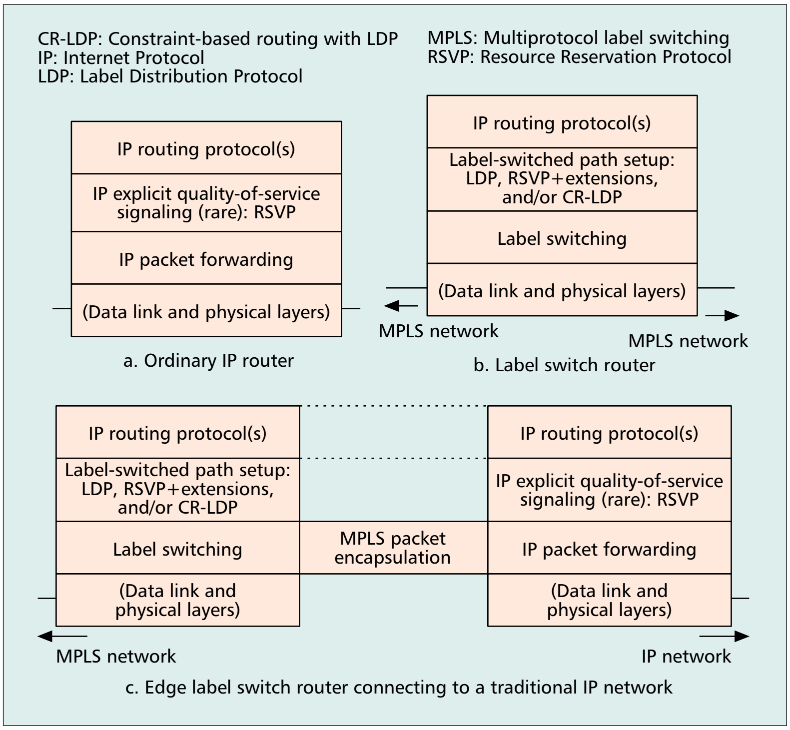

Figure 2. *Protocol stacks for IP and MPLS. These stacks are simplified for clarity; see the text.*

Many other protocols are used by routers. There are specific data link protocols for various media, and protocols to perform auxiliary functions. For example, routing protocols typically use the Transport Control Protocol or Universal

In both hop-by-hop routed MPLS and explicit routing, label distribution works in close conjunction with the IP routing protocol(s) in a network in order to set up LSPs.

Datagram Protocol as part of their communication process. However, these protocols are not used in the main function of routers, namely forwarding packets along the correct route to their destination. For a complete listing of all protocols in the IP stack, their standards status, and references to their RFC documents, see [8].

Terminology: IP Destination Prefixes — A destination prefix is a group of IP addresses that may be treated similarly for forwarding purposes. For example, the IPv4 destination prefix 10.1.1.0/24 represents the range of IP addresses in which the first three octets (24 bits) are 0000 1010, 0000 0001, and 0000 0001, with any values for the last octet. In other words, the destination prefix 10.1.1.0/24 represents any IPv4 address 10.1.1.*x*, where *x* is a value between 0 and 255 inclusive.

Protocols in the MPLS Network Core

IP Routing and Label-Switched Paths — As with ordinary routers, LSRs will run one or more IP routing protocols (Fig. 2b). However, instead of IP packet forwarding, LSRs use a different type of forwarding mechanism called *label switching*. Label switching relies on the setup of switched paths through the network; these are *label-switched paths* (LSPs). An LSP, in the simplest case, corresponds to an IP destination prefix that appears in the forwarding tables of several routers. For example, the interior gateway protocol instances in a succession of LSRs might create forwarding table entries for an IPv4 destination prefix 10.1.1.0/24. The LSRs then, in effect, compare their forwarding information to find that they all have a route to 10.1.1.0/24, and set up an LSP through themselves in order to carry all traffic for the destination prefix 10.1.1.0/24. Such setup of an LSP according to IP routing information is known as *hop-by-hop routed MPLS*.

Setup of LSPs is done by a process of *label distribution*. There are several valid combinations of protocols for label distribution:

- In general, all MPLS networks use the Label Distribution Protocol, which supports hop-by-hop routed MPLS.
- One alternative for explicitly routed MPLS is the Label Distribution Protocol in combination with the Resource Reservation Protocol plus some extensions.[2]
- The other alternative for explicitly routed MPLS is the Label Distribution Protocol in combination with some extensions, called *constraint-based LSP setup using LDP*.

Explicit routing is a mode of MPLS operation where LSPs are set up to override the normal hop-by-hop routed path chosen by an IP routing protocol for selected traffic. The normal application is in IP traffic engineering [3, 4]. Although explicit routing and traffic engineering override the routes from the IP routing protocol, they do not ignore the IP routing protocol. The IP routing protocol information must be consulted in order to ensure that the explicitly routed paths do not lead to looping routes. In addition, traffic engineering makes use of extensions to the interior gateway protocols to carry information about traffic loads in the network. In both hop-by-hop routed MPLS and explicit routing, label distribution works in close conjunction with the IP routing protocol(s) in a network in order to set up LSPs.

