Having reviewed the draft text of G.8113 and G.8114, and the various comments that have been drawn to our attention the IAB Ad Hoc T-MPLS team believes that the following issues need to be addressed before these documents can be given approval for publication.

Some of the issues below are duplicated, although most of the duplicates address the concern from a different angle, and hence need to be independently resolved.

In the week since we met to discuss this problem at IETF and produced a plan to resolve the issues we have not yet had time to trace down all of the reference trails and ensure that there are no issues raised by the reference material. We will continue to look at those and may surface more issues next week.

Regards

Stewart

Issue	Description
1	Our review of ITU-T Recommendation Y.1372/G.8113 (Y.17tor) and ITU-T Recommendation Y.1373/G.8114 (Y.17tom) indicates that the usage, behavior and semantics of IETF MPLS Reserved Label 14 as defined by RFC 3429 are being modified. For example, Section 10 of Y.1373/G.8114 defines a number of new OAM PDU types with associated functionality expanded upon throughout the document. Label 14 as specified in RFC 3429 points to Y.1711 (2002). Its use by the updates to Y.1711 (2002) (e.g., Y.1711 (2004), G.8112, Y.1373, etc.) MUST be addressed by the update to RFC 3429. Without such a document, RFC 3429 will become inconsistent in its description of how IETF MPLS Reserved Label 14 is being used, the stated policy for allocation of IETF MPLS Reserved Labels will not have been met and IETF Process will not have been followed.
2	Modification to RFC 3429 and documents that run on codepoints that RFC3429 allocates MUST be modified using the process defined in RFC 4929, "Change Process for Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Protocols and Procedures."
3	If ITU-T intends to publish G.8114, then it should be published without specifying the use of a specific reserved label and it should state that the method of identifying an OAM packet is for further study. The requirements for the new OAM alert mechanism should be brought to the IETF using the MPLS change process. The IETF Standard Process should then be used to define an appropriate alert mechanism for this protocol.

4	Draft new Recommendation Y.1373 only specifies the T-MPLS OAM packet formats, syntax and semantics of T-MPLS OAM packet fields. It does not contain procedures, forwarding, processing of MPLS frames and error handling. It is not self contained. For implementation of T-MPLS OAM, both OAM PDUs format specified in Y.1373 and procedures are required. The draft does not point to any Recommendation for procedures. For version 1 of T-MPLS, procedures are defined in Recommendation G.8121. It only supports OAM capabilities defined in Y.1711. For updating RFC 3429 which T-MPLS document will be referred? It is difficult to see how ITU-T can proceed with publication of G8114 until the complete set of corresponding elements of procedure are also ready for publication. This needs to be done to ensure compatibility.
5	Are Y.1711 and Y.1373 are different protocols or one is merely an extension to the other. The attached Annex to the contribution (SG 13 contribution from Juniper and Cisco to Jan 2008 meeting) provides a table comparing both protocols. From the table it is clear they are different protocols. In updating RFC 3429, MPLS Reserved label 14 will be assigned to which protocol - Y.1711 or T-MPLS OAM (Y.1373 and updated G.8121)? If it is allocated to T-MPLS what will happen to Y.1711? Whatever approach is taken it is important that the Y.1711 MPLS OAM is not deprecated without the application of due diligence to ensure that existing Y.1711 implementations can continue to operate correctly.
6	Is there an amendment proposed to change the reference to Y.1711 in G8110? If not then Y.1711 remains the MPLS OAM. Current MPLS OAM Recommendations Y.1711 (02/2004) – (in Force) specifies Operation and Maintenance for MPLS Networks Note that G8110/Y.1370 MPLS Layer network Architecture (01/2005) (In Force) refers to Y.1711 Clause 7.3.1 Reserved Label, Table 1 Reserved Label Values: 14 "OAM Alert Label – Label for MPLS OAM packets as described in ITU-T Rec. Y.1711. It is not used in G.809 model"
7	Annex B of Y.1373 provides "G.8112 (2006) Compatibility Considerations". This contribution identifies additional compatibility issues have been identified (see end of document) which need to be addressed.
8	The performance parameters specified in section 8 are not aligned with IP performance parameters defined in Y.1540.

