Network Working Group
Internet-Draft
Intended status: Informational
Expires: March 23, 2012

N. Sprecher
Nokia Siemens Networks

KY. Hong
Cisco Systems

September 23, 2011

The Reasons for Selecting a Single Solution for MPLS-TP OAM

draft-sprecher-mpls-tp-oam-considerations-01.txt

Abstract

The MPLS Transport Profile (MPLS-TP) is a profile of MPLS technology for use in transport network deployments. That is, MPLS-TP is a set of functions and features selected from the wider MPLS toolset and applied in a consistent way to meet the needs and requirements of operators of packet transport networks.

During the process of development of the profile, additions to the MPLS toolset have been made to ensure that the tools available met the requirements. These additions were motivated by MPLS-TP, but form part of the wider MPLS toolset such that any of them could be used in any MPLS deployment.

One major set of additions provides enhanced support for Operations, Administration, and Maintenance (OAM). This enables fault management and performance monitoring to the level needed in a transport network. Many solutions and protocol extensions have been proposed to address these OAM requirements, and this document sets out the reasons for selecting a single, coherent set of solutions for standardization.

Table of Contents

1. Introduction

The MPLS Transport Profile (MPLS-TP) is a profile of MPLS technology for use in transport network deployments. That is, MPLS-TP is a set of functions and features selected from the wider MPLS toolset and applied in a consistent way to meet the needs and requirements of operators of packet transport networks.

Operations, Administration, and Maintenance (OAM) plays a significant role in carrier networks, providing methods for fault management and performance monitoring in both the transport and the service layers, and enabling these layers to support services with guaranteed and strict Service Level Agreements (SLAs) while reducing their operational costs.

OAM provides a comprehensive set of capabilities that operate in the data plane. Network-oriented mechanisms are used to monitor the network's infrastructure in order to enhance the network's general behavior and level of performance. Service-oriented mechanisms are used to monitor the services offered to end customers. Such mechanisms enable rapid response to a failure event and facilitate the verification of some SLA parameters. Fault management mechanisms are used for fault detection and localization as well as for diagnostics and notification. Performance management mechanisms enable monitoring of the quality of service with regard to key SLA criteria (e.g., jitter, latency, and packet loss).

During the process of development of MPLS-TP, additions to the MPLS toolset have been made to ensure that the tools available meet the requirements. These additions were motivated by MPLS-TP, but form part of the wider MPLS toolset such that any of them could be used in any MPLS deployment.

One major set of additions provides enhanced support for OAM. Many solutions and protocol extensions have been proposed to address these OAM requirements. This document sets out the reasons for selecting a single, coherent set of OAM solutions for standardization.

1.1. Background

The ITU-T and IETF jointly commissioned a Joint Working Team (JWT) to examine the feasibility of a single, collaborative solution to support OAM requirements for MPLS transport networks for MPLS-TP. The JWT reported that it is possible to extend the MPLS technology to fully satisfy the requirements [RFC5317]. The investigation by the JWT laid the foundations for the work to extend MPLS, but a thorough technical analysis was subsequently carried out within the IETF with strong input from the ITU-T to ensure that the MPLS-TP OAM

requirements provided by the ITU-T and IETF would be met.

The report of the JWT [RFC5317] as accepted by the ITU-T was documented in [TD7] and was communicated to the IETF in a liaison statement [LS26]. In particular, the ITU-T stated that any extensions to MPLS technology will be progressed via the IETF standards process using the procedures defined in [RFC4929].

[RFC5317] includes the analysis that "it is technically feasible that the existing MPLS architecture can be extended to meet the requirements of a Transport profile, and that the architecture allows for a single OAM technology for LSPs, PWs, and a deeply nested network."

[RFC5654] in general, and [RFC5860] in particular, define a set of requirements for OAM functionality in MPLS-TP which are applicable to MPLS-TP Label Switched Paths (LSPs), Pseudowires (PWs), and MPLS-TP links. These documents are the results of a joint effort by the ITU-T and the IETF to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to enable the deployment of a packet transport network that supports the capabilities and functionalities of a transport network as defined by the ITU-T. The OAM requirements are derived from those specified by ITU-T in [Y.Sup4].

An analysis of the technical options for OAM solutions was carried out by a design team (the MEAD team) consisting of experts from both the ITU-T and the IETF. The team reached an agreement on the principles of the design and the direction for the development of an MPLS-TP OAM toolset. A report was subsequently submitted to the IETF MPLS Working Group at the Stockholm IETF meeting in July 2009. The guidelines drawn up by the design team have played an important role in the creation of a coherent MPLS-TP OAM solution.

The MPLS working group has modularized the function of MPLS-TP OAM allowing for separate and prioritized development of solutions. This has given rise to a number of documents each describing a different part of the solution toolset. At the time of writing, the most important of these documents have completed development within the MPLS working group and are advancing through the IETF process towards publication as RFCs. These documents cover the following OAM features:

- Continuity Check
- Connection Verification
- On-demand Connection Verification
- Route Tracing
- Remote Defect Indication
- Packet Loss Measurement

- Packet Delay Measurement
- Linear Protection Control
- Lock Instruction
- Loopback Testing
- Fault Management

The standardization process within the IETF allows for the continued analysis of whether the OAM solutions under development meet the documented requirements, and facilitates the addition of new requirements if any are discovered. It is not the purpose of this document to analyze the correctness of the selection of specific OAM solutions. This document is intended to explain why it would be unwise to standardize multiple solutions for MPLS-TP OAM, and to show how the existence of multiple solutions would complicate MPLS-TP development and deployment making networks more expensive to build, less stable, and more costly to operate.

