

MPLS

Multiprotocol Label Switching

MPLS Overview

Multiprotocol Label Switching (MPLS) brings the speed of Layer 2 switching to Layer 3. It allows routers to eliminate the complex lookup process based on the destination IP address, and make forwarding decisions based on the contents of a simple label. In addition to speed, MPLS offers two key advantages: it supports Quality of Service (QoS) and Virtual Private Networks (VPNs).

MPLS Operation

An MPLS network permits the definition of explicit paths, which are predefined routes through networks. MPLS routes packets along these pre-configured paths, called Label Switched Paths (LSPs).

Standard routing protocols, such as OSPF (Open Shortest Path First) and BGP (Border Gateway Patrol), determine these routes in advance, and then build tables that define the routes in each router. Each packet carries a label that indicates which exact route it should follow. In addition to supporting paths determined by OSPF and BGP, MPLS supports paths defined by a variety of constraints, including, but not limited to available bandwidth,

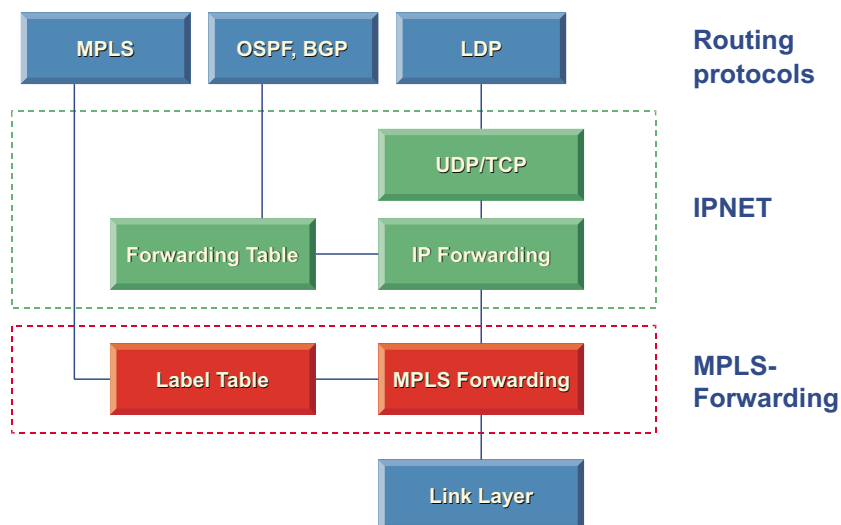
packets' priority settings, and the policy-based server's directives. In this way, MPLS supports QoS by constraint-based routing.

Ingress routers at the edge of the MPLS network use a packet's destination address to determine the appropriate LSP. The ingress routers label each packet with the selected LSP. MPLS routers rely exclusively on these labels for packet forwarding. Prior to forwarding, the label may be replaced by another label for use at the next hop router. At the egress router, the MPLS label is removed before packets are forwarded as standard IP traffic.

When an LSP fails, the router upstream from the failure signals the outage to the ingress router. The ingress

router establishes a new LSP and directs traffic to the new path. This rerouting process can be time-consuming and prone to failure. Fast re-routing enables the upstream router to route rapidly around the failure to the downstream router. The upstream router then signals the outage to the ingress router, thereby maintaining connectivity until a new LSP is established.

The Label Distribution Protocol (LDP) establishes MPLS LSPs. Using LDP, MPLS routers exchange label assignments to create LSPs. LDP allows MPLS routers to develop a labeling methodology, so each device is able to use the appropriate label to forward traffic to and through another device.



Interpeak MPLS block diagram.

Interpeak MPLS Features

Interpeak MPLS Forwarding

The forwarding plane may assume the role of an ingress, transit, or egress node. It supports label tagging and stacking, label lookups, and forwarding of labeled packets.

Ingress nodes prepare packets for entry into MPLS tunnels by labeling and sending the packet to an MPLS tunnel interface, where it is forwarded to the next hop in the label's LSP.

Transit nodes send labeled packets from incoming to outgoing interfaces, using information stored in Incoming Label Maps, implemented as Radix trees for fast lookups.

Egress nodes remove the label and pass the packet to the TCP/IP stack for further processing. The Network Layer destination of the packet is either pre-defined or explicitly set by a label operation. If pre-defined, the packet is either terminated in the TCP/IP stack or forwarded. If an explicit destination is used, it acts as the gateway address for further forwarding in the network layer.

All label operations can be added or removed dynamically. This allows for dynamic configuration of label handling without disruption of other networking traffic.

Integration with IPNET

Carriers are increasingly building Virtual Private Networks (VPNs) using the guidelines in the IETF's RFC 2547, which require routers to segregate routing information by VPN. In particular, organizations with multiple locations use MPLS and BGP to build VPNs over an IP backbone. MPLS forwards the packets over the backbone between the sites. BGP extracts the in-

MPLS integrated directly into IPNET supports:

- **MPLS over Ethernet devices**
- **Generic MPLS labels**
- **IPv4 and IPv6 over MPLS**
- **Label stacking**
- **Common label operations**
- **Delivery of PDUs to IPNET**
- **Ingress, transit or egress node**

Interpeak MPLS features.

ternal routes of each site's network and distributes this routing information to all VPN sites.

In such VPNs, the provider edge routers segregate routing information by VPN and maintain a forwarding table for each attached VPN. Each connected network no longer needs its own router, but can have its own secure forwarding table maintained by the TCP/IP stack. Called virtual routing, this system presents significant cost savings because many physical routers can be eliminated.

IPNET, Interpeak's TCP/IP stack, can support the extensive requirements

of large-scale RFC 2547 VPN deployments: the number of forwarding tables is virtually unlimited and the tables support private overlapping addresses.

IPv6

RFC 2547 guidelines may be applied to both IPv4 and IPv6 MPLS deployments. As carriers build more IPv6 networks, developers can continue to use the dual-mode IPv4/IPv6 stack. IPNET is a future-proof routing stack, a true dual-mode IPv4/IPv6 stack which supports numerous IPv4/IPv4 transition mechanisms.

- **MPLS software data plane**
- **Pre-integrated with commercial MPLS control plane products**
- **Runs on Linux, VxWorks, INTEGRITY, OSE, etc.**

Interpeak Secure Networking Software

Interpeak provides state-of-the-art networking solutions specifically designed for embedded systems. The company's embedded networking and security software is currently used in thousands of applications across the globe.

Headquartered in Stockholm, Sweden, Interpeak operates through a global network of distribution channels and has its own sales and field application force dispersed in strategic locations worldwide, including the USA, Europe, and Asia. For additional information, please visit our homepage www.interpeak.com.

All Interpeak products are trademarks or registered trademarks of Interpeak AB. Other brand and product names are trademarks or registered trademarks of their respective holders. The information in this document has been carefully reviewed, and is believed to be accurate and reliable. However, Interpeak AB assumes no liabilities for inaccuracies in this document. Furthermore, Interpeak AB reserves the right to change specifications embodied in this document without prior notice.

Version 1.02-r5. Copyright © 2005, Interpeak AB. All rights reserved.