

# The Coming of Age of MPLS

Rolf Winter, NEC Labs Europe

## ABSTRACT

Large operators have embraced multiprotocol label switching, deploying it in their backbone networks to enable a number of services and applications such as virtual private networks to just name one. Since the inception of the respective IETF working group, MPLS has accumulated a number of features and has been extended to be applicable in new contexts such as optical networks in the form of generalized MPLS. Currently, MPLS is being further extended to finally mature into a technology from which to build a full-fledged packet transport network that fulfills a large number of traditional transport network requirements. Having celebrated the early teens of MPLS in 2009, the IETF might well find itself celebrating MPLS's coming of age with the completion of the MPLS Transport Profile (MPLS-TP). This article is a short tutorial on what MPLS-TP is, how it came about, and what it promises to deliver in the future.

## INTRODUCTION

During the 74th Internet Engineering Task Force (IETF) meeting in 2009, the 12th birthday of multiprotocol label switching (MPLS) [1] was celebrated (counting from the formation of the IETF working group). Still in its teens, MPLS has found its way into many networks where it enables virtual private networks (VPNs), Border Gateway Protocol (BGP)-free core networks, traffic engineering, and other applications that have made it very popular among Internet service providers (ISPs). According to Request For Comments (RFC) 5218 [3], which defines metrics and criteria for assessing the success of a protocol or the lack thereof, MPLS can be officially declared a success. The two basic metrics that define success are wide deployment and the fact that a protocol meets its original design goals. According to RFC 5218, MPLS can even be called a "wild" success as it has exceeded its original design goals and is used in places that were not conceived when MPLS was being designed.

The foundation of the success of MPLS is its simplicity and flexibility. The only information MPLS adds to a packet are 20 bits worth of label information, a time-to-live (TTL) field, three bits to mark a packet's service class, and a bottom of stack delimiter bit.

The label, similar to an IP address, is used for forwarding a packet. Unlike an IP address, though, the label does not have any structure or format. Also unlike an IP address, a label only has local significance and changes at every hop

in the network. Relabeling a labeled packet is called *label swapping*, adding a label is called *label pushing*, and removing a label is called *label popping*. These operations are familiar expressions when talking about stacks, and indeed labels can be stacked. This is the MPLS way of building hierarchies. Together, these are the essential operations of the MPLS data plane.

Forwarding based on labels rather than IP addresses gives little extra benefit on its own. Much of its power, and difference from IP for that matter, stems from the way traffic is forwarded. IP follows the destination-based forwarding paradigm, which essentially means all forwarding is done based on the destination address in the IP header. MPLS, however, groups IP packets into so-called forwarding equivalence classes (FECs) at the boundary of the MPLS domain, which means all packets that are part of the same FEC receive the same forwarding treatment (i.e., are forwarded along the same label switched path, LSP). Conceptually, an FEC could be anything and is not limited to information in the IP header, which enables many of the applications mentioned before.

As for the control plane, MPLS does not mandate a single control protocol — another of MPLS' strengths, as this translates into flexibility. The choice of the control plane protocol, such as Label Distribution Protocol (LDP) [4] or Resource Reservation Protocol — Traffic Engineering (RSVP-TE) [5], depends on many factors. Some of these are network size, the particular role MPLS is supposed to play in the network, and the capabilities and applications that need to be supported, just to name a few.

This is MPLS in a nutshell; and, as already said, MPLS can doubtlessly be called a success, and a good indicator of its success is that MPLS has been used for applications and in ways that were not foreseen when MPLS was first designed. Another good indicator of MPLS's success is the way it has been extended to be used in other domains. Generalized MPLS (GMPLS) is a good example of this where the MPLS control plane is equipped with the means to enable forwarding based on time slots or wavelength, something not considered when MPLS was designed initially. Over time MPLS has also been extended to add important features that make it more robust, such as fast reroute. Another success story involving MPLS is the Pseudowire technology, which is used, for example, as a network convergence and migration technology where layer 2 frames (e.g., Ethernet) are carried over an MPLS network. And this list goes on.

