Analysis of the Similarities and Differences between MPLS Label Distribution Protocols RSVP and CR-LDP

Muhammad Asif[1], Zahid Farid[2], Muhammad Lal[3], Junaid Qayyum[4]

[1]Department of Information Technology and Media (ITM), Mid Sweden University
Sundsvall 85170, Sweden.

[2]School of Electrical and Electronics Engineering, Universiti Sains Malaysia (USM)
Penang 14300, Malaysia.

[3, 4] Department of Information Technology, Gandhara University of Sciences
Peshawar 25000, Pakistan.

Abstract

MPLS is a new technology that offers to open up the Internet by providing many additional services to applications using IP. MPLS forwards data using labels that are attached to each data packet. These labels must be distributed between the nodes that comprise the network. Many of the new services that ISPs want to offer rely on Traffic Engineering functions. There are currently two label distribution protocols that provide support for Traffic Engineering: Resource ReSerVation Protocol (RSVP) and Constraint-based Routed Label Distribution Protocol (CR-LDP). Although the two protocols provide a similar level of service, the way they operate is different, and the detailed function they offer is also not consistent. Hardware vendors and network providers need clear information to help them decide which protocol to implement in a Traffic Engineered MPLS network. Each protocol has its champions and detractors, and the specifications are still under development. Recognizing that the choice of label distribution protocol is crucial for the success of device manufacturers and network providers; this White Paper explains the similarities and important differences between the two protocols, to help identify which protocol is the right one to use in a particular environment. Data Connection's DC-MPLS family of portable MPLS products offers solutions for both the RSVP and CR-LDP label distribution protocols.

Keywords MPLS, ISPs, Traffic Engineering, Resource ReSerVation Protocol (RSVP), Constraint-based Routed Label Distribution Protocol (CR-LDP), Data Connection's DC-MPLS.

1. Introduction

CR-LDP is a set of extensions to LDP specifically designed to facilitate constraint-based routing of LSPs. Like LDP, it uses TCP sessions between LSR peers and sends label distribution messages along the sessions. This allows it to assume reliable distribution of control messages. [1] The basic flow for LSP setup using CR-LDP is as shown in Figure 1.

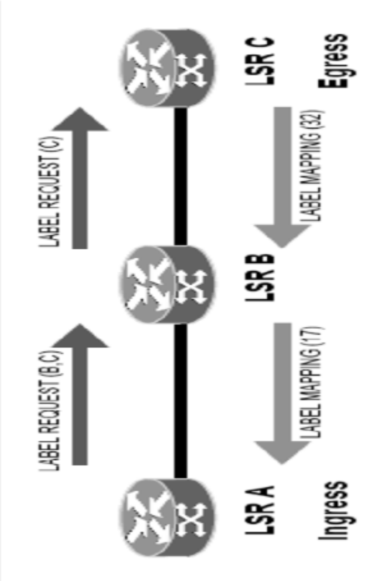

Figure 1 (CR-LDP LSP Setup Flow)

- The Ingress LSR, LSR A, determines that it needs to set up a new LSP to LSR C. The traffic parameters required for the session or administrative policies for the network enable LSR A to determine that the route for the new LSP should go through LSR B, which might not be the same as the hop-by-hop route to LSR C. LSR A builds a LABEL_REQUEST message with an explicit route of (B,C) and details of the traffic parameters requested for the new route. LSR A reserves the resources it needs for the new LSP, and then forwards the LABEL_REQUEST to LSR B on the TCP session.

- LSR B receives the LABEL_REQUEST message, determines that it is not the egress for this LSP, and forwards the request along the route specified in the message. It reserves the resources requested for the new LSP, modifies the explicit route in the LABEL_REQUEST message, and passes the message to LSR C. If necessary, LSR B may reduce the reservation it makes for the new LSP if the appropriate parameters were marked as negotiable in the LABEL_REQUEST.

- LSR C determines that it is the egress for this new LSP. It performs any final negotiation on the resources, and makes the reservation for the LSP. It allocates a label to the new LSP and distributes the label to LSR B in a LABEL_MAPPING message, which contains details of the final traffic parameters reserved for the LSP.