Label Encapsulation and Packet Forwarding — When a packet is sent on an LSP, a label is applied to the packet (Fig. 3). There are several distinct means of carrying labels on packets. In the generic label encapsulation, which may be used with any link type, the label is carried in an extra header applied to each packet between the data link header and the layer 3 header. On link types that inherently support the concept of a virtual channel, the label may be carried as virtual channel information, typically in a layer 2 header. On asynchronous transfer mode (ATM) links, the label may be carried as the virtual circuit identifier and/or virtual path identifier applied to each cell of the packet. On frame relay links, the label may be carried as the data link connection identifier applied to each packet. A similar scheme has also been proposed for synchronous optical network links. MPLS over optical switches similarly uses virtual channels, with each distinct optical wavelength used on a link constituting a different label [7].

Furthermore, two distinct types of MPLS can be defined:

- Packet-based MPLS is where the LSRs have full packet-handling capacity, and can examine layer 3 headers on packets.[3] Packet-based LSRs typically also have IP router function. Edge LSRs are typically a type of packet-based LSR.
- Switch-based MPLS is the form of MPLS where LSRs, known as *switch-based LSRs*, forward packets by means of layer 2 switching or optical switching. Switch-based LSRs have little or no capability to examine layer 3 headers.

Any type of label encapsulation can be supported by either class of LSR. For example, packet-based LSRs may have optical links using distinct lightpath wavelengths as labels; or there may exist switch-based LSRs dealing with the generic MPLS label encapsulation, but without the ability to deal with layer 3 headers. However, most switch-based LSRs used to date have been based on ATM switches. The use of switch-based MPLS on ATM switches is also referred to as *ATM MPLS*.

On each link, each distinct label value corresponds to a distinct LSP. The effect of carrying label information with each packet is most easily understood in switch-based MPLS. Because the MPLS protocols associate a virtual channel with an LSP at each hop, a virtual channel connection is set up end to end across the network to support the LSP. With packet-based MPLS, the generic label encapsulation allows labeling to occur over almost any link type, not just those that inherently support virtual channels. Packet-based MPLS equipment implements MPLS forwarding by mapping incoming label values to next-hop interfaces and outgoing label values; these mappings define a type of switching operation. Again, a switched path across the network is established across the network; this is the LSP. The combination of the protocols in the MPLS protocol stack allows LSPs to be established in conjunction with IP routing, and across any type of link.

[2] It is, in principle, possible to run an MPLS network using the Resource Reservation Protocol plus extensions, but without the use of the Label Distribution Protocol. However, it is not clear whether any MPLS implementation supports this.

[3] Packet-based MPLS is also known as frame-based MPLS. The latter term is avoided here because of possible confusion with frame relay.

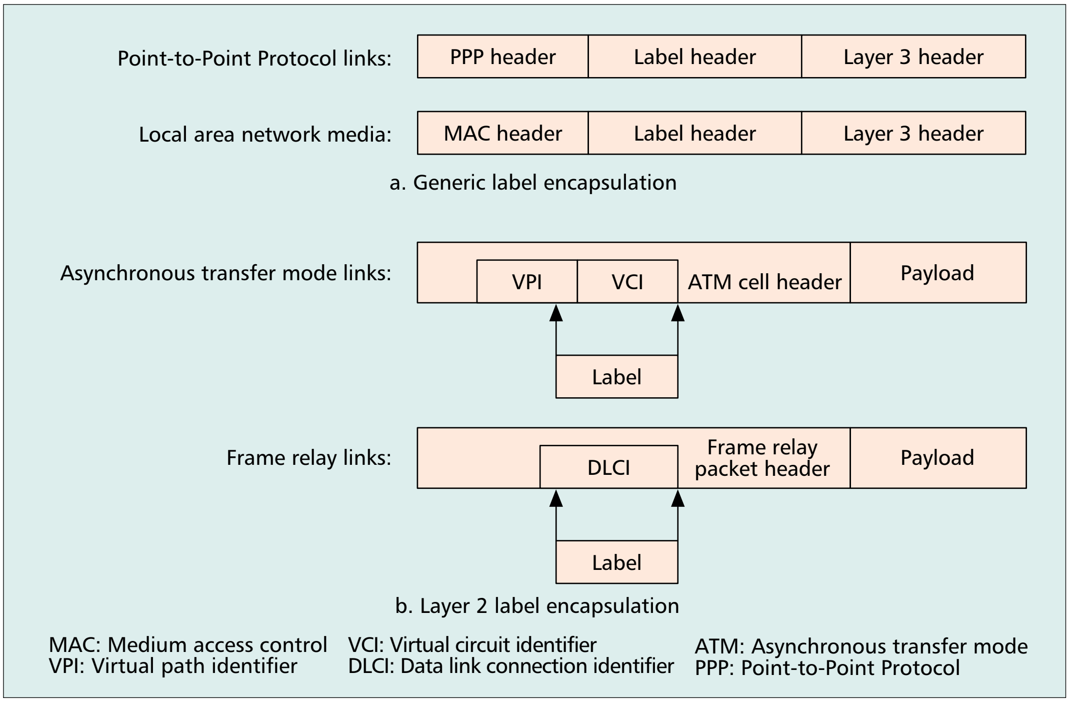

Figure 3. *Carrying a label on a packet.*

For the purposes of packet forwarding, the "traditional IP" and "MPLS" parts of the edge LSR protocol stack are linked by the process of MPLS encapsulation.