With this wealth of features and its wide applicability, one might think that MPLS is about to go into maintenance mode in terms of standardization. But this is far from the truth. Currently, MPLS is further extending its reach and applicability in the form of the MPLS Transport Profile (MPLS-TP) [6] which is about to transform MPLS into a true packet transport network (PTN) technology.

## THE MPLS TRANSPORT PROFILE

The MPLS Transport Profile is a very special IETF effort. It all started as T(ransport)-MPLS within the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) as a connection-oriented packet switched technology for transport networks. Only later did it move to the IETF, where it was relabeled — pun intended — to MPLS-TP. The reason was that T-MPLS was not fully compatible with the currently specified IP/MPLS (in this document, we refer to the set of pre-MPLS-TP specifications as *IP/MPLS* where this distinction is relevant; otherwise, we just use the term *MPLS*), and when the IETF realized this was ongoing, it coordinated with the ITU-T and evaluated the options to continue this effort — either continue in the ITU-T as something that is not MPLS or bring it to the IETF, which has the design authority over MPLS, and develop it there. The latter option was chosen, and it was agreed that the ITU-T will reflect the changes made by the RFCs-to-be in their existing Recommendations.

Both ITU-T and IETF are quite different worlds or schools of thought where packet-switched meets circuit-switched and quality of service (QoS) meets best effort, but for MPLS-TP the participants of both standards bodies work together on its specification. This turns out to not always be easy, as sometimes these different worlds collide. Just to be able to talk to each other requires a document that translates the technical terms both communities use. While there are still a few lingering problems, the specifications are progressing. Specific solutions are still under discussion, but the requirements and some of the architecture and framework documents have already become RFCs.

That leaves the question of what MPLS-TP actually is. Broadly speaking, MPLS-TP is the effort to make MPLS a technology that can be used to create a PTN as defined by the ITU-T. It is the attempt to evolve traditional transport network technology (e.g., based on synchronous optical network/synchronous digital hierarchy [SONET/SDH]) to be more packet-friendly while retaining the characteristics of traditional transport technology such as high degree of availability, QoS support, the operational look and feel, among others.

The above is achieved by reusing MPLS and Pseudowire mechanisms as much as possible, where necessary extending MPLS and Pseudowires, and, if really necessary, designing new mechanisms to meet ITU-T's PTN requirements. This list of PTN requirements is indeed long, and a number of RFCs [7–9] specify them in different areas such as operations, administration, and

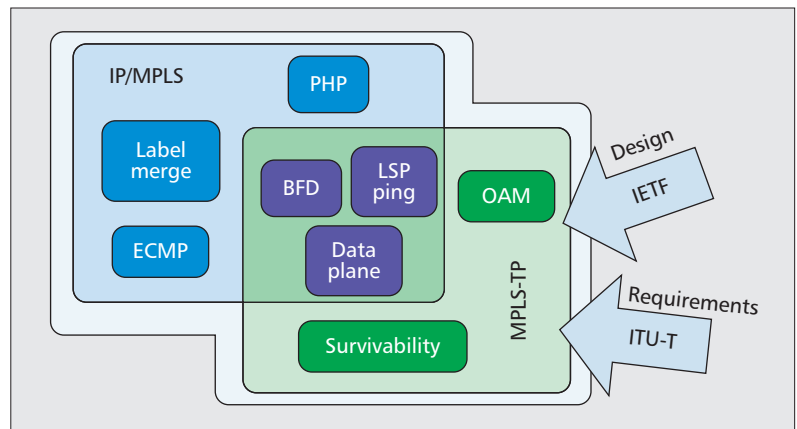


Figure 1. The “new” MPLS.

maintenance (OAM), security, and QoS. This does not always translate to new functionality that needs to be added to the MPLS toolkit. As Fig. 1 depicts, MPLS-TP does not even make use of the full set of MPLS functions. It is rather a necessary and sufficient subset of MPLS with added functionality where current MPLS is not up to the job (note that the figure is not complete; nor are the proportions to be interpreted as a representation of the actual amount of functionality).

In the following sections we take a closer look at what MPLS-TP does not use or require from the existing IP/MPLS toolkit, what the intersection of IP/MPLS and MPLS-TP consists of, and finally, the new tools MPLS-TP adds to the overall MPLS machinery.