- LSR B receives the LABEL_MAPPING and matches it to the original request using the LSP ID contained in both the LABEL_REQUEST and LABEL_MAPPING messages. It finalizes the processing at LSR A is similar, but it does not have to allocate a label and forward it to an upstream LSR because it is the ingress LSR for the new LSP.

Generic RSVP uses a message exchange to reserve resources across a network for IP flows. The Extensions to RSVP for LSP Tunnels enhances generic RSVP so that it can be used to distribute MPLS labels. RSVP is a separate protocol at the IP level. It uses IP datagrams (or UDP at the margins of the network) to communicate between LSR peers. It does not require the maintenance of TCP sessions, but as a consequence of this it must handle the loss of control messages. [2] The basic flow for setting up an LSP using RSVP for LSP Tunnels is shown in Fig.2 below.

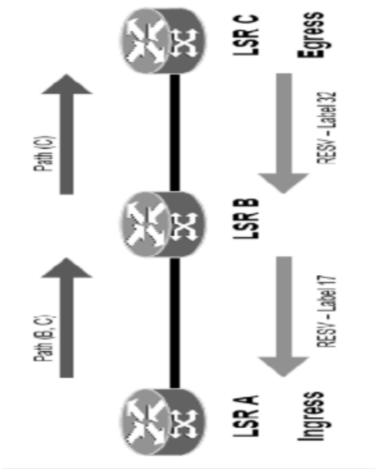

Figure 2 (RSVP LSP Setup Flow)

- The Ingress LSR, LSR A, determines that it needs to set up a new LSP to LSR C. The traffic parameters required for the session or administrative policies for the network enable LSR A to determine that the route for the new LSP should go through LSR B, which might not be the same as the hop-by-hop route to LSR C. LSR A builds a Path message with an explicit route of (B,C) and details of the traffic parameters requested for the new route. LSR A then forwards the Path to LSR B as an IP datagram.

- LSR B receives the Path request, determines that it is not the egress for this LSP, and forwards the request along the route specified in the request. It modifies the explicit route in the Path message and passes the message to LSR C.

- LSR C determines that it is the egress for this new LSP, determines from the requested traffic parameters what bandwidth it needs to reserve and allocates the resources required. It selects a label for the new LSP and distributes the label to LSR B in a Resv message, which also contains actual details of the reservation required for the LSP.

- LSR B receives the Resv message and matches it to the original request using the LSP ID contained in both the Path and Resv messages. It determines what resources to reserve from the details in the Resv message, allocates a label for the LSP, sets up the forwarding table, and passes the new label to LSR A in a Resv message.

- The processing at LSR A is similar, but it does not have to allocate a new label and forward this to an upstream LSR because it is the ingress LSR for the new LSP.

2. Comparative Analysis

The key differences between CR-LDP and RSVP are the reliability of the underlying transport protocol and whether the resource reservations are done in the forward or reverse direction. From these points come many of the other functional differences. [3] The table below summarizes the main technical similarities and differences between CR-LDP and RSVP for LSP Tunnels. The sections that follow explain in greater detail the implications of these technical differences between the protocols.

