# 136
TECHNOLOGY
WHITE PAPER

The Role of Signaling in MPLS

sig·nal[1], *noun.* **1.** a sign or event fixed or understood as the occasion for prearranged combined action; as, a signal for a fire drill. **2.** a sign given by gesture, mechanical device, etc. to convey command, warning, or other information; as, a red light is a stop signal. *verb.* **1.** to make known or communicate (information) by signals; as, the vessel signaled its arrival.

Gary Holland, Riverstone Networks

ABSTRACT

In order to understand the role of signaling in network communications, and its use in Multi-Protocol Label Switching (MPLS), it helps to first step back and broadly consider what signaling is. In general terms, signaling is the communication of information using pre-defined symbols. Humans have used signaling throughout history, conveying specific information through means ranging from drum beats and smoke signals to colored flags and lights.

All of us encounter signals in our everyday lives. For drivers, the familiar red, yellow, and green traffic signals control the flow of vehicles along the roadways. Signals are also used in specialized contexts. For example, a sea captain can use a variety of signals to communicate that a ship is in distress, including the well-known Morse code S.O.S. as well as signal flags, which indicate conditions ranging from a person overboard to a disabled ship. Likewise, for the weekend sailor, the color and position of lights on a vessel indicate whether that vessel is under power, sailing, or at anchor.

Other captains can use this signaling information to respond accordingly – to come to the aid of the ship in distress or, in the case of the weekend sailor, to steer clear of a vessel at anchor. Signaling, then, is the process of using signals to communicate information, typically with the intent of affecting some action. Let's see how it has been used in a networking context.

SIGNALING IN NETWORK COMMUNICATION

The use of Morse code in conjunction with the telegraph can be viewed as an early form of network communication signaling. For most of us, the network communication system that the term "signaling" is most closely associated with is the telephone system. Signaling is used widely throughout the telephone system to control everything from telephone handset functions to switch-to-switch interactions.

For example, signaling is used to set up, maintain, and tear down each phone call. From the caller's perspective, signaling provides everything from dial tone to the ringing of the called person's phone (or a busy signal in those homes with teenagers). Signaling is used in switch-to-switch communication for functions such as defining and assigning trunk paths and adapting to the current state of the system, including routing around failures.

[1] *"Webster's New Twentieth Century Dictionary, Unabridged;" Second Edition*

Early on, telephone signaling was based on a voltage system, where conditions such as a phone being on-hook or off-hook were a function of the current on the line. Similarly, a phone rang when voltage was applied to the phone line. Over time, this method of signaling gave way to bit-based signaling, whereby either two or four bits were used to encode signaling information. However, the growing sophistication of the telephone system outstripped the capabilities of the bit-based system, which was limited in the number of functions it could support.

While bit-based signaling is still used in some parts of the telephone system, it has largely been superseded by message-based digital signaling. As the name implies, message-based signaling relies upon the use of messages to convey signaling information. A message-based system can be easily expanded to accommodate new functions simply by creating new messages. Signaling System 7 (SS7), which is widely deployed throughout the telephone system for communication among switches, is a message-based signaling mechanism. SS7 provides signaling for circuit-related services, such as call control, as well as non-circuit related services, including billing services.

Asynchronous Transfer Mode (ATM) is another well-known communications technology that relies on message-based signaling for its operation. For example, an ATM device uses signaling to establish a connection with another ATM device. The signaling packet contains the ATM address of the desired ATM endpoint and may also contain Quality of Service (QoS) parameters required for the connection. Every switch along the path to the destination examines the signaling packet and determines if it can support the requested traffic parameters on both its ingress and egress interfaces. If it can support these parameters, the switch forwards the signaling packet to the next switch and also sets up a virtual connection as the signaling packet is forwarded.

If any switch along the path cannot accommodate the requested traffic contract, the request is rejected and a rejection message is sent back to the initiating ATM device. However, if all switches along the forwarding path can support the requested QoS and the endpoint can as well, the endpoint responds with an accept message and the connection is opened.