1.2. The Development of a Parallel MPLS-TP OAM Solution

It has been suggested that a second, different OAM solution should also be developed and documented in an ITU-T Recommendation. Various arguments have been presented for this duplication of effort, including:

- Similarity to OAM encodings and mechanisms used in Ethernet
- The existence of two distinct MPLS-TP deployment environments called Packet Switched Network (PSN) and Packet Transport Network (PTN).
- The need for similar operational experience in MPLS-TP networks and in pre-existing transport networks (especially SONET/SDH networks).

The first of these was discussed within the IETF's MPLS working group where precedence was given to adherence to the JWT's recommendation to select a solution that reused as far as possible pre-existing MPLS tools. Additionally, it was considered that consistency with encodings and mechanisms used in MPLS was of greater importance.

The second argument has not been examined in great detail because substantive evidence of the existence of two deployment environments has not been documented or presented. Indeed, one of the key differences between the two environments is the choice of MPLS-TP OAM solution which makes a circular argument.

The third argument contains a very important point: network operators want to achieve a smooth migration from legacy technologies such as SONET/SDH to their new packet transport networks. This transition can be eased if the new networks offer similar OAM features and can

be managed using tools with similar look and feel. The requirements specifications [RFC5654] and [RFC5860] capture the essential issues that must be resolved to allow the same look and feel to be achieved. Since the OAM solutions developed within the IETF meet the documented requirements, Network Management Systems (NMS) can easily be built to give the same type of control of MPLS-TP networks as is seen in other transport networks. Indeed, it should be understood that the construction of an NMS is not dependent on the protocols and packet formats within the OAM, but on the high-level features and functions offered by the OAM.

This document does not debate the technical merits of any specific solution. That discussion, and the documentation of MPLS-TP OAM specifications, was delegated by the IETF and ITU-T to the MPLS working group and can be conducted using the normal IETF, consensus-driven process. [OAM-Overview] presents an overview of the OAM mechanisms that have already been defined and that are currently being defined by the IETF, as well as a comparison with other OAM mechanisms that were defined by the IEEE and ITU-T.

This document focuses on an examination of the consequences of the existence of two MPLS-TP OAM solutions.

2. Terminology and References

2.1. Acronyms

This document uses the following acronyms:

CESoPSN	Circuit Emulation Service over Packet Switched Network
IEEE	Institute of Electrical and Electronics Engineers
ITU-T	International Telecommunications Union - Telecommunication Standardization Sector
JWT	Joint Working Team
LSP	Label Switched Path
MPLS-TP	MPLS Transport Profile
NMS	Network Management Systems
OAM	Operations, Administration, and Maintenance
PSN	Packet Switched Network
PTN	Packet Transport Network
PW	Pseudowire
PWE3	Pseudowire Emulation Edge to Edge
SAToP	Structure-Agnostic Time Division Multiplexing over Packet
SLA	Service Level Agreements
TDM	Time Division Multiplexing
TDMoIP	Time Division Multiplexing over IP

3. Motivations for a Single OAM Solution in MPLS-TP

This section presents a discussion of the implications of the development and deployment of more than one MPLS OAM protocol. The summary is that it can be seen that there are strong technical, operational, and economic reasons to avoid the development and deployment of anything other than a single MPLS OAM protocol.

3.1. MPLS-TP is an MPLS Technology

MPLS-TP is an MPLS technology. It is designed to apply MPLS to a new application. The original proposers of the concept assumed that the transport variant of MPLS would always exist in a disjoint network, and indeed with their first attempt at the technology (T-MPLS) had a number of significant incompatibilities with MPLS that were irreconcilable. When it was established that co-existence in the same layer network could and would occur, T-MPLS development was stopped and the development of MPLS-TP was begun. In MPLS-TP, MPLS was extended to satisfy the transport network requirements in a way that was both compatible with MPLS as has already been deployed, and MPLS and with the IETF envisioned it would develop in the future.

Given this intention for compatibility, it follows that the MPLS-TP OAM protocols should be consistent with the existing, deployed MPLS OAM, and should be designed according to the design philosophies that have lead to the current widespread adoption of MPLS. Key elements here are scalability and hardware independence, i.e. that the tradeoff between scaling to large numbers of monitored objects and the performance of the monitoring system should be a matter for vendors and operators to resolve, and that the tradeoff should be a soft transition rather than a cliff. Furthermore there should be no requirement to execute any component (other than packet forwarding) in hardware to achieve usable performance.

3.2. MPLS-TP is a Convergence Technology

It is possible to argue that using MPLS for Transport is only a stepping stone in the middle of a longer transition. Quite clearly all communication applications are being moved to operate over the Internet protocol stack of TCP/IP/MPLS, and the various layers that have existed in communications networks are gradually being collapsed into the minimum necessary set of layers. Thus, for example, we no longer run IP over X.25 over HDLC over multi-layered Time Division Multiplexing (TDM) networks.