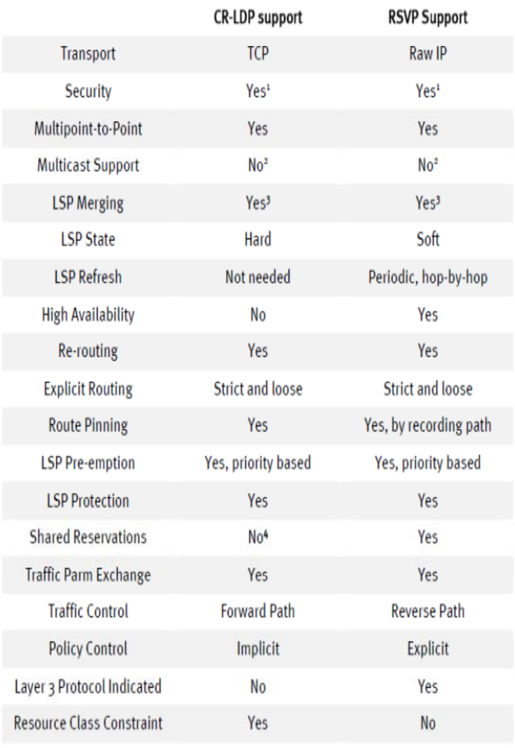

	CR-LDP support	RSVP Support
Transport	TCP	Raw IP
Security	Yes[1]	Yes[1]
Multipoint-to-Point	Yes	Yes
Multicast Support	No[2]	No[2]
LSP Merging	Yes[3]	Yes[3]
LSP State	Hard	Soft
LSP Refresh	Not needed	Periodic, hop-by-hop
High Availability	No	Yes
Re-routing	Yes	Yes
Explicit Routing	Strict and loose	Strict and loose
Route Pinning	Yes	Yes, by recording path
LSP Pre-emption	Yes, priority based	Yes, priority based
LSP Protection	Yes	Yes
Shared Reservations	No[4]	Yes
Traffic Parm Exchange	Yes	Yes
Traffic Control	Forward Path	Reverse Path
Policy Control	Implicit	Explicit
Layer 3 Protocol Indicated	No	Yes
Resource Class Constraint	Yes	No

Notes:

1. CR-LDP inherits any security applied to TCP. RSVP cannot use IPSEC but has its own authentication. See “Security” below.

2. Multicast support is currently not defined for any of the existing label distribution protocols.

3. CR-LDP does not allow explicit sharing, but see “LSP Modification” below for details of changing the allocated resources.

2.1. Availability of Transport Protocol

The most obvious difference between CR-LDP and RSVP is the choice of transport protocol used to distribute the label requests. RSVP uses connectionless raw IP (or UDP packets at the margins of the network). CR-LDP uses UDP to discover MPLS peers, and uses connection-oriented TCP sessions to distribute label requests. Many operating systems are packaged with full IP stacks including UDP and TCP, but sometimes TCP is not available. On some platforms access to raw IP is restricted. [4] Some existing ATM switches might not already incorporate an IP stack at all and one must be added to support either CR-LDP or RSVP. The availability and accessibility of the transport protocols may dictate which label distribution protocol is used, but is unlikely to be a major factor in the choice made by most MPLS equipment suppliers. RSVP requires that all received IP packets carrying RSVP messages are delivered to the RSVP protocol code without reference to the actual destination IP address in the packet. This feature may require a minor modification to the IP implementation.

2.2. Security

TCP is vulnerable to denial of service attacks, where the performance of the TCP session can be seriously impacted by unauthorized access to the network. This could impact CR-LDP. Authentication and policy control are specified for RSVP. This allows the originator of the messages to be verified (for example using MD5) and makes it possible to police unauthorized or malicious reservation of resources. Similar features could be defined for CR-LDP but the connection-oriented nature of the TCP session makes this less of a requirement. TCP itself could make use of MD5. IPSEC is a series of drafts from the IETF to provide authentication and encryption security for packets transported over IP. If IPSEC support is available in the IP stack it can be used by CR-LDP simply as part of the normal TCP/IP processing. RSVP targets its Path messages at the egress LSR, not at the intermediate LSRs. This means that IPSEC cannot be used because the intermediate LSRs would find themselves unable to access the information in the Path messages.

2.3. Multipoint Support

Multipoint-to-point LSPs allow label switched paths to merge at intermediate LSRs, reducing the number of labels required in the system and sharing downstream resources. This approach works particularly well in packet-switched networks, but requires non-standard hardware in cell-switched networks such as ATM to prevent interleaving of cells. CR-LDP and RSVP support multipoint-to-point LSPs. [5] Point-to-multipoint (multicast) IP traffic is not addressed by the current version of the MPLS Architecture, so it is not supported by CR-LDP or Labels RSVP. Generic RSVP was originally designed to include resource reservation for IP multicast trees, so it may be easier to extend to support multicast traffic in the future. However, this is an area for further study in both protocols.