PROTOCOLS IN MPLS EDGE DEVICES

Edge LSRs usually connect to both MPLS networks and traditional IP networks. Consequently, edge LSRs typically run all the protocols found on a traditional IP router, including IP packet forwarding, in addition to MPLS. This is illustrated in Fig. 2c.[4] IP routing protocols make no distinction between MPLS and traditional IP networks, which means a single routing protocol process in an edge LSR may support both MPLS and traditional IP links. For the purposes of packet forwarding, the "traditional IP" and "MPLS" parts of the edge LSR protocol stack are linked by the process of MPLS encapsulation.

In simple operation, MPLS encapsulation involves the application (or "pushing") of a label onto an IP packet, which is then forwarded across the network on the corresponding LSP. When the packet reaches the edge of the MPLS network, the label is then removed (or "popped") from the packet, which is then forwarded as an ordinary IP packet.

MPLS also allows more than one label to be carried on each packet: a "stack" of labels of arbitrary size may be carried using the generic MPLS label header. Multiple labels are used in MPLS traffic engineering [3, 4], and virtual private networks [5, 6].

STEPS IN DESIGNING MPLS NETWORKS

The previous discussion has shown that MPLS networks and other IP networks have many protocols in common. Their common use in MPLS and other IP networks means that MPLS shares many design steps in common with all IP routed networks. Consequently, most of the design steps for an MPLS network are in common with those of other IP networks. A typical design process is as follows:

1. Locate and design the points of presence (PoPs).
2. Dimension the backbone links in the network.
3. Design the IP routing.
4. Dimension the MPLS label space.
5. Configure any MPLS applications or advanced features, such as traffic engineering and virtual private networks.

The final design step is an ongoing process:

6. Refine and extend the design once the network is operational.

With the exception of steps 4 and 5, most of these steps are not specific to MPLS networks, and apply to IP routed networks of any type. These steps are described below, with emphasis on those considerations that are new to MPLS. Traffic engineering and virtual private networks (step 5) are not covered in this article, which concentrates on design issues common to all MPLS networks. The design issues unique to MPLS traffic engineering and virtual private networks would require more space to describe than is available here.

POINT OF PRESENCE DESIGN AND CHOICE OF EQUIPMENT

The design of PoPs for all IP networks, including MPLS networks, is constrained by:

- The choice of access link type(s) to be supported by the network.
- The choice of core link type(s) to be used in the network.
- Requirements for reliability, including use of warm or hot-standby redundant trunk cards, redundant processors, etc. These may constrain the choice of equipment to be used in a PoP.

[4] *In contrast is a recently proposed MPLS application where an MPLS network provides generic packet transport service, similar to the service of a traditional frame relay or ATM network. In such a service, IP packet forwarding is not used in the edge LSRs, which instead use a more rudimentary forwarding mechanism: all packets from a specified customer link are forwarded onto a particular LSP, and vice versa. In such a network, IP routing protocols run only in the MPLS network core, and not on the customer links.*