Increasingly the entire point of transport networks is to support the transmission of TCP/IP/MPLS. Using MPLS to construct a transport network is a relatively short-term stepping-stone toward running IP

and MPLS directly over fiber optics. MPLS has been deployed in operational networks for approximately a decade, with the existing MPLS OAM techniques widely deployed in operational networks. Service providers are not going to stop using the MPLS based OAM techniques that they have been using for years, and no one has proposed that they would. Thus, the question is not which OAM to use for transport networks; the question is whether service providers want to use two different sets of OAM tools in the future converged network. When we arrive at our destination of TCP/IP/MPLS running directly over fiber, the operators will use MPLS OAM tools to make this work.

3.3 There is an End-to-End Requirement for OAM

The purpose of OAM is usually to execute a function that operates end-to-end on the monitored object (such as an LSP or PW). Since LSPs and PWs provide edge-to-edge connectivity and can transition network operator boundaries, the OAM must similarly operate across network operator boundaries. This is particularly the case with the continuity check and connection verification functions that are needed to test the end-to-end connectivity of LSPs and PWs. It is, therefore, necessary that any two equipments that could ever be a part of an end-to-end communications path have a common OAM. This necessity is emphasized in the case of equipment executing an edge function, since with a global technology such as MPLS it could be interconnected with an edge equipment deployed by any other operator in any part of the global network.

This is a design paradigm that has guided the IETF in the development of its exiting network layer connectivity OAM protocols. For each network layer protocol there is only one ping, trace route, or fast connectivity protocol, and amongst these there is a common design approach.

This leads to the conclusion that for any network layer protocol every equipment needs to be able to execute or to inter-work with a canonical form of the OAM. As we shall demonstrate, inter-working between protocols adds significant complexity, and thus a single default OAM is strongly preferred.

3.4. The Complexity Sausage

A frequent driver for the replacement of an established technology is a perception that the new technology is simpler, and thus of greater economic benefit to the user. In an isolated system this may be the case, however when, as is usually case with communications technologies, simplification in one element of the system introduces a (possibly non-linear) increase in complexity elsewhere.

This creates the "squashed sausage" effect, where reduction in complexity at one place leads to significant increase in complexity at a remote location. When we drive complexity out of hardware by placing complexity in the control plane there is frequently an economic benefit, as illustrated by MPLS itself. However, when we drive OAM complexity out of one network element at the cost of increased complexity at a peer network element we loose out economically and, more importantly, we loose out in terms of the reliability of this important network functionality.

Furthermore, due to the need to ensure compatibility of an inter-working function between the two solutions (in order to achieve end-to-end OAM) we create a situation where neither of two disjoint protocol developments is able to make technical advances. Such a restriction is unacceptable within the context of the Internet.

3.5. Inter-Working is Expensive and Has Deployment Issues

The issue of OAM inter-working can easily be illustrated by considering an analogy with people speaking different languages. Inter-working is achieved through the use of an interpreter. The interpreter introduces cost, slows down the rate of information exchange, and may require transition through an intermediate language. There is considerable risk of translation errors and semantic ambiguities. These considerations also apply to computer protocols, particularly give the ultra-pedantic nature of such systems. In all cases, the availability of a single working language dramatically simplifies the system, reduces cost, and speeds reliable communication.

If two MPLS OAM protocols were to be deployed we would have to consider three possible scenarios:

1) Isolation of the network into two incompatible and unconnected islands.
2) Universal use of both OAM protocols.
3) Placement of inter-working (translation) functions or gateways.

We have many existence proofs that isolation is not a viable approach, and the reader is referred to the early T-MPLS discussions for examples. In summary, the purpose of the Internet is to achieve an integrated universal connectivity. Partition of the network into incompatible and unconnected islands is neither desirable nor acceptable.

Universal deployment of both OAM protocols requires the sum of the costs associated with each protocol. This manifests as implementation time, development cost, memory requirements, hardware

components, and management systems. It introduces additional testing requirements to ensure there are no conflicts when both protocols are run on a common platform. It also requires code and processes to deal with the negotiation of which protocol to use and conflict resolution (which may be remote and multi-party) when an inconsistent choice is made. In short, this option results in worse than double costs, increases the complexity of the resulting system, risks the stability of the deployed network, and makes the networks more expensive and more complicated to operate.

The third possibility is the use of some form of inter-working function. This is not a simple solution and inevitably leads to cost and complexity in implementation, deployment, and operation. Where there is a chain of communication (end-to-end messages passed through a series of transit nodes) a choice must be made about where to apply the translation and inter-working.

- In a layered architecture, inter-working can be achieved through tunneling with the translation points at the end-points of the tunnels. In simple network diagrams this can look very appealing and only one end-node is required to be able to perform the translation function (effectively speaking both OAM languages). But in the more complex reality of the Internet, nearly every network node performs the function of an end-node, and so the result is that nearly every node must be implemented with the capability to handle both OAM protocols and to translate between them. This turns out be even more complex than the universal deployment of both protocols discussed above.

- In a flat architecture, interworking is performed at a "gateway" between islands implementing different protocols. Gateways are substantially complex entities that usually have to maintain end-to-end state within the network (something that is against one of the fundamental design principles of the Internet) and must bridge the differences in state machines, message formats, and information elements in the two protocols. The complexity of gateways make them expensive, fragile and unstable, hard to update when new revisions of protocols are released, and difficult to manage.

To deploy an inter-working function it is necessary to determine whether the OAM protocol on the arriving segment of the OAM is identical to the OAM protocol on the departing segment. Where the protocols are not the same, it is necessary to determine which party will perform the translation. It is then necessary to route the LSP or PW through a translation point that has both sufficient translation capacity and sufficient data bandwidth and adequate path diversity. When an upgraded OAM function is required, the problem

changes from a two party negotiation to an n-party negotiation with commercial and deployment issues added to the mix.

