

Use of Ethernet Control Word **RECOMMENDED**

draft-ietf-pals-ethernet-cw-03

Stewart Bryant, Andy Malis & Ingas Bagdonas

IETF 101 March 2018

Status

- The authors took the feedback from IETF 100 and applied it to the text.
- We got some feedback.
- We want use this opportunity to get any other feedback the WG has on the new text, and assess whether this is now ready for WG Last Call.

Text Added to end of (4) Recommendation

Where the application of ECMP to an Ethernet PW traffic is required, and where both the ingress and the egress PEs support RFC6790 [RFC6790] (ELI) or both the ingress and the egress PEs support RFC6391 [RFC6391] (FAT PW), then either method may be used. The use of both methods on the same PW is not normally necessary and should be avoided unless circumstances require it. In the case of multi- segment PWs, if ELI/EL is used then it should be used on every segment of the PW. The method by which usage of ELI/EL on every segment is guaranteed is out of scope of this document.

Added to the End of (5) ECMP

The PW label is pushed before the LSP label. As the EL/ELI labels are part of the LSP layer rather than part of the PW layer, they are pushed after the PW label has been pushed.

MAC Header
LSP Label
Entropy Label Indicator
Entropy Label
PW Label
PW Control Word
Payload

Changed Text at End of 7 (Operations)

OLD

To remove this problem in the long term, and hence to reduce the operational cost of investigating problems associated with the incorrect forwarding of Ethernet packets over PWs not using the CW, it is RECOMMENDED that equipment that does not support the CW be phased out of operational use.

Changed Text at End of 7 (Operations)

Instead of including a payload type in the packet, MPLS relies on the control plane to signal the payload type that follows the bottom of the label stack. Some LSRs attempt to deduce the packet type by MPLS payload inspection, in some cases looking past the PW CW. If the payload appears to be IP or IP carried in an Ethernet header they perform an ECMP calculation based on what they assume to be the five tuple fields. However deduction of the payload type in this way is not an exact science, and where a packet that is not IP is mistaken for an IP packet the result can be packets delivered out of order. Misordering of this type can be difficult for an operator to diagnose. Operators therefore need to be careful when enabling capability that allows information gleaned from packet inspection past the PW CW to be used in any ECMP calculation.

Changed Text at End of 7 (Operations)

Instead of including a payload type in the packet, MPLS relies on the control plane to signal the payload type that follows the bottom of the label stack. **Some LSRs attempt to deduce the packet type by MPLS payload inspection, in some cases looking past the PW CW.** If the payload appears to be IP or IP carried in an Ethernet header they perform an ECMP calculation based on what they assume to be the five tuple fields. However deduction of the payload type in this way is not an exact science, and where a packet that is not IP is mistaken for an IP packet the result can be packets delivered out of order. Misordering of this type can be difficult for an operator to diagnose. Operators therefore need to be careful when enabling capability that allows information gleaned from packet inspection past the PW CW to be used in any ECMP calculation.

Changed Text at End of 7 (Operations)

Instead of including a payload type in the packet, MPLS relies on the control plane to signal the payload type that follows the bottom of the label stack. Some LSRs attempt to deduce the packet type by MPLS payload inspection, in some cases looking past the PW CW. **If the payload appears to be IP or IP carried in an Ethernet header they perform an ECMP calculation based on what they assume to be the five tuple fields.** However deduction of the payload type in this way is not an exact science, and where a packet that is not IP is mistaken for an IP packet the result can be packets delivered out of order. Misordering of this type can be difficult for an operator to diagnose. Operators therefore need to be careful when enabling capability that allows information gleaned from packet inspection past the PW CW to be used in any ECMP calculation.

Changed Text at End of 7 (Operations)

Instead of including a payload type in the packet, MPLS relies on the control plane to signal the payload type that follows the bottom of the label stack. Some LSRs attempt to deduce the packet type by MPLS payload inspection, in some cases looking past the PW CW. If the payload appears to be IP or IP carried in an Ethernet header they perform an ECMP calculation based on what they assume to be the five tuple fields. **However deduction of the payload type in this way is not an exact science, and where a packet that is not IP is mistaken for an IP packet the result can be packets delivered out of order.** Misordering of this type can be difficult for an operator to diagnose. Operators therefore need to be careful when enabling capability that allows information gleaned from packet inspection past the PW CW to be used in any ECMP calculation.

Changed Text at End of 7 (Operations)

Instead of including a payload type in the packet, MPLS relies on the control plane to signal the payload type that follows the bottom of the label stack. Some LSRs attempt to deduce the packet type by MPLS payload inspection, in some cases looking past the PW CW. If the payload appears to be IP or IP carried in an Ethernet header they perform an ECMP calculation based on what they assume to be the five tuple fields. However deduction of the payload type in this way is not an exact science, and where a packet that is not IP is mistaken for an IP packet the result can be packets delivered out of order. **Misordering of this type can be difficult for an operator to diagnose. Operators therefore need to be careful when enabling capability that allows information gleaned from packet inspection past the PW CW to be used in any ECMP calculation.**

The End