2.4. Scalability

The scalability of a protocol should be considered in terms of the network flows it uses, the resources needed to maintain the protocol state at each node, and the CPU load on each node. All of this must be considered in the context of the way in which MPLS is to be used in the network. If trunk LSPs are to be used across the network to connect key edge points, there will be less demand on scalability than using one LSP per flow, or setting up LSPs based on the routing topology. [6] The ability to merge LSPs also has a clear impact on scalability requirements, because data flows may be able to share resource allocations, and the number of labels needed in the network is reduced.

2.5. High Availability

- Availability is a measure of the percentage of time that a node is in service. Equipment vendors typically claim high availability for their boxes when they attain availability levels in the region of 99.999% ("5-nines").

- High Availability is a matter of detecting failures and handling them in a timely manner without any – or with only minimal – disruption to service. Detection and survival of link failures is covered in the following sections. This section is concerned with detection of and recovery from local failures, specifically hardware and software fault-tolerance and the use of online software upgrades to minimize system downtime.

- Survival of LSPs across software failure, and provision of online software upgrades in an MPLS system, are software implementation issues and should be addressed by any vendor serious about the provision of networking solutions. Tolerance of hardware faults relies on hardware detection and reporting of failures, on the availability of backup hardware, and on a suitably designed software implementation.

- Because RSVP is designed to run over a connectionless transport, it lends itself well to a system that must survive hardware failures or online software upgrades. Any control steps that are lost during the failover to the replacement backup system can be recovered by the state refresh processing that is built into RSVP.

- CR-LDP, on the other hand, assumes reliable delivery of control messages and so is not well placed to survive failover. Additionally, it is particularly hard to make TCP fault tolerant (a problem familiar to BGP implementers), with the result that a failover to a backup TCP stack results in the loss of the TCP connections. [7] This is interpreted by CR-LDP as a failure in all of the associated LSPs, which must subsequently be re-established from the ingress LSR.

- Metaswitch is researching ways to extend CR-LDP to allow it to survive online software upgrades and hardware faults.

- Until such extensions are added to CR-LDP, RSVP implementations will be able to provide better solutions for highly available MPLS networks.

2.6. Link and Peer Failure Detection

Where two LSRs are directly connected using a point-to-point link technology, such as ATM, the failure of LSPs can usually be detected by monitoring the state of the interfaces for the LSP. If, for example, an ATM link suffers loss of signal, both CR-LDP and RSVP can use the interface failure notification to detect the failure of the LSP. If two LSRs are connected over a shared medium, such as Ethernet, or are indirectly connected over a WAN cloud, for example using an ATM PVC, they may not necessarily receive a link failure notification from the link hardware.[8] LSP failure detection then relies on techniques inherent in the signaling protocols. So long as normal signaling traffic is flowing nothing else is necessary, but in stable state, additional processing is required to detect a failure.

- CR-LDP uses an exchange of LDP HELLO and KEEPALIVE messages to validate that the LSR peer and link are still active. Although TCP has a built-in keepalive system, this is typically too slow to respond to link and peer failures for the demands of MPLS LSPs.

- In RSVP, Path and Resv refresh messages serve to provide background traffic that indicates that the link is still active. However, to keep the per-LSP refresh traffic in a relatively stable network to a minimum, the refresh timer would be set quite high. To address this problem, an extension has been added to Labels RSVP so that RSVP HELLO messages can be exchanged to prove that the link and peer LSR are still active. The failure detection techniques and speed are therefore similar for both CR-LDP and RSVP, provided that RSVP uses the HELLO extensions. MPLS failure detection is much faster for directly attached LSRs.

2.7. Re-Routing

A strictly specified explicit route cannot be re-routed except by the ingress LSR (initiator). Consequently, failure at some point of an LSP must be reported to the ingress, effectively bringing down the whole LSP. However, a loosely specified portion of an explicit routed LSP, and any part of a hop-by-hop routed LSP, may be re-routed if

- a failure of a link or neighbor is detected (this is called Local Recovery)
- a better route becomes available
- the resources for the LSP are required for a new, higher priority LSP (this is called Pre-emption).