Note that when an end-to-end LSP or PW is deployed, it may transit many networks and the OAM might require repeated translation back and forth between the OAM protocols. The consequent loss of information and potential for error is similar to the children's game of Chinese Whispers.

The IETF has, with good reason, a history of strongly opposing proposals to inter-work protocols.

3.6. Selection of a Single OAM Solution When There is a Choice

When there is a choice of protocols for deployment or operation, a network operator must make a choice. Choice can be a good thing when it provides for selection between different features and functions, but it is a burden when the protocols offer essentially the same functions, but are incompatible.

In this case, the elements of the choice include:
- Which protocol will continue to be developed by its SDO?
- Which protocol is most stable in implementations?
- How to test and evaluate the two protocols?
- Which vendors support and will continue to support which protocol?
- What equipment from different vendors is compatible?
- Which management tools support which protocols?
- What protocols are supported by peer operators and what interworking function is needed?
- Which protocols are engineers experienced with and trained in?
- What are the consequences of a wrong choice?
- Will it be possible to migrate from one protocol to another in the future?
- How to integrate with other function already present in the network?
- How to future-proof against the inclusion of new function in the network?

At the very least, the evaluation of these questions constitute a cost and introduce delay for the operator. The consequence of a wrong choice could be very expensive, and it is likely that any comparative testing will more than double the lab-test costs prior to deployment.

From a vendor's perspective, the choice is even harder. A vendor does not want to risk not offering a product for which there is considerable demand, so both protocols may need to be developed leading to more than doubled development costs. Indeed, code

complexity and defect rates have often been shown to increase worse than linearly with code size, and (as noted in a previous section) the need to test for co-existence and interaction between the protocols adds to the cost and complexity.

It should be noted that, in the long-run, it is the end-users who pay the price for the additional development costs and any network instability that arises.

3.7. Migration Issues

Deployment of a technology that is subsequently replaced or obsoleted often leads to the need to migrate from one technology to another. Such a situation might arise if an operator deploys one of the two OAM protocol solutions and discovers that it needs to move to use the other one.

When the migration is between versions of a protocol, it may be that the new version is defined to support the old version. When the implementation is in software, upgrades can be managed centrally. However, neither of these would be the case with MPLS-TP OAM mechanisms, and hardware components would need to be upgraded in the field using expensive call-out services.

Hardware upgrades are likely to be service-affecting even with sophisticated devices with redundant hardware components. Furthermore, since it would be impractical to upgrade every device in the network at the same time, there is a need for either a significantly large maintenance period across the whole network (something that would inevitably drive all customers away!) or for a rolling plan that involves upgrading nodes one at a time with new hardware that has dual capabilities. Such hardware is, of course, more expensive and more complex to configure than hardware dedicated to just one OAM protocol.

Additionally, the transition phase of migration leads to dual mode networks as described in Section 6.3. Such networks are not desirable because of their cost and complexity.

In short, the potential for future migration will need to be part of the deployment planning exercise when there are two OAM protocols to choose between. This issue will put pressure on making the "right" choice when performing the selection described in Section 3.6.

4. Examples of Inter-Working Issues in the Internet

It is, of course, right to observe that there are a number of instances of multiple protocols serving the same purpose that have

arisen within the Internet. It is valuable to examine these examples to understand what issues they have caused and how they have been mitigated.

4.1. ISIS/OSPF

ISIS and OSPF are two competing link state IGP routing protocols that derive from the same root document and which, for political and personality reasons, were never reconciled prior to wide-scale deployment. It is an accident of history that one of these protocols did not gain overwhelming deployment and so force the other into retirement.

The existence of these two widely deployed and highly functional competing IGPs more than doubles the cost of link state IGP maintenance and deployment in the Internet. This is a situation that will almost certainly continue for the lifetime of the Internet. Although the Internet is clearly successful and operates well, the existence of these two IGPs forces all serious router vendors to implement both protocols (doubling the protocol cost of all routers even when an operator only wants to deploy one of the protocols), forcing an operator to make an active choice between IGPs during deployment, and requiring a gateway function between the islands of protocol use.

A mitigating factor in this specific case is that, owing to the way networks are partitioned for administrative and scaling reasons, there already existed a gateway routing protocol called BGP that propagates a summarized form of the IGP reachability information through-out the Internet. BGP means that there is actually no requirement for ISIS and OSPF to inter- work: there is no need for a translation function, and the two IGPs can continue to exist without impacting the function of the Internet. Thus, unlike the situation with MPLS OAM, the choice of IGP protocol is truly a local decision.

4.2. Time Division Multiplexing Pseudowires

The IETF's PWE3 working group has published the specification of three different TDM PW types. This happened after considerable effort to reach a compromise failed to reduce the set of options.

- SAToP is a relatively simple design. It is a Draft Standard and is the mandatory to implement, default, mode of operation.

- CESoPSN and TDMoIP are more complex approaches with different degrees of bandwidth efficiency optimized for different applications.

In this case all implementations must include the default mode of operation (SAToP). This means that end-to-end operation is guaranteed: an operator can select equipment from any vendor in the knowledge that they will be able to build and operate an end-to-end TDM PW service.