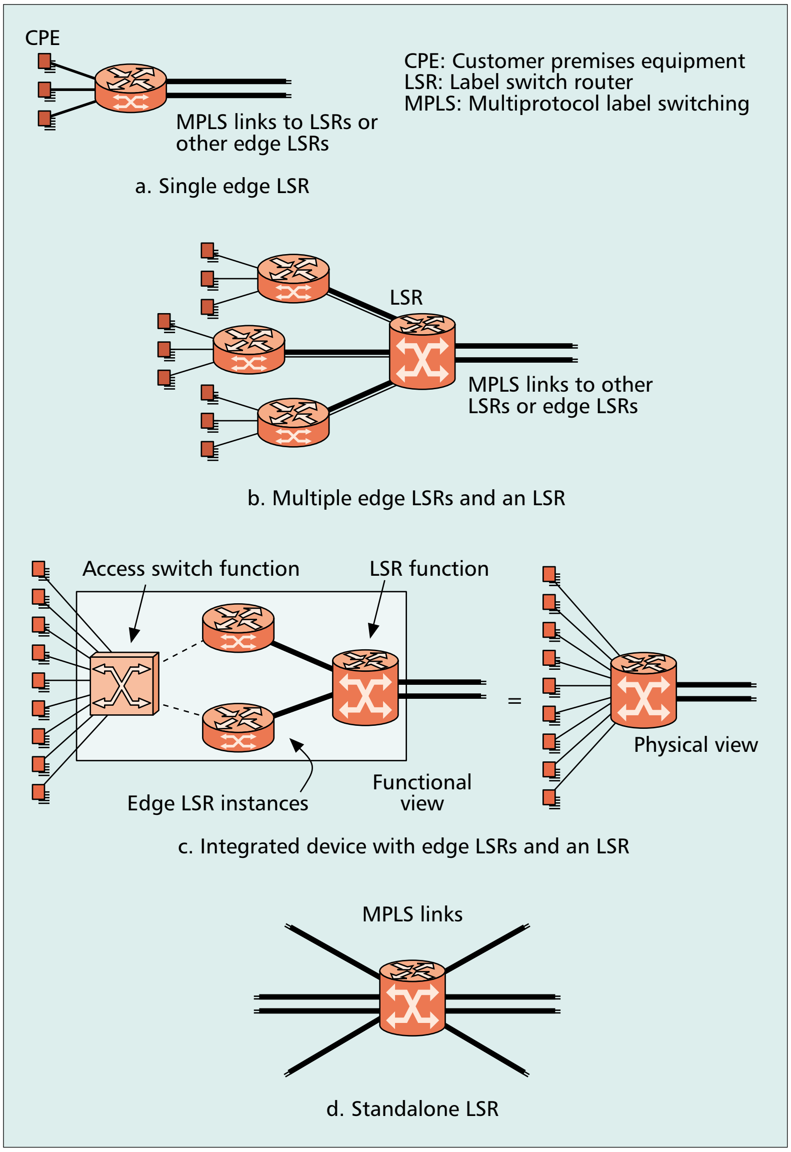

Figure 4. *PoP designs for MPLS.*

- Requirements for services beyond IP or MPLS, and availability of equipment meeting these requirements. For example, several commercial edge LSRs also have the capability of providing traditional ATM and frame relay edge switch function, and others offer only IP and MPLS edge function.
- Location of PoPs, which is largely determined by where the cities are.
- The population of end-user sites surrounding each location.

Once these are taken into account, MPLS equipment offers many different options for PoP designs. Some typical PoP designs are shown in Fig. 4.

A small PoP might consist of just a single edge LSR with a number of access lines from less than 10 to thousands. However, the population of user lines for a PoP often leads to a requirement for a PoP to have several edge LSRs. In a PoP with several edge LSRs, the use of an extra LSR as an aggregation device has several scalability advantages:

- It reduces the number of routing protocol peerings out of the PoP, assisting routing protocol scalability.
- It will assist in preventing local traffic from having to leave the PoP. In other words, traffic for which both the source and edge destination edge LSR are at the same PoP will not be required to leave the PoP.
- It will often reduce the number of links required from the PoP.
- In an MPLS network using ATM or frame relay switching and label merging (see “Choosing Core LSRs” below), it reduces the number of virtual circuits required from the PoP.

All except the last point also apply to a PoP in a traditional IP network with analogous structure.

At least two vendors produce devices that integrate the functions of an access switch, several edge LSRs, and an LSR. These provide the same functions as using several discrete devices, but with the practical advantages of having a single device: less space, somewhat lower power requirements, and so forth.

Choosing Core LSRs — LSRs without edge links will often be used, either as part of an edge PoP or as standalone devices. Aside from the lack of edge links, the design considerations are broadly the same as those listed above for all MPLS PoPs. Some specific considerations for core LSRs are:

- The number, as well as type, of core links to be supported: It is more important for a core LSR to support many core links than for an edge LSR.
- Label merging capability: All packet-based LSRs inherently support merging LSPs to a common destination, but such capability is less common for switch-based LSRs. Those switch-based LSRs that do support label merging are typically based on ATM or frame relay switches supporting virtual circuit merging. In an ATM or frame relay network, label merging capability brings a significant scalability advantage by reducing the number of virtual circuits required.

Redundancy and Reliability — In a broad sense, the reliability issues with IP and MPLS equipment are the same as for other telecommunications equipment. A network typically must be robust to device and link failures. However, the overall reliability of IP and MPLS networks is typically less dependent on the reliability of individual devices than in circuit or virtual circuit networks.

Some circuit (or virtual circuit) switches in connection-oriented networks have the ability to reroute circuits around failed links or nodes. Rerouting of circuits in a connection-oriented network typically takes a significant duration. This is due to the requirement to recalculate routes, often on a connection-by-connection basis, and reprogram some hardware for each changed connection. Consequently, rerouting in a connection-oriented network is a last resort, to be used only after redundant links or nodes have failed to prevent an outage. In IP routed networks, on the other hand, large aggregates of traffic can be rerouted simultaneously, and less (or no) hardware programming is typically required. Consequently, the reliability of many IP routed networks has relied primarily on rerouting and network-level redundancy (i.e., having a choice of several routes between each source and destination). Although IP routers are gaining similar redundancy features to circuit switches, network designers are still free to rely more on network-level redundancy in IP networks than in circuit or virtual circuit networks.

Packet-based MPLS has redundancy options similar to those of IP routing. Furthermore, there is recent work on protection switching in MPLS, which involves setting up a single LSP as an alternate to a failed link that might carry many LSPs. This capability, also known as *fast reroute*, is a new type of protection switching, and is unique to MPLS networks [9].

Rerouting may be somewhat slower in switch-based MPLS networks than in packet-based MPLS or IP routed networks, because switch-based MPLS equipment typically requires more hardware programming to reroute LSPs. In addition, ATM, frame relay, and optical switching often do not support the pushing and popping of extra labels, as required for rerouting of traffic from a physical link onto an LSP. Together with slower connection programming speeds, this means that switch-based MPLS is not as capable as packet-based MPLS of supporting fast reroute. Consequently, switch-based MPLS may sometimes have to rely more on device-level redundancy than packet-based MPLS.