Clearly, the types of signaling messages used will vary from one network communication system to another. Likewise, the functions that signals support will vary. In connection-oriented networks, for example, signaling is generally used to establish, monitor, and tear down connections. Overall, signals fall into a handful of basic categories and include supervisory signals, information signals, address signals, control signals, and alerting signals.

These signals can be put to use for a variety of purposes, such as admission control (which traffic is allowed onto the network), resource allocation (how much bandwidth or other network resources a particular traffic flow is entitled to), and accounting. And with message-based mechanisms, signaling can also be used to enable new services through the definition of new messages.

SIGNALING IN AN IP NETWORK

Historically, the Internet Protocol (IP) has not relied on signaling for its operation. Designed to permit communication among different types of underlying networks, IP is a protocol that defines an unreliable, connectionless delivery mechanism.[2]

IP encompasses three functions: 1) it defines the basic unit of data transfer across a TCP/IP network; 2) IP software performs the routing function to choose a path over which data will be sent; and 3) it includes a set of rules regarding unreliable packet delivery, including how packets should be processed, errors handled, and packets discarded. In essence, IP is responsible for sending and receiving data, nothing more. IP does not encompass signaling, per se.

[2] *Page 91 "Internetworking with TCP/IP" by Douglas E. Comer, Volume 1, Third Edition*

The Transmission Control Protocol (TCP), on the other hand, has some characteristics that could be construed as forms of signaling. TCP was designed to accommodate large transfers of data, or streams of data, between two application programs. As a transport protocol, TCP's main function is to provide reliable data exchange; that is, it ensures that data is delivered error free, in order, with no loss or duplication.

TCP accomplishes this task by establishing connections between application processes, identified with a port, on the sending and receiving computers. Many of TCP's connection management functions – such as establishing a connection, maintaining that connection for the transport of data, and terminating the connection – are communication activities that we typically associate with signaling protocols.

TCP not only establishes a logical connection between a pair of application processes, it also provides services over this connection. For example, as part of establishing a connection, TCP allows security and priority characteristics to be specified for that connection. It also provides flow control to regulate the flow of data across a connection. This ability to specify the parameters associated with a communication exchange is another characteristic of signaling. So while TCP may not be generally referred to as a signaling protocol, we can see that it can be viewed as one.

In contrast, the User Datagram Protocol (UDP) adds little to IP other than the concept of ports, which enables one application to send data to another application program. UDP uses IP to carry messages between computers, and so provides an unreliable, connectionless delivery service.

One IP-related protocol that is generally accepted as a signaling protocol is the Resource Reservation Protocol (RSVP). RSVP is a signaling protocol that was specifically designed to obtain QoS services from an IP network. The Internet Engineering Task Force (IETF) initially defined RSVP as the signaling mechanism for its Integrated Services (IntServ) QoS model. It has since evolved for other uses, including to signal explicit route setup in an MPLS environment. Before we delve more deeply into MPLS, however, it's useful to understand what RSVP is (and isn't) and how it operates.

OVERVIEW OF RSVP SIGNALING

The IETF designed the IntServ framework to provide end-to-end flow-based QoS to applications. IntServ includes a set of QoS service definitions – what we might think of as "gold," "silver," and "bronze" services, for example. It also outlines the mechanisms that network elements need, such as queuing schemes and buffering, to support these services. Lastly, IntServ spells out that applications need a way to communicate their QoS requirements to network elements.

In an IntServ/RSVP environment, RSVP is the mechanism that applications use to signal their resource requirements to network devices. RSVP is a message-based signaling protocol and can carry a variety of information in its messages. In an IntServ context, RSVP relies on IntServ-specific parameters to convey QoS information. In addition, RSVP can carry the authentication, accounting, and policy information needed to manage these QoS services.

However, as we noted earlier, message-based signaling mechanisms can be used for various applications simply by creating new messages. As a message-based signaling protocol, RSVP is not limited to carrying QoS-related messages. Rather, RSVP messages can carry different information to meet different application requirements. This "generic" aspect of RSVP has enabled it to be used for applications other than QoS signaling, such as signaling MPLS explicit routes. Because of RSVP's role in MPLS signaling, it's worthwhile to take a high-level look at how RSVP works in a QoS context.