If an operator wishes to deploy a TDM PW optimized for a specific application they may select equipment from a vendor offering CESoPSN or TDMoIP in addition to SAToP. Provided that all of their equipment and their management system can handle the optimized approach, they can run this in their network, but the operator has to carry the cost of selecting, purchasing, configuring, and operating the extended mode of operation.

This situation is far from ideal, and it is possible that long-distance, multi-operator optimized TDM PWs cannot be achieved. However, the existence of a default mode implemented in all devices helps to reduce pain for the operator and ensures that simpler end-to-end operation is always available. Additionally, the growth of other protocols is acting to diminish the use of long distance TDM circuits making this a self limiting problem.

4.3. Codecs

The n-squared codec inter-working problem was brought to the attention of the IETF by the ITU-T when the IETF started its work on a royalty-free codec suitable for use in the Internet. Every time a new codec is deployed, translation between it and all other deployed codecs must be performed available within the network, each participating node must be able to handle the new codec. Translation between codec is expensive and can lead to reduced quality.

This problem seriously constrains the addition of new codecs to the available set, and new codecs are only designed and released when there is a well established need (such as a major difference in functionality).

The application layer of the Internet is, however, predicated on a business model that allows for free or shared software (such as in open source developments), and is only possible with the existence of a royalty free codec. This, together with the specific characteristics of transmission over lossy packet networks comprised requirements equivalent to a major difference in functionality, and led to work in the IETF to specify a new codec.

The complexity, economic, and quality costs associated with inter-working with this new codec will need to be factored into the

deployment model. This, in turn, may adversely effect its adoption and the viably of its use in the Internet.

4.4. MPLS Signaling Protocols

There are three MPLS signaling control protocols used for distributing labels to set up LSPs and PWs in MPLS networks: LDP, RSVP-TE, and GMPLS.

The application domain for each of these is different, and unlike the OAM situation, there is limited requirement for inter-working between the protocols. For example, although one provider may use LDP to set up LSPs while its peer uses RSVP-TE, connectivity between the two providers usually takes place at the IP layer.

It should be noted that the IETF initially worked on another signaling protocol called CR-LDP with variants applicable to MPLS and to GMPLS. The development of this protocol was allowed to progress in parallel with RSVP-TE. However, once it was possible to determine that the solution preferred by the community of vendors and operators was RSVP-TE, the IETF terminated all further work on CR-LDP. No translation function or gateway point interfacing RSVP-TE to CR-LDP was ever proposed.

4.5. IPv4 and IPv6

If there were ever an example of why protocol inter-working is to be avoided if at all possible, it is the transition from IPv4 to IPv6.

The reasons for introducing IPv6 into the Internet are well rehearsed and don't need discussion here. The need for the transition to IPv6 arises from the expansion of the network size beyond the wildest dreams of the creators of the Internet, and the consequent depletion of the IPv4 address space.

This transition has proved to be the hardest problem that the IETF has ever addressed. The invention and standardization of IPv6 was straight-forward by comparison, but it has been exceptionally difficult to migrate networks from one established protocol to a new protocol.

The early assumption that by the time the IPv4 address space was exhausted IPv6 would be universally deployed failed to materialize due to (understandable) short-term economic constraints. Early migration would have been simpler and less costly than the current plans. The Internet is now faced with the considerable complexity of implementing and deploying inter-working functions.

If anything can be learned from the IPv4/IPv6 experience it is that every effort should be applied to avoid the need to migrate or jointly operate two protocols within one network. Adding to the mix a number of issues caused by OAM inter-working of MPLS, one of the Internet's core protocols, would be most unwelcome and would complicate matters still further.

5. Other Examples of Inter-Work Issues

5.1. SONET and SDH

SONET and SDH were defined as competing standards that basically provided the same functionality (simultaneous transport of multiple circuits of differing origin within a single framing protocol). Eventually, SONET was adopted in the U.S., Canada, and Japan, and SDH in the rest of the world.

Significant confusion resulted from this situation. Equipment manufacturers needed to select the market segment they intended to address. The cost of chipsets for a limited market increased. Service providers needed to consider the merits of the two standards and their long-term role in the industry when examining their investment options. Only a limited number of equipment manufactures were available for selection.

A massive interworking function had to be implemented in order to provide global connectivity (e.g., through U.S. and Europe) and this resulted in an increase in operational overhead. As SONET was considered to be the variant, interworking had to be performed before the SDH-based segment was reached.

5.2. IEEE 802.16d and IEEE 802.16e

IEEE 802.16d and IEEE 802.16e were two different, incompatible iterations of the WiMAX standards. In addition to the issues described for SONET/SDH, developers who implemented IEEE 802.16a found that they could not re-use their equipment design when developing the IEEE 802.16d variant. This increased the cost of development and lengthened the time to market.

5.3. CDMA and GSM

CDMA and GSM are two competing technologies for mobile connectivity. In addition to all the undesirable effects described above, the existence of these two technologies adversely affected customers who used roaming when overseas. Sometimes, customers were obliged to obtain an additional device from their service providers in order to roam when travelling abroad (for example, when travelling from Europe

to the U.S).

6. Potential Models For Coexistence

This section expands upon the discussion in Section 3 by examining three questions:

- What does it mean for two protocols to be incompatible?

- Why can't we assume that the two solutions will never coexist in the same network?

- What models could we support for coexistence?

6.1. Protocol Incompatibility

Protocol incompatibility comes in a range of grades of seriousness. At the most extreme, the operation of one protocol will prevent the safe and normal operation of the other protocol. This was the case with the original T-MPLS where MPLS labels that could be used for data in a native MPLS system were assigned special meaning in T-MPLS such that data packets would be intercepted and mistaken for OAM packets.