	Section 8 of Y.1373 needs to be updated based on the liaison received from SG 12. Note to other reviewers – there are many RFCs that define IP performance metrics, we need to check for compatibility with those: RFC 2330, RFC3393,RFC3432,RFC4144
9	The MEL bits identified in the Common PDU Format and G.8112 compatible format of Y.1373 use the IETF specified MPLS EXP bit field. The use of MEL bits in place of EXP bits is not specified in the MPLS RFCs and needs to be resolved.
10	The use of MEP bits instead of TTL bits is not specified in the MPLS RFCs. The MEP bits need to be moved to another position in the OAM message.
11	Recommendation Y.1372 requires ALL OAM messages be authenticated, but Recommendation Y.1373 does not comply with these requirements. This needs to be resolved.
12	G8113 – Section 7 bullet 1 says: “T-MPLS provides a unique connection-oriented layer network and hence there will be failure modes that are only relevant to T-MPLS.” How can there be failure modes in T-MPLS to do not occur in PWE3/MPLS? The statement that these mode occur implies that T-MPLS departs from the PWE3/MPLS architecture. We need to identify these unique failure modes and confirm no departure from the IETF MPLS design. Some draft copies of G.8113 strike the term “unique”, but the issue remains – what are these failure modes? We a list of these modes and time to verify that they are not a departure from the IETF design. If there is a departure it will need to be resolved prior to publication of these documents.
13	G8113 – Section 7 bullet 3 says: Operators need the ability to determine T-MPLS availability and network performance, noting that network performance metrics are only meaningful when the connection is in the available state. This information may also be used for accounting and billing purposes to ensure that customers are not inappropriately charged for degraded services or service outages. Accounting and billing has significant implications for integrity and security. Where are these aspects of the OAM requirements specified, and how does the OAM protocol address them?
14	G8114 runs on the same label as Y.1711, but at significantly higher packet rate. There appears to be no requirement in G8113, that a G8114 system detect that it is misconnected to a Y.1711 system and either shut down or run at a packet rate that will not overload it.
15	G8113 Section 7 para 6 states:

	Ensure that any defect that results in misdirected customer traffic is detected/diagnosed and leads to appropriate actions, e.g., squelching of traffic where relevant. This is a catch all requirement – where is this broken down into a checklist that the protocol designer can use to ensure that they have fully addressed this problem? Furthermore it is not clear from the elements of procedure how this works. We understand – but have not seen in writing - that the MEP is used. However the MEP procedure seems to be related to the use of the label 14 TTL in a way that is not approved by the IETF MPLS WG.
16	G8113 Section 7 para 6 stated: Proactive maintenance actions also help drive down operating complexity by minimizing the opportunity for incorrect defect diagnosis, and (like the previous item) they also promote customer trust of an operator. This is a non-technical assertion made without apparent justification. It could be strongly argued that the IETF approach of using signalled connections minimises defect occurrence.
17	G8113 Section 8 para 5 stated: The following anomalies should be automatically detected and corresponding defect states, with well-defined entry/exit criteria and appropriate consequent actions, should be defined: – loss of connectivity; – misconnections; – unintended self-replication (e.g., looping); – lost packets; – mis-inserted packets (e.g., mis-insertion of a packet into a T-MPLS connection). Whilst it is clear that the OAM will identify connection failure, it is not clear that the OAM is able to identify that this is a loop – how does loop detection work? The only way that a loop can replicate packets is if there is multicast in the system. Where is multicast OAM loop detection mechanism support defined, and where are the requirements defined? How do you deduce mis-inserted packets? For example how would the OAM deduce that two TDM streams have not been cross connected in the network core? We assume that you propose to deduce a number of these defects from exact packet counting and synchronous OAM injection. However that requires onerous demands the data path that the IETF has been reluctant to sanction

	for IP/MPLS because they do not scale. A proposal for congestion monitoring that required synchronous packet support was discussed in the PWE3 WG and consensus was to reject this as impractical in many forwarding designs.
18	G8113 Section 8 para 27 stated: T-MPLS OAM should provide mechanisms to detect mis-delivered packets. The OAM can detect mis-connection, but it cannot detect mis-delivery without data-plane support. Even then if there is a change to the forwarding path for OAM packets, this cannot be assured. Please remember that the presence of synchronous traffic in the network may defeat packet counter mechanisms.
19	G8113 Section 8 para 15 stated: T-MPLS connectivity status assessment should not be dependent on the dynamic behaviour of client-layer traffic. You can do CV without data, but loss rate needs data which by definition is dynamic.
20	G8113 Section 8 para 19 stated: T-MPLS OAM shall be interoperable with the OAM described in G.8112 (2006). This needs further investigation given the concerns about the incompatibility of the various label 14 based OAM mechanism described elsewhere in this document.
21	G8113 Section 8 stated: 25) T-MPLS OAM should provide mechanisms to ensure that unauthorized access is prevented from triggering any service provider/network operator T-MPLS OAM function. 28) T-MPLS OAM messages shall be authenticated. This requires that OAM authentication is mandatory on all OAM messages, but this is not reflected in the design specs.
22	G8113 Section 10.1 requires traceroute – how is this provided? .
23	The Security Sections says: The following items are related to security aspects: • item 6) of clause 7; • items 5), 25), 26), 27) and 28) of clause 8. That is not a threat analysis, or a direction to the protocol specifications to provide a security solution! Please provide a security section. Can we draw your attention to section 10 of draft-ietf-pwe3-vccv-15.txt
24	Given that the proposed label 14 behavior which seems to require a change to the P router behaviour - which the IETF strongly objects to - how do you verify that the OAM path is the same as the data path and hence rely on deductions based on fate sharing.