RSVP signaling consists of a series of interactions between senders, receivers, and the network devices in between. Applications that want to request a specific level of service from a network can use RSVP to do so. Since the QoS parameters that an application may request need to match those of the receiver, it is the receiver that spells out what bandwidth, latency, and other QoS characteristics are needed during a particular data exchange. It's easy to see why RSVP is receiver oriented; if Computer A, which has a 100 Mbps Ethernet interface, wishes to send a stream of data to Computer B, which has only a 10 Mbps Ethernet interface, it would make little sense for Computer A to be in control of the resource reservation and request 100 Mbps for the communication exchange with Computer B.

Although RSVP is receiver oriented, the sending computer initiates an RSVP reservation process by transmitting a particular kind of RSVP message known as a PATH message. The PATH message is responsible for finding a path through the network for a specific data stream (also called a flow), and to bind this route for RSVP's use. PATH messages contain information about the data flow's source as well as characteristics about the traffic it wants to send. Network devices also use PATH messages to discover nearby RSVP-enabled devices and to advertise their QoS capabilities.

Once a PATH message arrives at the receiving computer, that computer generates reservation (RESV) messages that indicate what kind of QoS services the receiver needs. These RESV messages travel hop-by-hop back along the route established by the PATH messages. The network device at each hop determines whether it can provide the requested QoS. If it can, it allocates the appropriate resources. The RESV message continues along its reverse path until it arrives at the sending computer. At this point, the reservation process is complete, and the sending station can begin transmitting data along the reserved path.

To recap, RSVP is a signaling protocol that provides a way for applications to signal their service requirements to all the devices in a network that will handle the data flow between a pair of sending and receiving applications. RSVP messages basically contain a header and "objects" that inform each device along the communication path about a special type of processing that's required. In a QoS context, RSVP messages carry QoS-related objects; devices along the RSVP path act upon this QoS-related information by reserving buffer space, etc. By defining new objects, RSVP signaling can be used for purposes other than establishing QoS reservations, as we will see in our discussion of MPLS signaling.

SIGNALING IN AN MPLS CONTEXT

In MPLS, packets are forwarded based on a label rather than the IP address, as in traditional routing. At each hop along the forwarding path, an MPLS node uses the label to make forwarding decisions for a packet. Remember that MPLS is a label switching or swapping scheme. Each MPLS node maintains a table of label information; MPLS nodes use this table to look up the label on a packet coming into the node so they know which label to apply to the packet before forwarding it on to the next MPLS node. Each MPLS node removes the incoming label and appends an outgoing label.

Like TCP, MPLS deals with streams of data, and a key aspect of MPLS is that packets with the same label are forwarded the same way. For packets forwarded in a hop-by-hop fashion, for example, this means that all packets with the same label are forwarded to the same next hop. In this way, a Label Switched Path (LSP) is established.

LSPs are basically a concatenation of one or more label switched hops. They are extended through a network as each MPLS node swaps the incoming label for the outgoing label assigned to the next hop for that data stream. In essence, an LSP is established when each MPLS node along the path between the initial (ingress) MPLS node and the final (egress) MPLS node has a binding between an incoming label and an outgoing label.

On the one hand, forwarding with labels requires relatively simple functions – looking up a label in a table, swapping labels, and perhaps checking a Time-to-Live field (TTL). On the other hand, a control component is needed for critical functions, such as creating the bindings between labels and routes, distributing label information to MPLS nodes so they know which label to apply to a given stream of data, and withdrawing labels.

MPLS has a clear separation between data forwarding and the control mechanisms used for routing, label management, and other control functions. That is, packets are forwarded based on label switching, regardless of the underlying control mechanism. As a result, a variety of control mechanisms can be used with MPLS; signaling is a key aspect of these control mechanisms.

Specifically, MPLS supports a variety of signaling mechanisms for label distribution and LSP set up. To date, the mechanisms the IETF has defined for basic label distribution is the Label Distribution Protocol (LDP). The IETF has also defined enhancements to LDP for Constraint Route signaling (CR-LDP) and extensions to basic RSVP for Traffic Engineering (RSVP-TE). Both of these signaling protocols build on existing protocols to allow explicit set-up and traffic engineering of LSPs in an MPLS network.