LINK DIMENSIONING

The dimensioning of bandwidth of links in an MPLS network is essentially the same as for any other IP network. A typical process starts with an estimate of a PoP-to-PoP traffic matrix based on user populations at each PoP and an estimate of the proportion of traffic destined to the wider Internet as opposed to the network in question. Based on this matrix, an initial estimate of the required link capacities can be calculated, and then used to provision the network. MPLS does provide one important benefit with respect to link dimensioning: MPLS is an effective traffic engineering tool, which can be used to fit traffic to available link capacities [3, 4].

DESIGNING IP ROUTING FOR AN MPLS NETWORK

Every LSR uses an ordinary IP interior gateway routing protocol. As in an ordinary IP network, the interior gateway protocol determines the routes for traffic, which in MPLS is carried in LSPs. An important implication is that the interior gateway protocol sees an MPLS network as being exactly like an ordinary router network. As illustrated in Fig. 5, it is possible to have various viewpoints of an MPLS network:

- Physical viewpoint: This viewpoint represents the physical devices and links in a network.
- Functional viewpoint: Where a product has several functions, these can be shown separately. For example, Fig. 5 includes two MPLS edge devices of a type available today. Each of the two devices includes two functionally separate edge LSRs and an LSR. In addition, each of the edge devices includes a permanent virtual circuit switching function that is functionally separate from the LSR function.
- Routing viewpoint: This viewpoint is the network as it is seen by an IP routing protocol. It is derived from the functional viewpoint as follows:
  - –Layer 2 switches and permanent virtual circuit switching functions are invisible to IP routing. If a customer site is connected to a router by a permanent virtual circuit, the virtual circuit is seen by IP routing as a single-hop direct connection. For example, note the sites labeled a in Fig. 5c, and assume that these are connected to edge LSR b. Then, in the routing viewpoint, the sites are directly adjacent to router b.
  - –Each edge LSR or LSR constitutes a router in the routing viewpoint.

Designing IP routing in an MPLS network is generally the same process as designing IP routing for an ordinary IP network [10]. With reference to the routing viewpoint, a network can be divided into interior gateway protocol areas, route summarization can be designed, and so on.

IP Routing Issues Specific to MPLS Networks — There are a small number of routing issues specific to MPLS networks:

- MPLS affects the choice of interior gateway protocol. Specifically, MPLS traffic engineering uses a link-state protocol such as Open Shortest Path First or Interior System to Interior System to distribute information about the traffic loading on links. Distance-vector protocols such as the Routing Information Protocol do not support MPLS traffic engineering.
- Some IP routing protocol implementations allow unnumbered IP links. These are links without an IP address: another IP address on a router, such as a loopback address, is used to substitute for the link address for routing protocol purposes. The use of unnumbered IP links may be desirable in an MPLS network, particularly if it is switch-based. Unnumbered links reduce the number of IP destination prefixes known to the LSRs, and hence may reduce the number of labels that will be used in the network.
- Route summarization must not be done at a switch-based LSR. Multiple interior gateway protocol areas can be used in a switch-based MPLS network, as illustrated in Fig. 6. A switch-based LSR may be used as an *area border router*, but only if no summarization is done at the area border router.

It is common for area border routers to be configured to summarize routes. For example, an area border router may receive several routes from within one area to destinations such as 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24. In order to reduce routing table sizes, it may be desired to advertise them into another area by way of a single *summary* route (e.g., a route to 10.1.0.0/16.) However, such summarization may not be done if the area border router is also a switch-based

Rerouting may be somewhat slower in switch-based MPLS networks than in packet-based MPLS or IP routed networks, because switch-based MPLS equipment typically requires more hardware programming to reroute LSPs.

Many of the issues of designing MPLS networks are similar to those of designing ordinary IP networks. However, one issue in MPLS networks that does not occur in other IP routed networks is the provisioning of label space.