25	How does this mechanism detect that packets are being delivered to a multicast client that has no authority to receive them?
26	G8113-Section 7-para 2 It should be noted that there is no fixed hierarchy in T-MPLS and, in theory (at least), the nesting depth can be unlimited. However there seems to be a requirement to search the complete (arbitrary size) stack. This is not scalable. What about the limits placed on the system by the size of the MEG level field?
27	G8113-Section 7-para 5 Improve the availability performance. Experience in designing the Internet led us to conclude that you could not get the highest levels of availability without dynamic routing protocols. Similarly with para 7 Minimize the number of defects that are not detected automatically before a customer reports the problem. The surest way to ensure connectivity is through dynamic protocols.
28	G8113 Section 8 para 7 OAM functions should provide means to detect anomalies that impact the transport of user traffic in the network. T-MPLS OAM packets should be forwarded on the same route as the T-MPLS user packets are forwarded. Why is this not a MUST be forwarded? Why is there not an additional requirement that OAM packets must be indistinguishable from data packets to an on-path LSR to ensure that data and OAM packets share the same fate?
29	G8113 Section 8 para 8 A defect event in a given layer network should not cause multiple alarm events to be raised, nor cause unnecessary corrective actions to be taken, in any client-layer network. Given that TMPLS can be arbitrarily nested how does this work – i.e. how does a server layer know that it should initiate the repair rather than wait for a server layer to initiate the repair? What are the repair delay implications and the timer setting implications of this?
30	G8113 Section 8 para 9 T-MPLS OAM functions should be simple and easily configured (ideally automatically) to allow efficient scaling to large network sizes. We cannot find the definition of the OAM configuration signalling protocol. Where is it defined?

31	G8113 Section 8 para 9 The use of T-MPLS OAM functions should be optional for the operator. This sounds very dangerous in a statically configured network. Surely CV bound to a mechanism to inhibit traffic under mis-configuration or equipment malfunction should be mandatory?
32	G8113 Section 8 para 9 The design of T-MPLS OAM functions should ensure that a T-MPLS equipment that does not support T-MPLS OAM functions will be able to silently discard the T-MPLS OAM packets addressed to a T-MPLS (termination) connection point in this equipment, or let T-MPLS OAM packets pass through transparently without disturbing the user traffic or causing unnecessary actions. Surely the design needs to be such that a TMPLS node stops delivering unwanted packets to a device that does not understand them? Surely that should be discard and log so that the fault can be diagnosed?
33	G8113 Section 8 para 22 T-MPLS OAM shall provide a mechanism for Maintenance Communication Channels MCC. What are the security implications of this? What about the impact on the OAM message rate? Can this interfere with the primary function of the OAM? What are the message priority considerations? Same question for G8113 Section 8 para 23 – user extensions.
34	G8113 One-way packet delay measurements. With what accuracy? What are the time availability requirements/implications?
35	G.8114 specifies that an OAM packet is preceded with an MPLS header, with label 14 and the S-Bit=1. So clearly label 14 is at the bottom of the stack. Furthermore G8114 6.5 says : A receiving MIP should be able to identify OAM packets it has to process by recognizing them as being OAM packets (due to the presence of the OAM Alert label value) with the MEL field equal to 0, a Function type that is supported by the MIP and a data-plane identifier that identifies this MIP. MPLS behavior is either SWAP, POP, or PUSH (or some combination thereof). In no case do we swap; but also look at the next label (or any other label in case it happens to be 14. Furthermore, G.8110 tells us that T-MPLS can be applied recursively,