THE LABEL DISTRIBUTION PROTOCOL

LDP comes into play primarily for best-effort forwarding and is the signaling protocol designed for setting up LSPs on a hop-by-hop basis. LDP encompasses two major functions. First, it is a mechanism by which MPLS nodes can exchange information about label mappings. Second, it defines a set of procedures and messages that allow MPLS nodes to establish LSPs through a network that match hop-by-hop routed paths. It is this second function, more so than the first, that clearly delineates LDP as a signaling protocol.

Like all modern signaling protocols, LDP is message based. LDP defines a number of messages, including session messages that are used to establish, maintain, and terminate sessions between MPLS nodes. It also uses advertisement messages in order to create, change, and delete label mappings. In addition, it has notification messages to provide advisory and error information. All of these messages are transmitted over TCP connections. The only messages that LDP sends as UDP packets are discovery messages, which MPLS nodes use to announce their presence on the network.

As we noted above, LDP establishes LSPs on a hop-by-hop basis. That is, each MPLS node independently selects the next hop for a given stream of data based on the information in its label information table. However, the path that an LSP follows can also be explicitly defined based on some special handling requirement or constraint, such as the need for a path with a given QoS or for paths that allow the network operator to balance traffic across a network.

Network operators may want to use explicit routing for a variety of reasons. For example, explicit routing can be used for traffic engineering, enabling network operators to select paths for traffic with an eye toward balancing the load on links, routers, and switches across a network. Network operators also may want to apply resource reservations to explicitly routed paths to ensure that a particular traffic type, such as voice, gets the handling it needs as it traverses the network. Specific signaling protocols have been defined to establish explicitly routed LSPs. Let’s examine them more closely.

SIGNALING MECHANISMS FOR EXPLICIT ROUTING

The IETF has defined two signaling schemes for establishing explicitly routed or constraint-based LSPs in an MPLS environment. One relies on the use of RSVP, while the other approach is based on extensions to LDP. Both schemes allow for the setup of explicitly routed LSPs and for signaling of a set of parameters, such as bandwidth constraints, relating to those LSPs.

In RFC 3209, the IETF has defined extensions to RSVP that allow it to be used to establish explicitly-routed LSPs. Since support for traffic engineering was a key requirement that drove many of the new features in RSVP, this enhanced version of the protocol is referred to as RSVP-TE. RFC 3209 defines procedures for using RSVP messages to allocate and bind labels, to distribute these label bindings, and to establish LSPs.

RSVP-TE uses many of the same message types defined for RSVP, but with newly-defined objects. For example, labels are distributed using a special LABEL object in RESV messages. Similarly, labels are bound to specific LSP tunnels through the use of the LABEL_REQUEST object in PATH messages. An EXPLICIT-ROUTE object in PATH messages is used to set up explicit routes, and so on.

RSVP-TE encompasses a range of MPLS-related functions, including the capability to establish LSP tunnels with or without QoS reservations, reroute around an established LSP tunnel, and preempt an established LSP tunnel under administrative policy control. Because RSVP was originally designed to signal QoS requirements in an IP network, RSVP-TE is easily used to support QoS-based traffic forwarding. The RSVP-TE "PATH" message includes a field that encompasses the QoS parameters for bandwidth, burst limits, delay, jitter, etc. for each link along the requested LSP.

The second signaling protocol that the IETF has defined for explicit routing is based on extensions to LDP. The IETF has defined these extensions in an Internet draft titled "Constraint-Based LSP Setup Using LDP," and refers to the signaling protocol as CR-LDP (for constraint-based routing). Functionally, CR-LDP is quite similar to RSVP-TE: It allows for the specification of traffic parameters relating to an explicitly routed LSP and supports the preemption of existing LSPs to accommodate the resource requirements of a new LSP. As with RSVP-TE, CR-LDP relies on the definition of new objects that are carried in LDP messages. As in the RSVP-TE "PATH" message, the CR-LDP "LABEL REQUEST" message has fields defined for the explicit route and specific bandwidth parameter requests for each link along the requested LSP.