A lesser incompatibility arises where the packets of one protocol are recognized as "unknown" or "not valid" by implementations of the other protocol. In this case the protocols rules of one protocol require the packets of the other protocol to be discarded and may result in the LSP or PW being torn down.

The least level of incompatibility is where the packets of one protocol are recognized as "unknown" by implementations of the other protocol. In this case the protocols rules of one protocol allow the packets of the other protocol to be ignored, and they are either silently discarded or forwarded untouched.

These are obvious issues with all of these grades of incompatibility ranging from disruption or corruption of user data, through connection failure, to the inability to provide end-to-end OAM function without careful planning and translation functions.

6.2. Inevitable Coexistence

Networks expand and merge. For example, one service provider may acquire another and wish to merge the operation of the two networks. This makes partitioning networks by protocol deployment a significant issue for future-proofing commercial interactions. Although a network operator may wish to present difficulties to disincentivize

hostile take-over (a poison pill) most operators are interested in future options to grow their networks.

As described in Section 3.2, MPLS is a convergence technology. That means that there is a tendency for an ever-increasing number of services to be supported by MPLS, and for MPLS to be deployed in an increasing number of environments. It would be an unwise operator who deployed a high-function convergence technology in such a way that the network could never be expanded to offer new services or to integrate with other networks or technologies.

As described in Section 3.3, there is a requirement for end-to-end OAM. That means that where LSPs and PWs span multiple networks, there is a need for OAM to span multiple networks.

All of this means that, if two different OAM protocol technologies are deployed, there will inevitably come a time when some form of coexistence is required, no matter how carefully the separation is initially planned.

6.3. Models for Coexistence

Two models for co-existence can be considered:

- An integrated model based on the "ships-in-the-night" approach. In this model, there is no protocol translation or mapping. This model can be decomposed as:

  - Non-integrated mixed network where some nodes support just one protocol, some support just the other, and no node supports both protocols.

  - Partial integration where some nodes support just one protocol, some support just the other, and some support both protocols.

  - Fully-integrated dual mode where all nodes support both protocols.

- An "island" model where groups of similar nodes are deployed together. In this model there may be translation or mapping, but it is not always required. This model can be further decomposed:

  - "Islands in a sea" where connectivity between islands of the same type is achieved across a sea of a different type.

  - "Border crossings" where connectivity between different islands is achieved at the borders between them.

6.3.1. The Integrated Model

The integrated model assumes that nodes of different capabilities coexist within a single network. Dual-mode nodes supporting both OAM solutions may coexist in the same network. Interworking is not required in this model, and no nodes are capable of performing translation or gateway function (see Section 6.3.2 for operational modes including translation and gateways).

In this model, protocol messages pass "as ships in the night" unaware of each other, and without perturbing each other.

As noted above, there are several sub-models.

6.3.1.1. Mixed Network Without Integration

In a mixed network with no integration some nodes support one protocol and other nodes support the other protocol. There are no nodes that have dual capabilities.

All nodes on the path of an LSP or PW that are required to play an active part in OAM must support the same OAM protocol. Nodes that do not support the OAM protocol will silently ignore (and possibly discard) OAM packets from the other protocol, and cannot form part of the OAM function for the LSP or PW.

In order to provision an end-to-end connection that benefits from the full OAM functionality, the planning and path-computation tool must know the capabilities of each network node, and must select a path that includes only nodes of the same OAM protocol capability. This can result in considerably suboptimal paths, and may lead to the network being partitioned. In the most obvious case, connectivity can only be achieved between end-points of the same OAM capability. This leads to considerable operational complexity and expense, as well as the inability to provide a fully-flexible mesh of services.

In the event of dynamic network changes (such as fast reroute) or if misconnectivity occurs, nodes of mismatched OAM capabilities may become interconnected. This will disrupt traffic delivery because end-to-end continuity checks may be disrupted, and it may be hard or impossible to diagnose the problem because connectivity verification and route trace function will not work properly.

6.3.1.2. Partial Integration

In a partially integrated network, the network in Section 6.3.1.1 is enhanced by the addition of a number of nodes with dual capabilities. These nodes do not possess gateway or translation capabilities (this

is covered in Section 6.3.2), but each such node can act as a transit point or end-node for an LSP or PW that uses either OAM protocol.

Complexity of network operation is not eased, but there is greater connectivity potential in the network.

6.3.1.3 Dual Mode

Dual mode is a development of partial integration described in Section 6.3.1.2 such that all nodes in the network are capable of both OAM protocols. As in that Section, these nodes do not possess gateway or translation capabilities (this is covered in Section 6.3.2), but each such node can act as a transit point or end-node for an LSP or PW that uses either OAM protocol. Thus, every LSP or PW in the network can be configured to use either of the OAM protocols.

However, it seems unlikely that an operator would choose which OAM protocol to use on a per LSP or per PW basis. That would lead to additional complexity in the management system and potential confusion if additional diagnostic analytics need to be performed. This mode increases the complexity of implementation, deployment, and operation without adding to the function within the network (since both OAM solutions provide the same level of function), so this mode would not be selected for deployment except, perhaps, during migration of the network from one OAM protocol to the other.

6.3.2. Island Model

In the island model, regions or clusters of nodes with the same OAM capabilities are grouped together. Tools to interconnect the technologies are deployed based on layered networking or on interworking between the protocols. These lead to the two sub-models described in the Sections that follow.