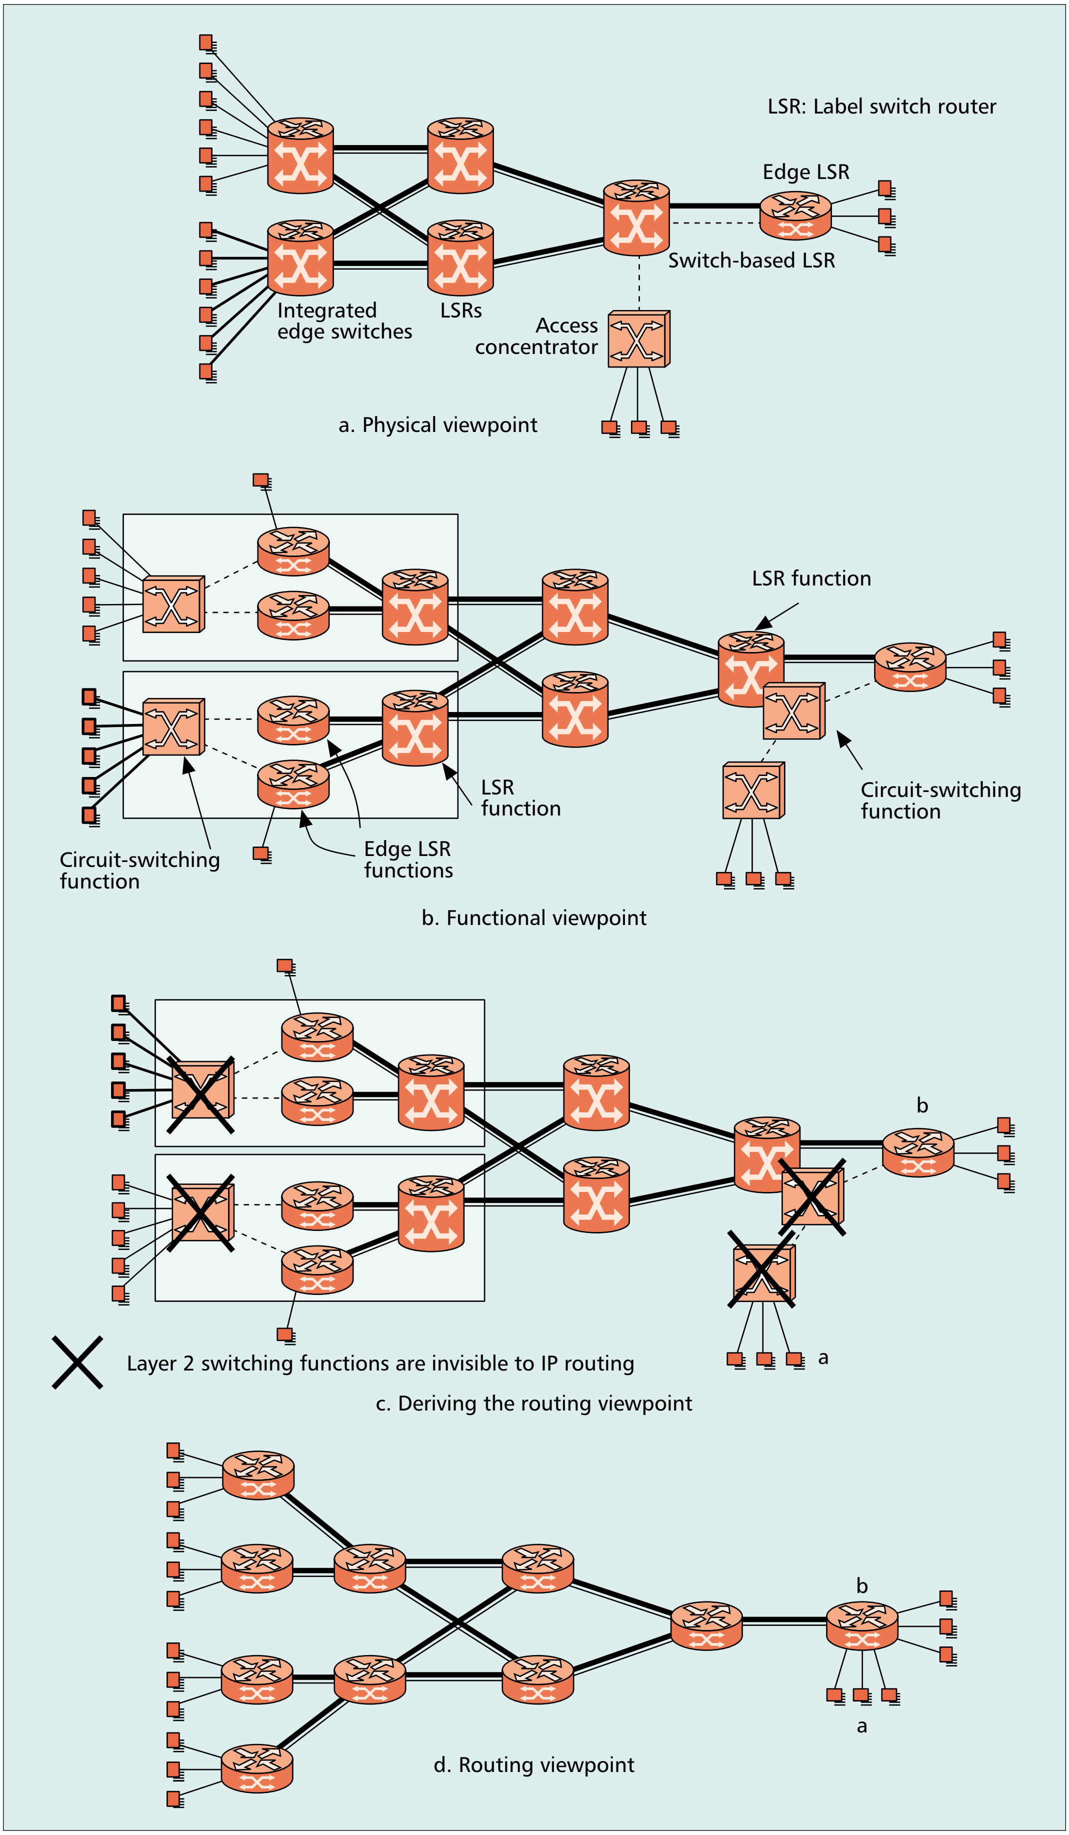

Figure 5. *Viewpoints of an MPLS network.*

LSR. If route summarization is required at the area border routers in a switch-based MPLS network, the area border routers must be packet-based LSRs. This requirement is not a contradiction, since it is possible for a packet-based LSR to support switch-based MPLS links.

