MULTIPROTOCOL LABEL SWITCHING (MPLS) - II

Routing in MPLS Domains

Figure 6.7 shows the label-switching paradigm in an MPLS network. An ingress LSR is an edge device that performs the initial packet processing and classification and applies the first label. An ingress LSR creates a new label. A core LSR swaps the incoming label with a corresponding next-hop label found from a forwarding table. At the other end of the network, another edge router, the egress LSR, is an outbound edge router and pops the label from the packet. It should be noted that multiple labels may be attached to a packet, forming a stack of labels. Label stacking enables multilevel hierarchical routing. For example, BGP labels are used for higher-level hierarchical packet forwarding from one BGP speaker to the other, whereas Interior Gateway Protocol (IGP) labels are used for packet forwarding within an autonomous system. Only the label at the top of the stack determines the forwarding decision.

Figure 6.5. Multiple layer 2 switching example in MPLS

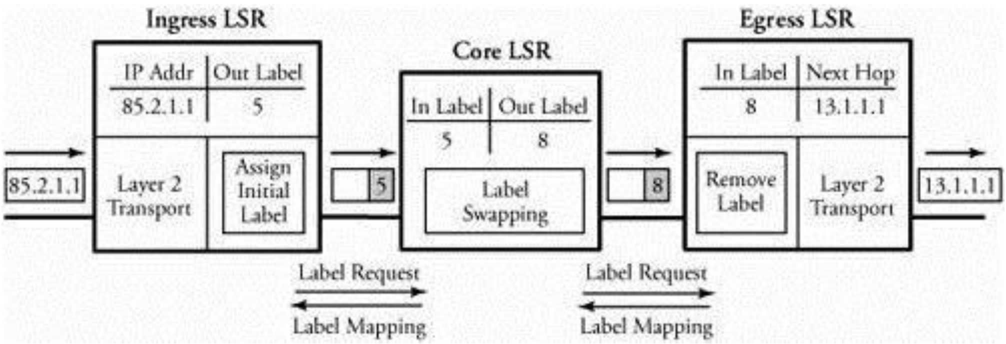

Once an IP packet enters an MPLS domain, the ingress LSR processes its header information and maps that packet to a forward equivalence class (FEC). At this point, a label switch path (LSP) through the network must be defined, and the QoS parameters along that path must be established. The QoS parameters define how many resources are to be used for the path and what queueing and discarding policy are to be used. For these functions, two protocols are used to exchange necessary information among routers: An intradomain routing protocol, such as OSPF, is used to exchange routing information, and the Label Distribution Protocol (LDP) assigns labels. At the end of the process, the router appends an appropriate label for FEC purposes and forwards the packet through.

Packet forwarding at the core LSR is based on a label-swapping mechanism. Once it receives a labeled packet, the core LSR reads the label as an index to search in the incoming label map table for the corresponding next-hop label. The label in the MPLS header is swapped with the out-label and sent on the next hop. This method of packet forwarding simplifies the routing process by replacing the longest-prefix match of IP routing with simple short-label exact-match forwarding. The real benefit of this method is that instead of processing IP headers for forwarding packets, routers process a short label. Once a packet arrives at the egress LSR, its MPLS header is decapsulated, and the stripped packet is routed to its destination.

In summary, an MPLS domain has three label manipulation instructions: An ingress LSR creates a new label and pushes it to the label stack of a packet, a core LSR swaps the incoming label with a corresponding next-hop label found from the forwarding table, and an egress LSR (outbound edge router) pops a label from the label stack. Only the label at the top of the stack determines the forwarding decision. The egress LSR strips the label, reads the IP packet header, and forwards the packet to its final destination.

Tunneling and Use of FEC

In an MPLS operation, any traffic is grouped into FECs. FEC implies that a group of IP packets are forwarded in the same mannerfor example, over the same path or with the same forwarding treatment. A packet can be mapped to a particular FEC, based on the following criteria:

- Source and/or destination IP address or IP network addresses
- TCP/UDP port numbers
- Class of service
- Applications

As mentioned earlier, labels have only local significance. This fact removes a considerable amount of the network-management burden. An MPLS packet may carry as many labels as required by a network sender. The process of labeled packets can always be performed based on the top label. The feature of label stack allows the aggregation of LSPs into a single LSP for a portion of the route, creating an MPLS tunnel. Figure 6.8 shows an IP packet moving through an MPLS domain. When the labeled packet reaches the ingress LSR, each incoming IP packet is analyzed and classified into different FECs. This traffic-classification scheme provides the capability to partition the traffic for service differentiation.

Figure 6.8. An IP packet labeled in an MPLS domain and tunneled to reach the other end of the domain

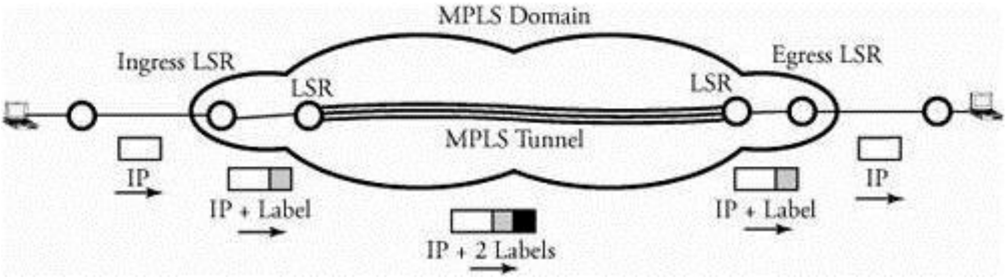

Route selection can be done either hop by hop or by explicit routing. With hop-byhop routing, each LSR can independently choose the next hop for each FEC. Hop-byhop routing does not support traffic engineering, owing to limited available resources. Explicit routing can provide all the benefits of traffic engineering. With explicit routing, a single LSR determines the LSP for a given FEC. For explicit routing, LSRs in the LSP are identified, whereas in an explicit routing, only some of the LSRs in an LSP are specified.

With the introduction of constraint-based routing, FEC can segregate the traffic into different levels of QoS, each with different service constraints, to support a variety of services, such as latency-based voice traffic and security-based VPN. At the beginning of the tunnel, an LSR assigns the same label to packets from a number of LSPs by pushing the label onto each packet's stack. At the other side of the tunnel, another LSR pops the top element from the label stack, revealing the inner label.

Source : http://elearningatria.files.wordpress.com/2013/10/cse-vi-computer-networks-ii-10cs64-notes.pdf