	that is tunnels may be nested. So that may imply looking even further in to the MPLS stack. This section implies a behavior not recognized in IETF MPLS. This section and all procedures and data types that rely on it should be removed from G8114.
36	G.8114 5.5 defines a MIP. This seems to be an LSR that interacts with an OAM packet that it is forwarding. This is a concept that is foreign to MPLS (other than via a router alert which is not a mechanism used by this protocol), and either it should be clarified how an LSR provides the required interaction within the capabilities provided at MPLS LSR conforming to RFC3031, or the term should be removed from the document.
37	G.8114 5.6 defines the MEL levels as 0..7, but later in the section it says that MEL 7 is dropped, so MEL 7 is not a MEL level and the MEL levels are 0..6 Also G8113 says that there can be arbitrary nesting. Isn't the restriction to 6 levels a contradiction with the requirements?
38	G.8114 6.1 says that Appendix 1 describes different network scenarios. Appendix 1 is not very clear, but does not seem to provide more than one scenario and so does not provide a sufficiently rich set of examples to fully understand and review the mechanisms proposed in this document.
39	G.8114 7.1.1 talks about mismatched CV periods. It needs to specify the action when there is a mismatch.
40	G.8114 section 7.4 needs to be edited to remove all references to MIPS performing a loopback since this seems to be based on a mechanism that in not conformant to the IETF MPLS architecture.
41	G.8114 7.5 talks about the continued transmission of packets to a MEP that is broken? At what rate? An exponential backoff should normally be applied in these circumstances, but does not seem to be specified.
42	G.8114 8.2 needs some text warning the reader about asymmetric path effects lest they are tempted to take one way = two way /2
43	Below Fig 10.2-6 there is a comment that 3 long words will be ignored. In Y.1711 they are specified as must be zero. What is the DEFINED behavior of a Y.1711 system receiving a type 1 function with these bits not zero.
44	G.8114 10.4 and subsections. References to MIPS needs to be removed – see issue 40
45	G.8114 section 11, needs to be replaced by a full threat analysis of the new features introduced by G.8114
46	We need more time to check out ref chain incl 1588 and 8121 2004. We will

	provide this as soon as possible.
47	G.8114 Section 8.1 defines proposes a mechanism to measure the service packet loss by inserting an OAM packet and that conveys the number of packets sent since the last OAM performance monitoring packet. This approach is a plane violation in that the OAM plan needs to get a packet count from the data plane and to stall the data plan in order to insert an OAM packet in the data packet flow. This is acknowledged in G.8114 to be a difficult problem needing hardware assistance. This method is not conformant to the IETF PWE3/MPLS architecture, since it places constraints on the relative sequencing of packets from identifiable different flows. There is however another approach that does not violate the isolation of the data and OAM planes and uses a mechanism that is fully specified in the IETF PWE3 design. T-MPLS carries the service packets using a PWE3 defined pseudowire encapsulation. All PWE3 defined pseudowire control words include a sequence number intended for the data plane to detect missing and out of order packets. When sequence number support is enabled the data plane is able to count the number of missing data packets and report this to the MIB. By sampling the MIB the management system is able to determine the packet loss rate. The losses are thus determined by the data plane using information added to the data plane for this explicit purpose in way that conforms to the IETF PWE3/MPLS architecture and which does not require interaction between data plane and OAM plane or the precision insertion of an OAM packet. Section 8.1 should therefore be removed and replaced by a mechanism than is based on the use of the existing data plane counter mechanism.
48	Y.1372 (Y.17tor) Clause 4 'abbreviations and acronyms' includes SSM (sync status message). This is not mentioned (that we can see) in the rest of the document so it has no need to be here. However, of more substantive relevance (and to Y.17tom) is the fact that the carriage of both the actual timing information (network clock) and the SSM protocol which refers to it must use a fate-sharing 'bottom-of-stack' solution, i.e. only p2p section layer information signals. These must not be carried by the OAM of path layer networks (any, not just T-MPLS).
49	Y.1372 (Y.17tor) **Clause 6 'reference networks' says:** "This Recommendation specifies the requirements for OAM functions that are applied to point-to-point and point-to-multipoint T-MPLS connections." It would be useful to be clarified what is meant/defined by a 'connection' here. Can we confirm that it is understood that the definition of a connection is a construct that: - has a single source of traffic units - has no internal choice of routing of traffic units - has no re-ordering of traffic units

	Note - The latter point here implies that a connection must only have a single traffic/QoS class. We therefore wish to clarify whether T-MPLS constructs respect this requirement. If not, we suggest guidance is sought from the Q12/SG15 experts (unified modelling) on what we should call/regard T-MPLS constructs.
50	Y.1372 (Y.17tor) **Clause 7 item 2 says**: "The T-MPLS nesting capability allows the creation of multiple T-MPLS layer network instances in their own right, within the framework of the T-MPLS technology. It should be noted that there is no fixed hierarchy in T-MPLS and, in theory (at least), the nesting depth can be unlimited" Is each nested T-MPLS instance a layer network in its own right, i.e. fully functionally decoupled (all 3 planes) from other client or server instances of T-MPLS? This needs clarifying (including things like how the S bit is handled when T-MPLS levels are stacked).....and if true, this is also not how MPLS works, so T-MPLS cannot be considered a profiled sub-set of MPLS.
51	Y.1372 (Y.17tor) Clause 8 item 10 says: "The use of T-MPLS OAM functions should be optional for the operator. A Network operator should be able to choose which OAM functions to use and which connections it applies them to." When one uses label-swapping techniques then misconnectivity defects have a chance of propagating all the way to trail termination points. In particular, 'important' traffic may be leaking into 'non-important' traffic and this cannot be detected unless consistent OAM is run on all the connections in the network. So in a label-swapping co-ps mode network the CV function (as a min) is not really optional and should be deployed on all connections, though other OAM functions (especially ‘on-demand’ type) could be considered optional.
52	Y.1372 (Y.17tor) Clause 8 20 says: "T-MPLS OAM should allow for a stateless layer network interworking (also known as service interworking) between T-MPLS (based on G.8110.1) and an Ethernet (based on G.8010) networks." This is a rather naive statement.....peer-partition interworking is much more than just the OAM component of the data-plane! Without going into major detail here (but we can do if required) we must not even countenance the peer-partition interworking of technologies that belong to different network modes as there will be a large functional mismatch across both data-plane and control-plane components. So this statement must be removed.
53	Y.1372 (Y.17tor) Clause 10.1 talks about using 'request/response' on-demand OAM functions for performing diagnostic tests. One has to remember that misconnectivity defects may not be bi-