- The previous rule also applies to the Border Gateway Protocol. *Autonomous system boundary routers* are those routers which communicate between administrative domains using the Border Gateway Protocol. Such routers often summarize routes. Consequently, a switch-based LSR may not be an autonomous system boundary router, but a packet-based LSR may be one.

The restrictions on summarization exist because summarization stops some LSPs being set up end to end. For example, assume that an area border summarizes reachability for 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24 with a single route for 10.1.0.0/16. Now assume that a packet with destination IP address 10.1.1.55 arrives with a label for 10.1.0.0/16. The area border router cannot label switch the packet. It must look past the label and examine the IP address to find that the packet should go on to 10.1.1.0/24. Since switch-based LSRs cannot examine layer 3 headers and IP addresses, they may not do IP route summarization.

Dimensioning MPLS Label Space

Many of the issues in designing MPLS networks are similar to those in designing ordinary IP networks. However, one issue in MPLS networks that does not occur in other IP routed networks is the provisioning of label space.

A certain number of LSPs are carried on each link in an MPLS network. Many LSRs support label merging, which means that LSPs to the same destination can be merged in the LSR, so only one label is required to that destination on each link. In a network using label merging and hop-by-hop routed MPLS, bounds on label usage are quite simple to calculate. For example, if labels in a network are required for 1000 distinct IP destination prefixes, and five distinct labels are required to each destination for carrying different classes of service, no more than 5000 labels will be required on each link.

Every LSR supports a limited number of labels on its links, and this limit is implementation-specific. The generic MPLS label encapsulation allows a 20-bit label range, which means that a link may support up to 1,048,576 (1M) distinct labels. However, many LSR implementations, particularly of switch-based LSRs, support far fewer than this. Consequently, an important part of MPLS network design is to calculate the label space requirements on each link, and compare this to the number of active labels supported on each link.

Label Usage and Hop-by-Hop Routed MPLS — In hop-by-hop routed MPLS, labels are required to enable LSPs to correspond to the known IP destination prefixes. A routing table, as built by an IP routing protocol, lists the destination prefixes known in a routing protocol area. The forwarding table will list destination prefixes of several types:

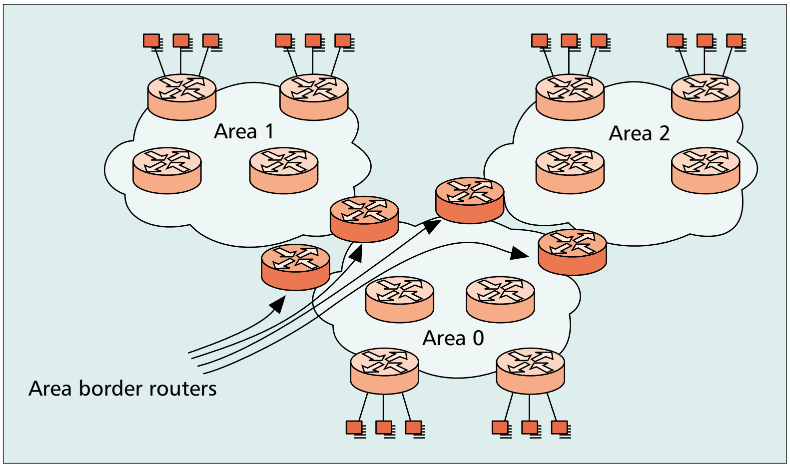

Figure 6. *The routing viewpoint of a network with multiple routing areas. There are restrictions on the use of label switch routers as area border routers.*

- The subnet address prefix of any numbered point-to-point link, or any other subnet, is a destination prefix. For this reason, it is sometimes desirable to use unnumbered links, to reduce the number of destination prefixes.
- Any other IP address of a device in the area contributes a destination prefix. Such addresses may include host addresses and loopback addresses on LSRs.
- There usually will be destination prefixes from outside the area, but advertised in the area by area border routers (or autonomous system boundary routers). If many address prefixes are summarized into a single address prefix before advertisement in the area, this counts as a single destination prefix.

Labels are not necessarily required for all destination prefixes in the forwarding table. MPLS implementations may allow filtering of the forwarding table when determining the destination prefixes for which labels are required. This might be used, for example, to prevent labels being set up to the destination prefixes of numbered links. In general, labels are required for a subset of the destination prefixes in the forwarding table, of size d.

In networks using label merging, a limit on the number of active labels required per link is straightforward to calculate, as shown above. The limit is cd, where c is the number of different labels required to each of the destination prefixes. $c = 1$ unless distinct labels are required for multiple service classes, in which case it will typically be 2–5. In networks without label merging, the limits are much larger.