### IP/MPLS BUT NOT MPLS-TP

As already said, MPLS-TP uses only a subset of the currently specified IP/MPLS functionality. There are various reasons for excluding certain functionality. For example, penultimate hop popping (PHP) is not supported as part of the transport profile because it would conflict with some of the OAM tools being specified. Basically, the removal of the label at the penultimate hop removes important information that is needed for certain OAM mechanisms to work. For similar reasons LSP merging is not permissible. Also, equal cost multipath (ECMP) is not supported as there is a strong requirement that OAM packets share their fate with data packets. When ECMP is employed, this can no longer be guaranteed.

Probably the biggest difference from IP/MPLS is that the support of IP is optional for the transport profile. In other words, even in the absence of IP, MPLS-TP needs to be able to work properly and provide its full functionality as specified. As IP/MPLS does rely on IP in various places, this is indeed a consideration with significant impact. But a typical transport network environment indeed does not rely on the presence of IP.

The same holds for a dynamic control plane. A transport network consists of long lasting connections, often managed via an operation support system (OSS) used to establish and configure paths. To retain the operational look and feel of transport networks, this way of running an MPLS-TP network must be supported. More advanced methods such as a dynamic con-

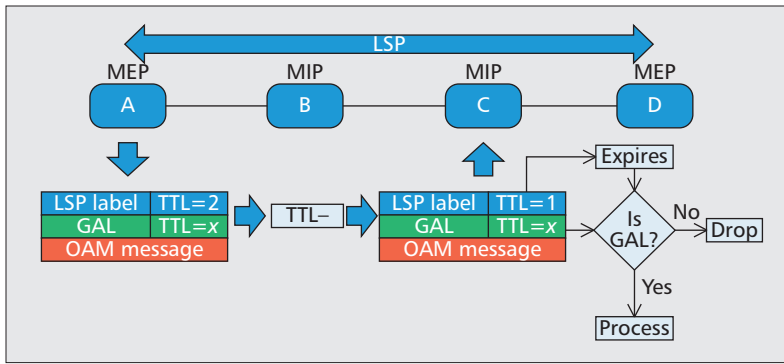


Figure 2. OAM operation using the generic alert label mechanism.

trol plane are therefore purely optional. These restrictions often require the extension of existing mechanisms, as seen in the next section.

### BOTH IP/MPLS AND MPLS-TP

Probably the most important element of the intersection between IP/MPLS and MPLS-TP is the compatibility requirement of the data plane. MPLS cannot make any modifications to the IP/MPLS data plane that would render the two incompatible. As mentioned before, certain functions such as PHP were excluded from the transport profile; however, this does not impact compatibility with the IP/MPLS data plane. As mentioned before, MPLS-TP is first and foremost a necessary and sufficient subset of MPLS.

Generally, the IETF is attempting to reuse as much as possible from IP/MPLS and other existing technology for the transport profile, which is reflected in the way a number of existing mechanisms are extended to fulfill the transport role. A good example is bidirectional forwarding detection (BFD). BFD is a simple failure and liveness detection protocol that is used between routers; for example, BFD is designed to detect failures much faster than native routing mechanisms, which often operate on a timescale of seconds. The transport network requirements demand functions such as connectivity check (CC), connectivity verification (CV), and remote defect indication (RDI), or, in other words, functions to check liveness, verify that one endpoint is connected to the expected endpoint(s), and a function to report defects to the other end. BFD is in principle capable of doing this, albeit small modifications are needed, and some options for BFD are not available in the transport profile. So the IETF decided to go with BFD to fulfill these MPLS-TP requirements. Similarly, LSP-Ping can fulfill similar tasks and more, such as route tracing, with some tweaking. While BFD is the proactive way of realizing these functions (i.e., in a preconfigured continuous way), LSP-Ping is the on-demand version of implementing them.