	directionally symmetric, i.e. whilst there might be a path to allow traffic to leak out of a connection there may not be a return path......indeed, it is very likely there will not be (this is not like a cl-ps mode network)! In any case, one should not assume defects know they need to be so well behaved and bi-directionally symmetric, so in general 'request/response' OAM in co mode layer networks should only be used under the assumption that the network is defect-free, e.g. as a quick check of connectivity at provisioning or a RTD test.
54	Y.1372 (Y.17tor) Clause 10.1 talks about *unidirectional* connections and yet also refers to 'dual-ended packet loss measurements'. If we are understanding this correctly, surely dual-ended loss measurements can only apply to bi-directional connections.
55	For consistency in the protocol definition in the G.8121 needs to define support of G8114's MEPs for tandem connections and MIPs. Where is this documented? We need to review this so that we can be sure that the design is consistent and may provide further comments on this review. Understanding the required LSR behavior to support MEPs for tandem connections and MIPs, and ensuring that this is consistent with the IETF MPLS architecture is a critical and gating factor in the IETF being able to endorse the design.
56	In G.8114, there seems to be technical issues with things like LB since the LBR TTL value is copied from received LBM and may never reach the originator if other MIPs were present between responding node and originating node, defeating the purpose of the LB operation.
57	And in G8114 annex C, it discusses interworking (mapping) ethernet oam and tmpls oam(also not in G8121). We assume the is an Ethernet AC connecting into T-MPLS connectivity in core and the need to signal defects between Ethernet and T-MPLS domains. Or in other terms, if two technologies peer instead of more common client-server relationship in which case it is tunneling instead of mapping. Where is this requirement specified in G.8113? If it is not a specified requirement it should not be called up in G.8114 and should not be a design constraint.
58	The timelines for the update of G.8121 are not aligned as already noted. The scope of the updates in G.8121 is not complete, therefore rendering some of the missing details difficult for the IETF to review and approve the design. We have technical issues with MIP functionality e.g. the TTL in LBR not being a value that would allow reception by LBM originator.

59	The functionality coverage between Y.1711 and G.8114 is quite different with minimal commonality. Also, the issue of positioning of Y.1711 MEPs and G.8114 MEPs such that same requirements to allow transparency for higher MEG is not discussed and IMO not possible since Y.1711 has no visibility into MELs and will block all OAM frames received.
60	With regards to performance monitoring, we have not reviewed all IETF documents, however, one issue is around the definition of Frame Delay Variation which seems to have two approaches, one called IPDV and other called PDV. IPDV deals with delay and variation between subsequent frames while PDV deals with the reference for delay being served by minimum delay in the measurement window. The ITU-T approach is based on IPDV we need to check further to understand where IETF stands on this discussion or metric definition. We must note though that IPDV offers the same information as PDV way of defining FDV or jitter metric.
61	This issue is covered elsewhere, but for clarity we restate it: G8110 is the MPLS Layer Network Architecture and it clearly points to Y.1711 as OAM for MPLS and references the OAM Alert label 14. As far as We know there is no update or Amendment to this Recommendation in the pipe. G.8110.1 which is the Architecture of T-MPLS currently points to Y.1711 but there is a revision that is prepublished that would change that to Y.1373 when approved and of course incorporating Label 14. Bottom line is there is a current inconsistency between G8110 (MPLS) and G8110 (TMPLS) in the ITU-T Recommendations with no indication this is being addressed.

`The following text relates to issue #`

The IETF sets significant store in backwards compatibility of the protocols used in the Internet, of which RFC3429 is one. It is therefore necessary that any extension to RFC3429 be fully backwards compatible. The following extract from a contribution by Juniper to ITU-T therefore needs to be resolved:

One of the issues that is coming up: that Y.1711 and Y.1373 are different protocols or one is merely an extension to the other. The attached Annex to the contribution provide table comparing both protocols. From the table it is clear they are two separate protocols.