If insufficient label space is available on a link, some required LSPs will be blocked. Such blockage will generally lead to a lack of complete connectivity in the network and inability to deliver some hop-by-hop routed traffic. Consequently, it is important to ensure that the number of labels available is sufficient.

Equipment Limits — Once the label space requirements have been analyzed, they can be compared to the capabilities of the equipment in the network. The nature of the limits on active

MPLS significantly extends the capabilities of IP networks by introducing support of VPNs and traffic engineering, and with integration of IP routing with various "later 2" switching technologies.

labels in commercial equipment will vary widely. There may be limits per LSR, per link, and/or of other types. For example, some LSRs support a certain number of active labels per line card, shared among several links. In LSRs supporting label merging, there may be limits on the number of multipoint-to-point trees of LSPs merging onto a single link, distinct from the number of labels supported on the link. Optical LSRs may have further constraints on exactly how labels are used. For example, some optical LSRs may not allow the label value (i.e., lightpath wavelength) to be changed when an LSP is switched from link to link. Other optical LSRs may allow changes only by a certain deviation in lightpath wavelength. However, dealing with such constraints in optical LSRs is a matter for further research. In any case, full details of the constraints on label usage should be obtained from the manufacturer(s) of the particular LSRs used in a network.

ONGOING REFINEMENT

Network design is an ongoing process. Once an MPLS network is deployed, it has ongoing design requirements that have much in common with ordinary IP networks. One important activity is to measure the actual PoP and link traffic, compare them against the assumptions used in the initial design, and make modifications as necessary. This process is repeated as traffic flows increase.

As customers are added and traffic increases, network changes will be required, including:

- Adding new edge LSRs to PoPs to cope with increasing numbers of customer links
- Adding new links to the network
- Adjusting traffic engineering parameters
- Checking for allocation of sufficient label space on links

CONCLUSIONS

MPLS significantly extends the capabilities of IP networks by introducing support of virtual private networks and traffic engineering, and with integration of IP routing with various layer 2 switching technologies. MPLS achieves these new capabilities by extending the IP protocol architecture, rather than replacing it. In general, all LSRs will use IP routing protocols. Much of the complexity of routed IP network configuration is in the configuration of the routing protocol(s), and the same is true for the design of MPLS networks.

Most design steps for an MPLS network correspond to those in designing an ordinary IP network, but have some extra considerations that have been described in this article. The most important extra considerations are that MPLS introduces new options for designing reliability and redundancy in networks, and MPLS networks have some constraints on IP routing protocol configuration. On the other hand, MPLS introduces a new design step that is not found in the design of other IP networks: ensuring sufficient provision for active labels on links. The advantages of MPLS stem from the introduction of a general concept of labels and label switching; consequently, MPLS function relies on having sufficient label space available.

REFERENCES

[1] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," RFC 3031, Jan. 2001.

[2] G. Armitage, "MPLS: The Magic Behind the Myths," *IEEE Commun. Mag.*, vol. 38, no. 1, Jan. 2000, pp. 124–31.

[3] D. Awduche, "MPLS and Traffic Engineering in IP Networks," *IEEE Commun. Mag.*, vol. 37, no. 12, Dec. 1999, pp. 42-47.

[4] G. Swallow, "MPLS Advantages for Traffic Engineering," *IEEE Commun. Mag.*, vol. 37, no. 12, Dec. 1999, pp. 54–57.

[5] E. Rosen and Y. Rekhter, "BGP/MPLS VPNs," RFC 2547, Mar. 1999; E. Rosen *et al.*, "BGP/MPLS VPNs," Internet draft, draft-rosen-rfc2547bis-03.txt, work in progress.

[6] K. Muthukrishnan and A. Malis, "Core MPLS IP VPN Architecture," RFC 2917, Sept. 2000.

[7] A. Banerjee *et al.*, "Generalized Multiprotocol Label Switching: An Overview of Routing and Management Enhancements," *IEEE Commun. Mag.*, vol. 39, no. 1, Jan 2001, pp. 144–50.

[8] Internet Engineering Task Force, "Internet Official Protocol Standards," RFC 2600. (Note that every 100th RFC document is an update to this document, and this document will automatically be superseded by RFC 2700, 2800, etc. as published.)

[9] V. Sharma, *et al.*, "Framework for MPLS-based Recovery," Internet draft, draft-ietf-mpls-recovery-frmwrk-02.txt, work in progress.

[10] B. Halabi, *Internet Routing Architectures*, New Riders Publishing 1997.

BIOGRAPHY

JEREMY LAWRENCE (jlawrenc@cisco.com) is a consulting engineer at Cisco Systems. He received a B.E.(Hons.) degree from the University of Canterbury, New Zealand, in 1991, and a Ph.D. from the University of Wollongong, Australia, in 1995. Since 1996 he has worked on MPLS issues at Cisco Systems. This work started with design of aspects of ATM switching support in the MPLS protocols, and design of MPLS support in the first carrier-class ATM label switch router. He has since carried out extensive work on the design of ATM MPLS networks, both in application to specific networks and with regard to the design process itself.