Re-routing is most easily managed from the ingress (including re-routing of strictly specified LSPs) and is supported by both CR-LDP and RSVP, though with slightly different characteristics. An LSR using RSVP can install a new route by simply refreshing the Path for an LSP to a different next-hop as soon as the alternate route is available/required. The old path can be left to time out because refreshes will no longer be sent. However, this wastes resources on the old path. “Make-before-break” is a mechanism whereby the old path is used (and refreshed) while the new path is set up, and then the LSR performing the re-routing swaps to using the new path and tears down the old path. This basic technique can be used to avoid double reservation of resources in both CR-LDP (using the modify value for the action flag on the LABEL_REQUEST) and RSVP (using shared explicit filters). Re-routing of loosely specified parts of LSPs at intermediate LSRs when a “better” route becomes available can lead to thrashing in unstable networks. To prevent this, a loosely specified part of a route may be “pinned”: In CR-LDP this is simply a matter of flagging the loose part of the explicit route as pinned. This means that once the route has been set up, it is treated as though it had been strictly specified and cannot be changed. In RSVP, pinning requires some additional processing. The initial route is specified with a loose hop. The Record Route object is used on the Path and Resv messages to feed back the selected route to the ingress. The ingress can use this information to re-issue the Path message with a strictly specified explicit route.

2.8. LSP Modification

LSP modification, for example, to change the traffic parameters for an LSP, is an equivalent operation to re-routing, though the change of route is optional for LSP modification. This means that the function is always present in RSVP and will be present in CR-LDP provided that the modify value of the action flag on a LABEL_REQUEST is supported by the implementations (this is a relatively recent addition to the CR-LDP drafts and so early implementations might not support modification of LSPs). Note that support for the modify value of the action flag in CR-LDP leads to increased data occupancy, bringing intermediate LSR occupancy up to a figure similar to that required at RSVP intermediate LSRs.

2.9. LAMBDA Networking

Lambda networking presents an interesting set of problems for an MPLS implementation. The full advantages of wavelength switching can only be encompassed if LSPs are switched in hardware without recourse to software. In this respect the lambda network is similar to an ATM network; the MPLS labels are identified with individual wavelengths. The number of wavelengths is, however, very small – too small for the likely number of LSPs transiting any one link. Additionally, the capabilities of an individual wavelength are far in excess of the normal requirements of an LSP, so that such a one-to-one mapping would be highly wasteful of network resources.

2.10. Traffic Control

Significantly, CR-LDP and RSVP perform resource reservation at different times in the process of LSP setup. CR-LDP carries the full traffic parameters on the LABEL_REQUEST. This allows each hop to perform traffic control on the forward portion of LSP setup. The traffic parameters can be negotiated as the setup progresses, and the final values are passed back on the LABEL_MAPPING allowing the admission control and resource reservation to be updated at each LSR. This approach means that an LSP will not be set up on a route where there are insufficient resources. RSVP carries a set of traffic parameters, the Tspec on the Path message.[8] This describes the data that is likely to use the LSP. Intermediate LSPs can examine this information and could make routing decisions based on it. However, it is not until the egress LSR is reached that the Tspec is converted to a Flowspec returned on the Resv message, which gives details of the resource reservation required for the LSP. This means that the reservation does not take place until the Resv passes through the network, with the result that LSP set up may fail on the selected route because of resource shortage. RSVP includes an optional function (adspec) whereby the available resources on a link can be reported on the Path message. This allows the egress LSR to know what resources are available, and modify the Flowspec on the Resv accordingly. Unfortunately, not only does this function require that all the LSRs on the path support the option, but it has an obvious window where resources reported on a Path message may already have been used by another LSP by the time the Resv is received. A partial solution for RSVP LSRs lies within the implementation, which could make a provisional reservation of resources as it processes the Path message. This reservation can only be approximate since it is based on the Tspec not the Flowspec, but it can considerably ease the problem. CR-LDP offers a slightly tighter approach to traffic control especially in heavily used networks, but individual RSVP implementations can provide a solution that is almost as good.

2.11. Policy Control

RSVP is specified to allow the Path and Resv messages to carry a policy object with opaque content. This data is used when processing messages to perform policy-based admission control. This allows Labels RSVP to be tied closely to policy policing protocols such as COPS (Common Open Policy Service) using the Internet draft “COPS Usage for RSVP”. By contrast, CR-LDP currently only carries implicit policy data in the form of the destination addresses, and the administrative resource class in the traffic parameters.