It is through signaling mechanisms that establish LSPs that MPLS achieves its ability to bring "connection orientedness" to IP. And because LDP, RSVP-TE, and CR-LDP are message-based signaling mechanisms, it's easy to see how they are extended to support additional constraints or parameters in establishing LSPs.

CONCLUSION

Clearly, signaling has played a significant role in human communication throughout history, and is a key tool for communications specialists in extending networking functionality. In an MPLS context, signaling enables MPLS nodes to exchange the information they need to establish LSPs, and to communicate the type of processing that a packet requires as it traverses an LSP.

For further reading on "Approaches to Signaling in MPLS," see:
www.riverstonenet.com/support/mpls/approaches_to_signaling_in_mpls.htm

ACRONYMS

10GbE:	10-Gigabit Ethernet
ABR:	Available Bit Rate
ACL:	Access Control List
ADM:	Add Drop Muxes
ANSI:	American National Standards Institute
API:	Application Program Interface
APS:	Automatic Protection Switching
ARP:	Address Resolution Protocol
ASIC:	Application-Specific Integrated Circuit
ASP:	Application Service Provider
ATM:	Asynchronous Transfer Mode
BGP:	Border Gateway Protocol
BOND:	Bandwidth on Demand
BPDU:	Bridge Protocol Data Units
CAGR:	Compound Annual Growth Rate
CAPEX:	Capital Expenditure
CAR:	Committed Access Rate
CBR:	Constant Bit Rate
CIR:	Committed Information Rate
CLEC:	Competitive Local Exchange Carrier (ie: MCI, Sprint, etc.)
CLI:	Command Line Interface
CMTS:	Cable Modem Termination System
CO:	Central Office
CORBA:	Common Object Request Broker Architecture
CoS:	Class of Service
CPE:	Customer Premise Equipment
CR-LDP:	Constraint-Based LDP
CSMA/CD:	Carrier Sense Multiple Access/ Collision Detection
CSP:	Content Service Provider
CSPF:	Constraint-based Shortest Path First
DHCP:	Dynamic Host Configuration Protocol
DiffServe:	Differential Services IETF Standard
DLL:	Data Link Layer
DOCSIS:	Data Over Cable System Interface Specification
DS1/DS3:	Digital Signal, Level 1 (1.54 Mbps) or 3 (44.7 Mbps)
DSCP:	DiffServ Code Point
DSL:	Digital Subscriber Line
DSLAM:	DSL Access Multiplexer
DXC:	Digital Cross Connects
E1/E2:	European Trunk 1/2 (2 Mbps/34.3 Mbps)
EBITDA:	Earnings Before Interest, Taxes, Depreciation, and Amortization
ECTA:	European Communities Trademark Association
EFM:	Ethernet in the First Mile
EMS:	Element Management System
ERP:	Enterprise Resource Planning
ESP:	Ethernet Service Provider
EtherLEC (or ELEC):	Ethernet Local Exchange Carriers
EXP:	Experimental bits
FCAPS:	Fault, Configuration, Accounting, and Security
FDDI:	Fiber Distributed Data Interface
FEC:	Forwarding Equivalence Class
FRAD:	Frame Relay access device
GARP:	Generic Attribute Register Protocol
GMPLS:	Generalized MPLS
GVRP:	GARP VLAN
HFC:	High Frequency Cable
HPR:	Hitless Protocol Restart
HPS:	Hitless Protection System
HSSI:	High Speed Serial Interface
IANA:	Internet Assigned Numbers Authority
IBGP:	Internal Border Gateway Protocol
IEEE:	Institute of Electrical and Electronic Engineers
IETF:	Internet Engineering Task Force
IGP:	Interior Gateway Protocol
Ion SDK:	Ion Software Developer's Kit
IP:	Internet Protocol
IS-IS:	Intermediate Systems to Intermediate Systems
ISP:	Internet Service Provider
ITU:	International Telecommunications Union