Another requirement on MPLS-TP is that OAM packets need to share their fate with regular user traffic. That means OAM packets need to travel in-band on the forwarding path, that is, in the data plane of the transport path (or a subpart of the transport path) to which the OAM operation applies. These OAM functions also need to operate in environments without IP and dynamic control plane, as these are optional in MPLS-TP, as mentioned before. All of the requirements

above can be met by the Pseudowire associated channel. Unfortunately, it has only been specified for Pseudowires, so the IETF generalized this mechanism to also be applicable to LSPs. For MPLS-TP, this generic associated channel (G-ACh) [10] is a central mechanism that enables OAM (and other control mechanisms) in transport network environments. In order to make nodes aware that a packet belongs to an associated channel, RFC 5586 also reserves a special label (the G-ACh label, GAL) as an exception mechanism that identifies packets as part of a control channel at the bottom of the label stack. To address a node in such an environment TTL expiry can be used: when the TTL expires, a node checks the bottom of stack label. In case it finds the GAL it will inspect the associated channel header to see what type of OAM function needs to be performed. This process is illustrated in Fig. 2. The four nodes A, B, C and D in the figure are part of an LSP, where A and D are the Label Edge Routers. A inserts a labeled packet into the LSP with a TTL of 2, thereby addressing the node that is two hops away on that LSP (node C). At node C the TTL expires and since the GAL is present, the packet is processed further to examine and react on the OAM message that the packet contains.

The message from all this is clear, IP/MPLS tools are not that far away from fulfilling the elaborate MPLS-TP requirements and the IETF is working busily to make them applicable for the transport profile.

### MPLS-TP BUT NOT IP/MPLS

The features IP/MPLS is missing in order to build a PTN that fulfills the ITU-T requirements are mainly in the areas of OAM, network management and survivability mechanisms.

When it comes to OAM, the ITU-T has a very specific OAM model or framework that was brought to the IETF [1]. This model introduces terminology, constructs, and a few limitations into the MPLS-TP OAM world. The most basic OAM functional component in this model is a maintenance entity (ME), which is the relationship between two points on a transport path. OAM operations always apply to an ME. The endpoints of an ME are referred to as maintenance entity group (MEG) endpoints (MEPs), and elements located between those endpoints are called MEG intermediate points (MIPs); the latter might or might not exist.

Although the above summary of the OAM framework is far from complete, the overall model is rather generic, and from the above it seems to simply introduce terminology — naming the functional components that participate in a certain OAM operation. But this framework comes with certain restrictions that impact how OAM can actually be implemented. As an example, MIPs are not allowed to initiate OAM packets on their own without external triggers. This functionality is reserved for MEPs, and MIPs can only react to OAM messages that have been sent to them or simply forward OAM messages. This seems like an artificial restriction, but it has been applied in transport networks for some time. MPLS-TP OAM solutions need to operate within this framework.

Loss and delay measurements are examples of missing OAM functions in IP/MPLS, and a current draft specifies two simple protocols that

enable these measurements which are carried over the G-ACh. As mentioned before, apart from OAM, survivability mechanisms also need to be added, that is, mechanisms that recover the network's ability to deliver packets correctly after a failure or the severe degradation of the network's performance. Optical transport network equipment has traditionally delivered a high benchmark in terms of reliability and the ability to recover from failure and MPLS-TP needs to play in that league. Therefore, e.g., linear protection and ring protection solutions for MPLS-TP are being actively worked on. However, for many survivability solutions there are currently multiple drafts that propose different mechanisms and many of these specifications are still moving targets.

While the new additions to MPLS as part of the MPLS-TP effort are enabling MPLS to play a transport network role, there is nothing that prevents many of these to be applied in an IP/MPLS environment; and one should not forget that everything that is added as part of MPLS-TP is still core MPLS, not a parallel technology. Whatever is being standardized under the label MPLS-TP together with IP/MPLS is the new MPLS.

## WHY MPLS-TP

In the previous sections, we have broadly covered what MPLS-TP is and how it fulfills transport network requirements. Before that we summarized how it came about. This really leaves the question of why one actually wants MPLS-TP.

The wish to introduce packet-oriented technology into transport networks comes from the fact that the traffic in today's network is predominantly IP traffic (i.e., packets). MPLS-TP promises to handle packet-based services more efficiently than traditional circuit-switched transport network technology. As MPLS is already deployed in many core networks today and is a connection-oriented packet-switched technology, MPLS was the best candidate as a starting point for a PTN. Using MPLS as the basis will also make interworking with the core network of large providers easier, and some operators are already thinking out loud about deploying MPLS not just in their core network but also in the aggregation and access network, some of which just might be MPLS-TP.