6.3.2.1. Islands in a Sea

One way to view clusters of nodes supporting one OAM protocol is as an island in a sea of nodes supporting the other protocol. In this view, tunnels are used to carry LSPs or PWs using one OAM type across the sea and into another island of a compatible OAM type. The tunnel in this case is an LSP utilizing the OAM protocol supported by the nodes in the sea. Theoretically an island can be as small as one node, and the strait between two islands can be as narrow as just one node.

Layering in this way is an elegant solution to operating two protocols simultaneously and is, of course, used to support different technologies (such as MPLS over optical). However, in such layering

deployments there is no simple integration of OAM between the layers, and the OAM in the upper layer must regard the tunnel as a single hop with no visibility into the OAM of the lower layer. Diagnostics within the upper layer are complicated by this "hiding" of the nodes along the path of the tunnel in the lower layer

In the scenarios described so far, both ends of each connection have been placed in islands of compatible OAM types. It is possible to achieve connectivity between a node in an island and a node in the sea if the end-point in the sea has dual capabilities (i.e., can be viewed as a single-node island).

A number of islands may lie along the path between end-points necessitating the use of more than one tunnel. To further complicate matters, the islands may lie in an inland sea so that in is necessary to nest tunnels.

Regardless of the scenario, operating tunnels/layers adds to the management complexity and expense. Furthermore, it should be noted that in an MPLS network there is often a call for any-to-any connectivity. That is, any node in the network may need to establish an LSP or a PW to any other node in the network. As previously noted, the end-points of any LSP or PW must support the same OAM type in the islands-in-a-sea model, so this tends to imply that all, or nearly all, nodes will end up needing to support both OAM protocols.

The use of tunnels can also degrade network services unless carefully coordinated. For example, a service in the upper layer may be provisioned with protection so that a working and backup path is constructed using diverse paths to make them robust against a single failure. However, the paths of the tunnels (in the lower layer) are not visible to the path computation in the upper layer with the risk that the upper layer working and is protection paths share a single point of failure in the lower layer. Traffic engineering techniques have been developed to resolve this type of issue, but they add significant complexity to a system that would be a simple flat network if only one OAM technology was used.

6.3.2.2. Border Crossings

Instead of connecting islands with tunnels across the sea, islands of different types can be connected direct so that the LSP or PW transits the series of islands without tunneling. In this case protocol translation is performed each time the LSP/PW crosses a border between islands that use a different OAM protocol.

In principle this makes for a straight-forward end-to-end connection.

However, protocol translation presents a number of issues as described in Section 3. The complexity is that in planning the end-to-end connection, gateways with protocol translation capabilities must be selected to lie on the path.

7. The Argument For Two Solutions

The decision to define and develop an alternative MPLS-TP OAM solution was based on several assertions:

- The IETF solution is taking too long to standardize
- Commonality with Ethernet solutions is beneficial
- There are two different application scenarios
- There is no risk of interaction between the solutions
- The market should be allowed to decide between competing solutions.

The following sections look briefly at each of these claims.

7.1. Progress of the IETF Solution

The MPLS-TP OAM work carried out within the IETF is the product of joint work within the IETF and ITU-T communities. That is, all interested parties share the responsibility for progressing this work as fast as possible. Since the work is contribution-driven, there is no reason to assume that consensus on the technical content of the work could be reached any faster.

Opening discussions on a second solution seems certain to increase the work-load, and will not will only slow down the speed at which consensus is reached.

It must be noted that, at the time of writing, work on MPLS-TP OAM within the IETF is very nearly complete with only one Internet-Draft remaining to be completed by working group.

7.2. Commonality with Ethernet OAM

Ethernet can be used to build packet transport networks and so there is an argument that Ethernet and MPLS-TP networks will be operated as peers. Examining the issues of end-to-end connections across mixed networks, many of the same issues as discussed in Section 6 arise. If a peer networking gateway model (see Section 6.3.2.2) is applied there is a strong argument to making the OAM technologies as similar as possible.

While this might be a valid discussion point when selecting the single OAM solution for MPLS-TP, it is countered by the need to achieve OAM consistency between MPLS and MPLS-TP networks. One might

make the counter argument that if there is a strong need to make MPLS-TP as similar as possible to Ethernet, it would be better to go the full distance and simply deploy Ethernet.

Furthermore, the approach of a second MPLS-TP OAM protocol does not resolve anything. Since MPLS-TP is not Ethernet, a gateway will still be needed, and this would constitute a second MPLS-TP OAM so additional gateways or interworking functions will be needed because coexistence is inevitable as described in the rest of this document.

7.3. Different Application Scenarios

It has been suggested that two different applications of MPLS-TP exist: Packet Switched Network (PSN) and Packet Transport Network (PTN). These applications have not been documented in the IETF and most of the support for the idea has come in discussions with a little documentation in the ITU-T [TD522].

One of stated differences between these applications lies in the OAM tools that are required to support the distinct operational scenarios. The OAM used in a PSN should be similar to that used in an MPLS network (and so should be the MPLS-TP OAM defined in the IETF) while the OAM used in a PTN should provide the same operational experience to that found in SONET/SDH and OTN networks.

The basic MPLS-TP OAM requirements in [RFC5654] make this point, saying:

> Furthermore, for carriers it is important that operation of such packet transport networks should preserve the look-and-feel to which carriers have become accustomed in deploying their optical transport networks, while providing common, multi-layer operations, resiliency, control, and multi-technology management.