IXC:	Inter-exchange Carrier
LAN:	Local Area Network
LDP:	Label Distribution Protocol
LEC:	Local Exchange Carrier
LER:	Label Edge Router
LFAP:	Lightweight Flow Accounting Protocol
LSP:	Label Switched Path
LSR:	Label Switched Router
MAC:	Media Access Control
MAN:	Metropolitan Area Network
MCDN:	Microcellular Data Network
MDI:	Media Dependent Interface
MDU:	Multiple Dwelling Unit
MIB:	Management Information Base
MLPPP:	Multi-Layer Point-to-Point Protocol
MMDS:	Multi-point Multi-channel Distribution System
MMF:	Multi-mode fiber
MPLS:	Multi-Protocol Label Switching
MSN:	Management Network Server
MSP:	Metropolitan Service Provider
MSTP:	Multiple Spanning Tree Protocol
MTBF:	Mean Time Between Failure
MTU:	Multiple Tenant Unit
MVST:	Multiple-VLAN Spanning Tree
NAT:	Network Address Translation
NEBS:	Network Equipment Building Systems
NMS:	Network Management System
OAM:	Operations, Administration, and Maintenance
OC-3/OC-12:	Optical Carrier 3/12 (155 Mbps/622 Mbps)
OPEX:	Operational Expenditure
OSI:	Open System Interconnection
OSPF:	Open Shortest Path First
OSS:	Operation and System Support or Operational Support Systems
PAT:	Port Address Translation
PCS:	Physical Coding Sublayer
PDU:	Protocol Data Units
PE:	Provider Edge
PEF:	Packet over Ethernet over Fiber
PES:	Packet over Ethernet over SONET-based optical
PEW:	Packet over Ethernet over WDM-based optical
PHY:	Physical Layer
PMA:	Physical Media Attachment
PMD:	Physical Media Dependent
PON:	Passive Optical Networking
POP:	Point of Presence
POS:	Packet over SONET
PPP:	Point to Point Protocol
PSTN:	Public Switch Telephone Network
PVC:	Private Virtual Circuit
PVST:	Per-VLAN Spanning Tree
QoS:	Quality of Service
RBOC:	Regional Bell Operating Company (ie: PacBell, etc.)
RED:	Random Early Discard
RFC:	Request for Comments
RMON:	Remote Monitoring
RPR:	Resilient Packet Ring
RRST:	Rapid Ring Spanning Tree
RSTP:	Rapid Spanning Tree Protocol
RSVP:	Resource Reservation Protocol
RSVP-TE:	RSVP-Traffic Engineering
rt-VBR:	real-time Variable Bit Rate
SAN:	Storage Area Network
SAP:	Session Announcement Protocol
SFP:	Short Formfactor Pluggable
SLA:	Service Level Agreement
SMF:	Single-mode fiber
SNMP:	Simple Network Management Protocol
SPoF:	Single Points of Failure
SONET:	Synchronous Optical Network
SONET/SDH:	Synchronous Optical Network/ Synchronous Digital Hierarchy
SPoF:	Single Points of Failure
SRP:	Spatial Reuse Protocol
STP:	Spanning Tree Protocol
T1:	Trunk 1 (1.544 Mbps)
TCO:	Total Cost of Ownership
TCP/IP:	Transmission Control Protocol/Internet Protocol
TDM:	Time Division Multiplexing
TLS:	Transparent LAN Service
TNM:	Telecommunication Network Management
TOS:	Terms of Service, or Type of Service
TTL:	Time-to-Live
UBR:	Undefined Bit Rate
UDP:	User Datagram Protocol
VAS:	Value-added Services
VBR:	Variable Bit Rate
VLAN:	Virtual LAN
VoD:	Video on Demand
VoIP:	Voice over IP
VPLS:	Virtual Private LAN Services
VPN:	Virtual Private Network
VRC:	Virtual Router Cluster
VRRP:	Virtual Router Redundancy Protocol
WAN:	Wide Area Network
WDM:	Wave Division Multiplexing
WFQ:	Weighted Fair Queuing
WRED:	Weighted Random Early Discard
XGMII:	10G Media Independent Interface

Riverstone Networks, Inc.
5200 Great America Parkway, Santa Clara, CA 95054 USA

877 / 778-9595 or **408 / 878-6500** or **www.riverstonenet.com**

Printed in the USA

v 1.3 4/02