2.12. Layer 3 Protocol

Although an LSP can carry any data, there are occasions when knowledge of the layer 3 protocol can be useful to an intermediate or egress LSR. If an intermediate LSR is unable to deliver a packet (e.g. because of a resource failure) it can return an error packet specific to the layer 3 protocol (such as ICMP for IP packets) to notify the sender of the problem. For this to work, the LSR that detects the error must know the layer 3 protocol in use. Also, at an egress, it may help the LSR to forward data packets if the layer 3 protocol is known. RSVP identifies a single payload protocol during LSP setup, but there is no scope within the protocol for CR-LDP to do this. Even RSVP is unable to help when more than one protocol is routed to a particular LSP. Recent discussions led by Metaswitch in the MPLS Working Group have considered options for identifying the payload protocol in CR-LDP, and for marking the payload packets so that their protocol can be easily determined.

2.13. QOS AND DIFF-SERV

CR-LDP and RSVP have different approaches to Quality of Service (QoS) parameters. The RSVP Tspec object carried on Path messages describes the data that will flow rather than the QoS that is required from the connection. Various RFCs and Internet drafts describe how to map from different QoS requirements to the Tspec (for example, RFC 2210 - The Use of RSVP with IETF Integrated Services). The CR-LDP specification is more explicit about how the information carried on a LABEL_REQUEST message is mapped for QoS. Support for Diff-Serv (IP Differentiated Services) is addressed by an Internet draft (draft-ietf-mpls-diff-ext), which defines extensions to LDP, RSVP and CR-LDP. If implemented, this draft extends the full function of Diff-Serv to an MPLS network.

2.14. Provision of VPNs

Virtual Private Networks (VPNs) are an important feature of the service provided by ISPs to their customers. VPNs allow physically private networks to be extended to encompass remote sites by connecting them through the Internet. A customer in these circumstances expects to be able to preserve their IP addresses (which might not be globally unique) and to have the security of their data guaranteed. MPLS can provide an excellent solution as described in RFC 2547. Both CR-LDP and RSVP are suitable MPLS signaling protocols for VPNs over MPLS.

2.15. Voice Over IP and Voice Over MPLS

Voice over IP (VoIP) is an exciting development in Internet technology. The concept of a single infrastructure for voice and data, providing faster, cheaper and value-added services, is very attractive. MPLS is set to be a major component in VoIP networks, offering connection-oriented paths with resource reservation through the connectionless Internet. Voice over MPLS (VoMPLS) is the term given to the transfer of voice traffic over an MPLS network. This could involve establishing LSP Tunnels to act as trunks for multiple calls, or setting up LSPs for the duration of individual calls. Alternatively, VoMPLS could mean sending voice samples as labeled MPLS packets without including IP headers. Whichever approach is used, both CR-LDP and RSVP are suitable MPLS signaling protocols.

2.16. MIB Management

Traffic Engineered LSPs can be managed at their ingress and inspected at their egress through the MPLS Traffic Engineering MIB. This MIB is currently in an early stage which slightly favors CR-LDP, but new drafts will be produced that fully support RSVP and CR-LDP.

2.17. Acceptance/Availability

There is currently no clear “winner” between RSVP and CR-LDP in terms of market acceptance. Although generic RSVP has been available for a number of years from a variety of equipment vendors, and in that sense is an established network protocol, the changes required to a generic RSVP stack to add support for Labels RSVP are non-trivial, and hence Labels RSVP is in many respects a new protocol. CR-LDP is based on ideas that have been implemented in proprietary networks for as long as ten years, but as an IETF protocol it is very new and somewhat unproven.