Thus, the look-and-feel of the OAM has been a concern in the design of MPLS-TP from the start, and the solutions that have been defined in the IETF were designed to comply with the requirements and to provide operational behavior, functionality and processes similar to those available in existing transport networks. In particular, the toolset supports the same controls and indications as those present in other transport networks, and the same management information model can be used to support the MPLS-TP OAM tools (in areas where the technology type is irrelevant).

It is important to note that the operational look-and-feel does not determine the way in which OAM function is achieved. There are multiple ways of achieving the required functionality while still providing the same operational experience and supporting the same

management information model. Thus, the OAM protocol solution does not dictate the look-and-feel, and the demand for a particular operational experience does not necessitate the development of a second OAM protocol.

7.4. Interaction Between Solutions

Section 3 of this document discusses how network convergence occurs and indicates that where two MPLS-TP solutions exist they are, in fact, very likely to appear either in the same network or at gateways between networks.

Indeed, since nodes offering either solution are likely to both be branded as "MPLS-TP", and since network interoperation (as described in Section 6) demands the existence of some nodes that are either dual-mode or act as protocol translators/gateways, there is considerable likelihood of the two OAM solutions interacting through design or through accident. When a node is capable of supporting both OAM protocols, it must be configured to support the correct protocol for each interface and LSP/PW. When a device has interfaces that offer different MPLS-TP OAM function, the risk of misconfiguration is significant. When a device is intended to support end-to-end connections, it may need to translate, map, or tunnel to accommodate both protocols.

Thus, the very existence of two OAM protocols within the common MPLS-TP family, makes copresence and integration most likely.

7.5. Letting The Market Decide

When two technologies compete it is common to let the market decide which one will survive. Sometimes the resolution is quite fast, and one technology dominates the other before there is widespread deployment. Sometimes it takes considerable time before one technology overcomes the other, perhaps because one technology has become entrenched before the emergence of the other, as in the in the case of MPLS replacing ATM. In more cases, however, the market does not select in favor of one technology or the other - as in many of the cases described in Sections 4 and 5 of this document, sometimes both technologies continue to live in the network.

Letting the market decide is not a cheap option. Even when the resolution is rapid, equipment vendors and early adopters pay the price of both technologies. When it takes longer to determine which technology is correct there will be a period of coexistence followed by the need to transition equipment from the losing solution to the winning one. In the cases where no choice is made, the network is permanently complicated by the existence of the competing

technologies.

In fact, the only time when allowing the market to decide can be easily supported is when the competing technologies do not overlap. In those cases, for example different applications in the user-space, the core network is not perturbed by the decision-making process and transition from one technology to the other is relatively painless. This is not the case for MPLS-TP OAM, and coexistence while the market determines the correct approach would be expensive, while the necessary transition after the decision has been made would be difficult and costly.

8. Security Considerations

This informational document does not introduce any security issues.

However, it should be noted that the existence of two OAM protocols raise a number of security concerns:

- Each OAM protocol must be secured. This leads to the existence of two security solutions each needing configuration and management. The increased complexity of operating security mechanisms tends to reduce the likelihood of them being used in the field and so increases the vulnerability of the network. Similarly, the existence of two security mechanisms raises the risk of misconfiguration.

- One OAM protocol may be used as a vector to attack the other. Inserting an OAM message of the other OAM protocol onto a link may cause the service to be disrupted and, because some nodes may support both OAM protocols, it may be possible to cause the disruption at a remote point in the network.

- Securing a network protocol is not a trivial matter for protocol designers. Duplicating design effort is unlikely to result in a stronger solution and runs the risk of diluting the effort and creating two less-secure solutions.

9. IANA Considerations

This informational document makes no requests for IANA action.

10. References

10.1. Normative References

[RFC5654] Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, February 2009.

[RFC5860] Vigoureux, M., Betts, M., and D. Ward, "Requirements for OAM in MPLS Transport Networks", RFC 5860, April 2009.

10.2. Informative References

[Y.Sup4] "ITU-T Y.1300-series: Supplement on transport requirements for T-MPLS OAM and considerations for the application of IETF MPLS technology", Y.Sup4, 2008.

[OAM-Overview] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Mechanisms", draft-ietf-opsawg-oam-overview, work in progress.

[RFC5317] Bryant, S. and L. Andersson, "Joint Working Team (JWT) Report on MPLS Architectural", RFC 5317, February 2009.

[RFC4929] Andersson, L. and A. Farrel, "Change Process for Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Protocols and Procedures", RFC 4929, June 2007.

[TD7] "TD7 (WP3/SG15): IETF and ITU-T cooperation on extensions to MPLS for transport network functionality", ITU TD7 (WP3/SG15), December 2008.

[TD522] "TD522 (Q12/SG15): Clarification of the PTN/solution X environment", ITU TD522 (Q12/15), February 2011.

[LS26] "LS: Cooperation Between IETF and ITU-T on the Development of MPLS-TP", COM15-LS26-E, December 2008.

Authors' Addresses

Nurit Sprecher
Nokia Siemens Networks
3 Hanagar St. Neve Ne'eman B
Hod Hasharon, 45241
Israel

Email: nurit.sprecher@nsn.com

Kyung-Yeop Hong
300 Beaver Brook Road
Boxborough,
Massachusetts 01719
USA

Email: hongk@cisco.com