Discussion

Recommendation G.8112 specifies as follows:

6.2.1.1   T-MPLS OAM

T-MPLS OAM is specified in [ITU-T Y.1711]. Figure 6-13 illustrates the set of T-MPLS OAM and their format.

The T-MPLS OAM header consists of a T-MPLS shim header with a reserved label value of 14 (OAM Alert). The DL and DT fields in the FDI and BDI OAM packets are set to all-zeroes by the transmitter and are to be ignored by the receiver. The LSP TTSI fields in the FDI and BDI OAM packets are set to all-zeroes by the transmitter and are to be ignored by the receiver. The format of the LSP TTSI field in the FFD/CV OAM packets within the scope of T-MPLS is for further study.

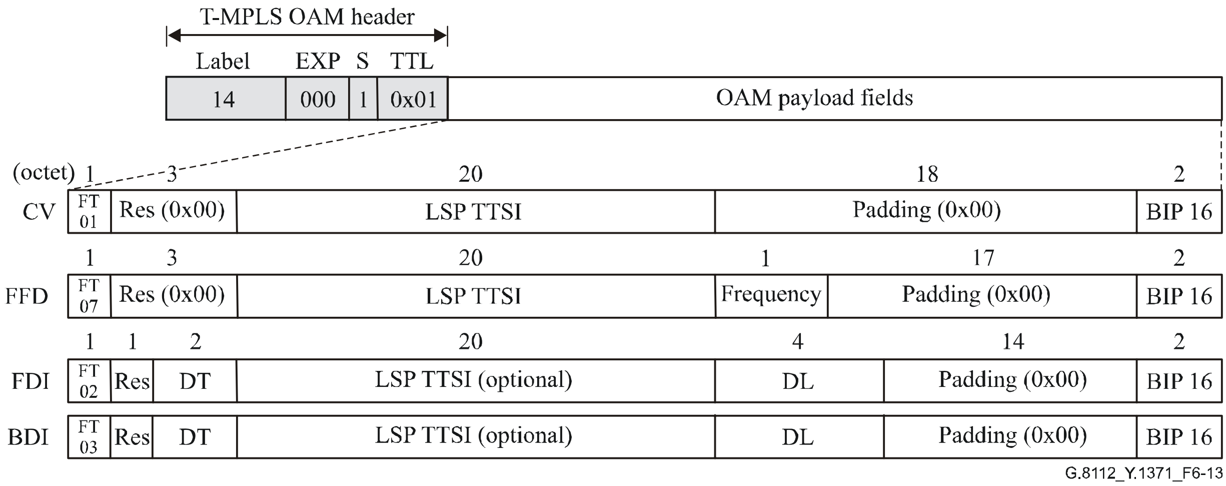

Figure 6-13 – Y.1711 defined T-MPLS OAM

(Note: This figure is actually taken from G.8112)

Procedures for T-MPLS OAM

Recommendation Y.1711 is self contained. It specifies procedures, frame formats, OAM PDUs and error handling. The OAM mechanisms defined in this Recommendation assume common forwarding of the LSP payload and Y.1711 PDUs.

Draft new Recommendation Y.1373 only specifies the T-MPLS OAM packet formats, syntax and semantics of T-MPLS OAM packet fields. It does not contain procedures for forwarding and processing of MPLS frames and error handling. It is not self contained.

For implementation of T-MPLS OAM, both the OAM PDUs format specified in Y.1373 and the procedures are required. The draft Y.1373 does not point to any Recommendation for these procedures.

For version 1 of T-MPLS, the procedures are defined in Recommendation G.8121. G.8121 only supports the OAM capabilities defined in Y.1711. G.8121 has not been updated with the procedures required to support Y.1373. In order to verify the PDUs format, these missing procedures are required.

Backward compatibility with G.8112

Annex B of Y.1373 provides “G.8112 (2006) Compatibility Considerations”. The list only identifies some of the issues. This contribution identifies additional compatibility issues.

2.2.1 TTSI coding

Recommendation Y.1711 defines the TTSI in section 6.1.4. The structure of the LSP Trail Termination Source Identifier (TTSI) is defined by using a 16-octet LSR ID IPv6 address followed by a 4-octet LSP Tunnel ID.

In Recommendation G.8112 section 6.2.1.1 specifies "The format of the LSP TTSI field in the FFD/CV OAM packets within the scope of T-MPLS is for further study".

Y.1373 defines the Data plane identifier as TLV in section 10.1.2.3.

It is not clear how backward compatibility is supported.