2.18. Interoperability

There are two interoperability issues to be addressed. Do two implementations support a compatible set of options, and do they interpret the specifications in the same way? The option sets are functions of the flexibility of the protocol. RSVP has more implementation options than CR-LDP and so is perhaps at more risk. However, the protocol is specified to allow interworking between implementations that support different function sets. An IETF MPLS draft (draft-loa-mpls-cap-set) provides a list of capability sets to allow implementations to identify the functions that they provide. Interoperability testing is clearly the only way to prove that two implementations interwork correctly. Interoperability forums are being set up in many places including the University of New Hampshire InterOperability Labs (UNH IOL), George Mason University in Washington DC with the support of UUNET, EANTC in Berlin, NetWorld & Interop events. All of these forums will include work on RSVP and CR-LDP. Participation in interoperability events is a clear requirement for all MPLS software vendors. Testing for hardware vendors will be a combination of involvement in interoperability events, in-house testing with competitors’ equipment, and collaborative work with other vendors. Hardware vendors have a right to expect the support of their software suppliers during interoperability testing.

Summary

CR-LDP and Labels RSVP are both good technical solutions for setting up and managing Traffic Engineered LSPs. Early versions of both protocols had some functional omissions, but these are being fixed by subsequent Internet drafts so that the level of function provided by each protocol is similar. Some key differences in the structure of the protocols and the underlying transport mean that the support that the protocols can provide will never converge completely. These differences and the differences in speed and scope of deployment will be the main factors that influence vendors when they are selecting a protocol. The choice between RSVP and CR-LDP should be guided by the function of the target system. What LSP setup model will be used? How stable are the LSPs – do they represent permanent trunks or short-duration calls? How large is the network and how complex is it? Is this a stand-alone network or must the components interwork with other hardware and other networks? A final consideration must be the robustness of the hardware solution. What level of fault tolerance is required? How important is high availability?

- Two informational Internet drafts may help guide the choice of protocol.
- Applicability Statement for Extensions to RSVP for LSP-Tunnels
- Applicability Statement for CR-LDP

References

[1] AT&T Knowledge Ventures. (2007, July 25). *Transitioning to an MPLS Network*. Retrieved November 19th, 2007, from http://www.business.att.com/nx_resource.jsp?repoid=Topic&rtype=Whitepaper&rvalue=eb_fpoc_navigating_to_mpls_enabled_networks&repoitem=vpns&segment=ent_biz

[2] AT&T Knowledge Ventures. (2007, August 31). *Understanding VPN Technology Choices:Comparing MPLS, IPSec and SSL*. Retrieved November 19th, 2007, from http://www.business.att.com/nx_resource.jsp?repoid=Topic&rtype=Whitepaper&rvalue=understanding_vpn_technology_choices&repoitem=vpns&segment=ent_biz&guid=4BFDAE84-C61B-416F-886A-F606E9678B1C;08905D72-1FE7-450C-8EA5-B5F1565DD558

[3] Cisco Systems, Inc. (2004). *Managed VPN – Van Wijnen and Versatel*. Retrieved November 19th, 2007, from http://www.cisco.com/en/US/netsol/ns465/networking_solutions_customer_profile0900aecd801aa3f5.html

[4] Cisco Systems, Inc. (2005). *From Frame Relay to IP VPN: Why to Migrate, Why to Out-Task*. Retrieved November 18th, 2007, from http://www.cisco.com/en/US/netsol/ns458/networking_solutions_white_paper0900aecd8017a894.shtml

[5] Layer 2 MPLS VPN. (2007, October 20). In *Wikipedia, The Free Encyclopedia*. Retrieved 18:03, November 18, 2007, from http://en.wikipedia.org/w/index.php?title=Layer_2_MPLS_VPN&oldid=165912463

[6] Multiprotocol Label Switching. (2007, November 7). In *Wikipedia, The Free Encyclopedia*. Retrieved 18:04, November 18, 2007, from http://en.wikipedia.org/w/index.php?title=Multiprotocol_Label_Switching&oldid=169803565

[7] Pseudo-wire. (2007, November 17). In *Wikipedia, The Free Encyclopedia*. Retrieved 18:01, November 18, 2007, from http://en.wikipedia.org/w/index.php?title=Pseudowire&oldid=172121304

[8] Juniper Networks. Traffic Engineering for the New Public Network. http://www.omimo.be/magazine/00q4/2000q4_p054.pdf