As you go all the way up from lambdas to IP routing, at every layer you add not only functionality but also cost. IP/MPLS, with all the features it provides, is a quite costly technology; however, MPLS-TP is removing many of the costly bits such as the reliance on IP and, with it, on IP routing protocols and so forth. These functions are often simply not needed in a transport context; as such, MPLS-TP also promises many of the MPLS advantages at a lower price. If some of these optional capabilities are really needed, there is nothing that prevents use of them, such as using GMPLS as the control plane — obviously at a cost.

## CONCLUSION

This article has only scratched the surface of what MPLS-TP really is and has probably not well reflected the huge amount of energy behind the standardization of MPLS-TP. Two communities,

ITU-T and IETF, busily — but not always smoothly — work on the specification of new and extension of existing mechanisms to create a profile of MPLS that can be used in a transport network context. While individual solutions that fulfill the long list of requirements are still being discussed, the general framework of MPLS-TP has already been finished.

MPLS-TP uses a subset of the existing MPLS toolkit which represents the minimal amount of functionality that is needed in a transport network environment. As such, IP and control plane support is not required for MPLS-TP. Central management via an OSS and static configuration is the basic operational mode of MPLS-TP, which therefore retains the look and feel of today's transport networks. With the features added as part of the MPLS-TP effort, it will deliver the reliability, QoS support, and richness of OAM mechanisms that are characteristics of transport technology at the same time as it will support packet-based services more efficiently than traditional transport networks.

Not even yet fully standardized, other fora have already discovered MPLS-TP and are looking into how it can be used for their purposes. For example, the Broadband Forum has started some early stage work on it. Also, first interoperability tests of OAM mechanisms have been performed (e.g., at Isocore) by a number of equipment vendors.

With all this energy, a broad interest in the technology and the promises MPLS-TP makes, it appears that MPLS-TP, just like IP/MPLS before it, will be a success. Once MPLS has matured into a technology that can be used to construct packet transport networks, the IETF has every reason to celebrate MPLS' coming of age.

## ACKNOWLEDGMENT

Rolf Winter is partly funded by Trilogy, a research project supported by the European Commission under its Seventh Framework Programme.

## REFERENCES

- [1] MPLS Working Group, <http://datatracker.ietf.org/wg/mpls/>
- [2] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," IETF RFC 3031, Jan. 2001.
- [3] D. Thaler and B. Aboba, "What Makes for a Successful Protocol?," IETF RFC 5218, July 2008.
- [4] L. Andersson, I. Minei, and B. Thomas, "LDP Specification," IETF RFC 5036, Oct. 2007.
- [5] D. Awduche et al., "RSVP-TE: Extensions to RSVP for LSP Tunnels," IETF RFC 3209, Dec. 2001.
- [6] M. Bocci et al., "A Framework for MPLS in Transport Networks," IETF RFC 5921, July 2010.
- [7] B. Niven-Jenkins et al., "Requirements of an MPLS Transport Profile," IETF RFC 5654, Sept. 2009.
- [8] M. Vigoureux, D. Ward, and M. Betts, "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks," IETF RFC 5860, May 2010.
- [9] K. Lam, S. Mansfield, and E. Gray, "Network Management Requirements for MPLS-based Transport Networks," IETF RFC 5951, Sept. 2010.
- [10] M. Bocci, M. Vigoureux, and S. Bryant, "MPLS Generic Associated Channel," IETF RFC 5586, June 2009.

## BIOGRAPHY

ROLF WINTER (rolf.winter@neclab.eu) is a senior researcher at NEC Laboratories Europe where he is involved in networking research and standardization, the latter mainly in the context of IETF, where he also chairs the LEDBAT working group. He received his Ph.D. from the Freie University Berlin in 2006. Currently, he is involved in the Trilogy project, where he is working in the areas of routing, resource control, and economics of the Internet of the future.

*With all this energy, and broad interest in the technology and the promises MPLS-TP makes, it appears that MPLS-TP, just like IP/MPLS before it, will be a success. Once MPLS has matured into a technology that can be used to construct packet transport networks, the IETF has every reason to celebrate MPLS' coming of age.*