2.2.2 Ignoring the information by G.8112

Figure 10.2-6 specifies CV_v1 PDU format processed in a G.8112 (2006) compatible MEP. It further states "A G.8112 (2006) compatible MEP will ignore the information in the horizontal and vertical shaded fields, i.e. octet 6 [bits 1…5], octet 7 and octets 29…40".

G.8112 clearly states what fields should be ignored by the receiver. The fields in the figure 10.2.6 are not specified to be ignored. It is an implementation option how it be treated and as a result this may cause a backwards compatibility problem.

2.2.3 Function types not supported

Section B.4 says Y.1372 will never generate or receive FFD packets. It is the responsibility of management not to enable FFD generation in a G.8112 end-point.

Y.1711 says "The provisioning of FFD on an LSP will provide capability of fast failure detection (by default FFD OAM is not provisioned). It is recommended that FFD will be generated at the LSP ingress at a rate of 20 per second. This will provide failure detection on the LSP in order of 100s of milliseconds".

If this capability is not provisioned how is the capability provided?

What will happen if the FFD is enabled in G.8112?

2.2.4 Frequency of sending CV packets

In G.8112 the CV flow is generated at the LSP's source LSR with a nominal frequency of 1/s and terminated at the LSP's sink LSR.

Section 7.1.1 defines the transmission period. 7 different values are specified for transmission period, the default values are recommended based on the application area for which CC is being used. When a transmission period other than the default value for an application area is used, the behaviour of the intended application is not guaranteed.

To support backward compatibility only one value can be used. How are other capabilities like protection switching supported? FFD has better granularity which is not supported.

2.2.5 Intermediate processing

Recommendation Y.1711 defines in section 6.1.2.

Y.1373 section 5.6 specifies "In order to distinguish OAM packets of nested MEGs, each MEP tunnels incoming OAM packets by incrementing the MEL in the source direction and decrementing it in the sink direction."

Section 6.5 specifies the OAM label TTL field processing, which is different from standard TTL processing.

2.2.6 OAM packet processing

In Y.1711 the OAM packet processing includes:

All OAM packets must have a minimum payload length of 44 octets

BIP16 processing must be performed on all OAM packets prior to being able to reliably pass their payload for further processing. Any OAM packets that show a BIP16 violation upon reception processing should be discarded.

If the G.8112 equipment is between two equipment supporting Y.1373, it does not pass transparently other OAM packets. A large number of packets may be dropped.

3 Proposal

This contribution proposes the following:

This contribution identifies additional compatibility issues. The list only identifies some of the issues. These issues must be resolved for backward compatibility.

The procedures to be used for draft new Recommendation Y.1373 must be specified or a pointer to a Recommendation that provides the procedures must be provided.

Comparison of Recommendation Y.1711 and draft new Recommendation Y.1373

Y.1711	Y.1373
1 Scope (not complete) This Recommendation is designed primarily to support point-to-point and multipoint-to-point explicit routed LSPs (ER-LSPs) with limited applicability to LSPs that employ penultimate hop popping (PHP). The OAM mechanisms defined in this Recommendation assume common forwarding of the LSP payload and Y.1711 PDUs. In some situations this may not be true, such as when the LSP payload is load balanced across a plurality of parallel paths while still appearing as a single trail to the ingress and egress. LSRs introducing variations in connectivity are responsible for ensuring that the availability behaviour of Y.1711 per ingress-egress pair is preserved.	1 Scope (not complete) This Recommendation specifies the mechanisms required to operate and maintain the network and service aspects of the transport MPLS (T-MPLS) layer network. It also specifies the T-MPLS OAM packet formats, syntax and semantics of T-MPLS OAM packet fields. The T-MPLS OAM mechanisms as described in this Recommendation apply to both point-to-point T-MPLS connections and point-to-multipoint T-MPLS connections. The architectural basis for this Recommendation is the T-MPLS architectural specification G.8110.1. The OAM functions of the server layer network used by the T-MPLS layer network are not within the scope of this Recommendation. The OAM functions of the client layer network(s) of the T-MPLS layer network are not within the scope of this Recommendation either. Juniper Comment: It does not say it is the next version of Y.1711. Y.1711 is only reference in some code points etc.
5.3 OAM payload The payload of an OAM packet is composed of the OAM Function Type, the specific OAM function type data and a common BIP16 error detection mechanism. All OAM packets must have a minimum payload length of 44 octets to facilitate ease of processing and to support minimum packet size requirements of current L2 technologies (e.g., Ethernet). This is achieved by padding the specific OAM type data field with all 0s when necessary. All padding bits are reserved for possible future standardization.	There are two PDU formats: Common OAM PDU format G.8112 compatible OAM PDU format

5.4 Handling of errored OAM packets Each OAM packet uses a BIP16 (in the last two octets of the OAM payload area) to detect errors. The BIP16 remainder is computed over all the fields of the OAM payload, including the Function Type and the BIP16 bit positions (which are all pre-set to zero for initial calculation purposes). The BIP16 generator polynomial is G(x) = x16 + 1. BIP16 processing must be performed on all OAM packets prior to being able to reliably pass their payload for further processing. Any OAM packets that show a BIP16 violation upon reception processing should be discarded.	The PDU format is based on the Function type. If the function type is G.8112 compatible, only then will BIP16 be checked. This procedure defeats the purpose of BIP16. Further no minimum length check. This will cause major difference between to protocols. Y.1711 module will drop all packets less than 44 octets. If the length check is OK, it checks BIP16. All new PDUs will fail this test.
6.1.1 Stack encoding OAM packets are differentiated from normal user-plane traffic by an increase of one in the label stack depth at a given LSP level at which they are inserted. Therefore, they maintain this label stack difference of one (from normal user-plane traffic) as they traverse any lower layer server LSPs. Label The OAM Alert Labelled header is added before (i.e., below) the normal user-plane forwarding labelled header at the LSP trail source point. EXP The OAM packets can be used on both E-LSPs and L-LSPs. The coding of the EXP field should be set to all 0s in the OAM Alert Labelled header and to whatever is the "minimum loss-probability PHB" in the preceding normal user-plane forwarding header for that LSP. This is to ensure the OAM packets have a PHB which ensures the lowest drop probability [5]. OAM capabilities defined in the future may require different encoding of the EXP field.	For OAM Alert label, EXP and TTL fields are modified.

S bit The S bit is set only in the OAM Alert Labelled header. TTL The TTL field should be set to 1 in the OAM Alert Labelled header. The reasons for this are: • OAM packets should never travel beyond the LSP trail termination sink point at the LSP level they were originally generated (noting that they are not examined by intermediate label-swapping LSRs, and are only observed at LSP sink points). • the TTL of the immediately prior normal user-plane forwarding header is used to mitigate against damage from looping packets.	
6.1.2 Intermediate/penultimate processing OAM packets are transparent to intermediate LSRs, including the penultimate LSRs.	In Y.1373 section 5.6 specifies "In order to distinguish OAM packets of nested MEGs, each MEP tunnels incoming OAM packets by incrementing the MEL in the source direction and decrementing it in the sink direction." Section 6.5 specifies the OAM label TTL field processing, which is different from standard TTL processing.
6.1.4 TTSI (Trail Termination Source Identifier) structure The structure of the LSP Trail Termination Source Identifier (TTSI) is defined by using a 16-octet LSR ID IPv6 address followed by a 4-octet LSP Tunnel ID.	In Recommendation G.8112 section 6.2.1.1 specify "The format of the LSP TTSI field in the FFD/CV OAM packets within the scope of T-MPLS is for further study". In Y.1373 defines Data plane identifier as TLV in section 10.1.2.3.
Defect location Defect location is a 4-byte field. The identity of the network in which the defect has been detected should be encoded in the Defect Location (DL) in the form of	This field is set to all zeros which is a reserve field.

an Autonomous System (AS) number. RFC 1930 [13] defines the AS number as being 2 bytes long. Section 6.8 of Y.1711 specifies Defect type entry/exit criteria and consequent actions. This function can not be provided in G.8112.	
Function types not supported Section B.4 says Y.1372 will never generate or receive FFD packets. It is the responsibility of management not to enable FFD generation in a G.8112 end-point. Y.1711 says "The provisioning of FFD on an LSP will provide capability of fast failure detection (by default FFD OAM is not provisioned). It is recommended that FFD will be generated at the LSP ingress at a rate of 20 per second. This will provide failure detection on the LSP in order of 100s of milliseconds". If this capability is not provisioned how the capability is provided? What happen if the FFD is enabled in G.8112?	
Frequency of sending CV packets In G.8112 the CV flow is generated at the LSP's source LSR with a nominal frequency of 1/s and terminated at the LSP's sink LSR. Section 7.1.1 defines the transmission period. 7 different values are specified for transmission period, the default values are recommended based on the application area for which CC is being used. When a transmission period other than the default value for an application area is used, the behaviour of the intended application is not guaranteed. To support backward compatibility only one value can be used. How other capabilities like protection switching supported? FFD has better granularity which is not supported.	
Ignoring the information by G.8112 Figure 10.2-6 specifies CV_v1 PDU format processed in a G.8112 (2006) compatible MEP. It further states "A G.8112 (2006) compatible MEP will ignore the information in the horizontal and vertical shaded fields, i.e. octet 6 [bits 1…5], octet 7 and octets 29…40". In G.8112 clearly states what fields should be ignored by the receiver. The fields in the figure 10.2.6 are not specified to be ignored. It is an implementation option how it be treated. It may cause backward compatibility problem.	