MPLS Multipath Extensions

Current Name:
Multipath Extensions for MPLS Traffic Engineering
draft-villamizar-mpls-tp-multipath-te-extn-02

Will be Renamed to:
Multipath Extensions for MPLS Traffic Engineering
draft-villamizar-mpls-multipath-te-extn-00

Curtis Villamizar (OCCNC)

The two MPLS Multipath drafts will be renamed shortly, removing TP from the draft name. Emphasis will be placed on carrying LSP with strict packet ordering requirements (of which MPLS-TP is an example). The Infinera IPR no longer applies since Entropy Label is used rather than Infinera's forwarding method and the IPR will not be carried over to the new drafts.

MPLS Multipath Extensions - Purpose - Overview

- "Use of Multipath with MPLS-TP and MPLS"
  aka draft-villamizar-mpls-tp-multipath-03
  - mentions limitations of **without** extensions
- This draft addresses those limitations by providing extensions to IGP-TE and RSVP-TE.

MPLS Multipath Extensions - Purpose - Specifics

- Identify LSP which require strict packet ordering (for example MPLS-TP LSP)
- Identify LSP which contain other LSP which require strict packet ordering (for example, MPLS-TP client LSP in MPLS server LSP) which MUST NOT traverse multipath hops which lack Entropy Label support.
- Provide an estimate of largest contained microflow in a given LSP with only a requirement not to reorder microflows. - see later slide
- Identify capabilities of multipath links for the purpose of path computation. - see later slide

Multipath Extensions - Microflow Estimate #1

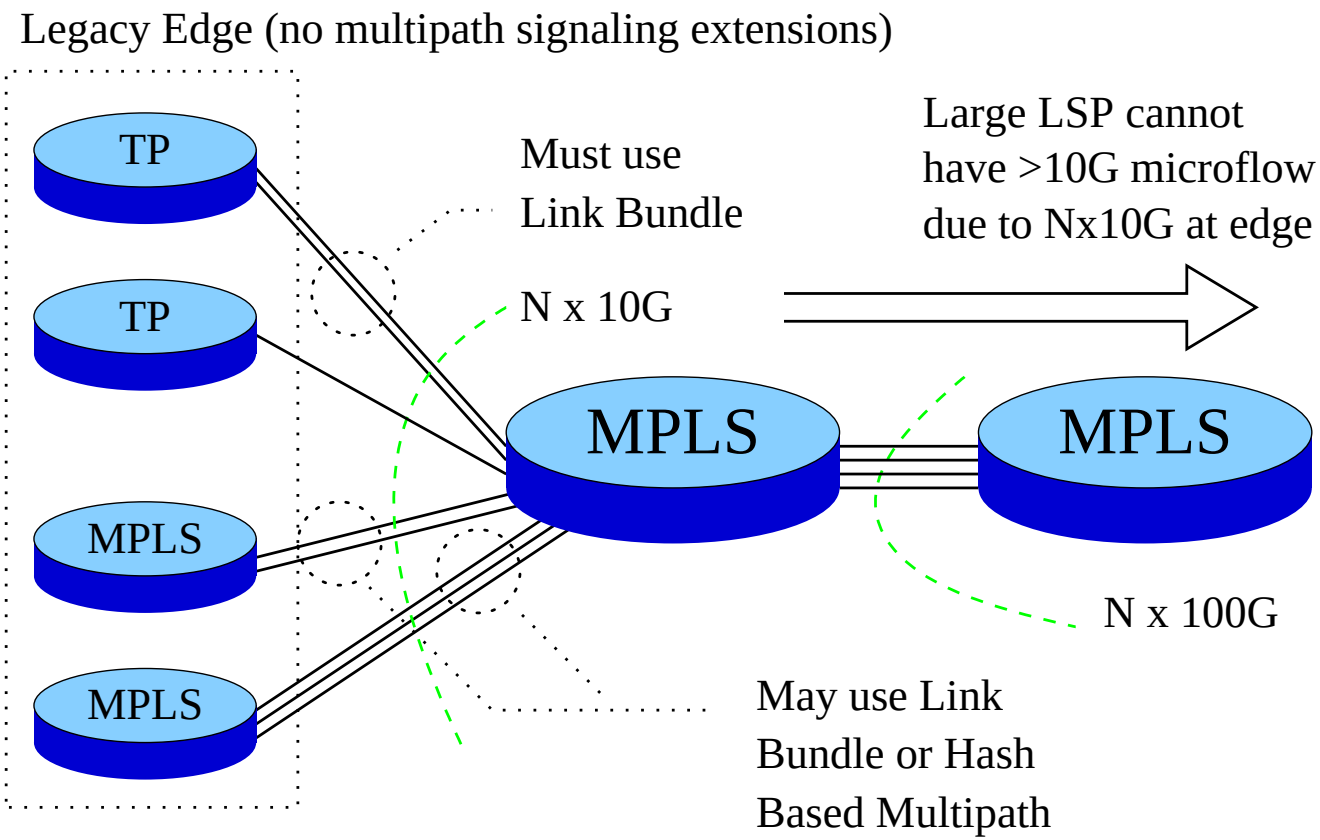

Multipath Extensions - Microflow Estimate #2

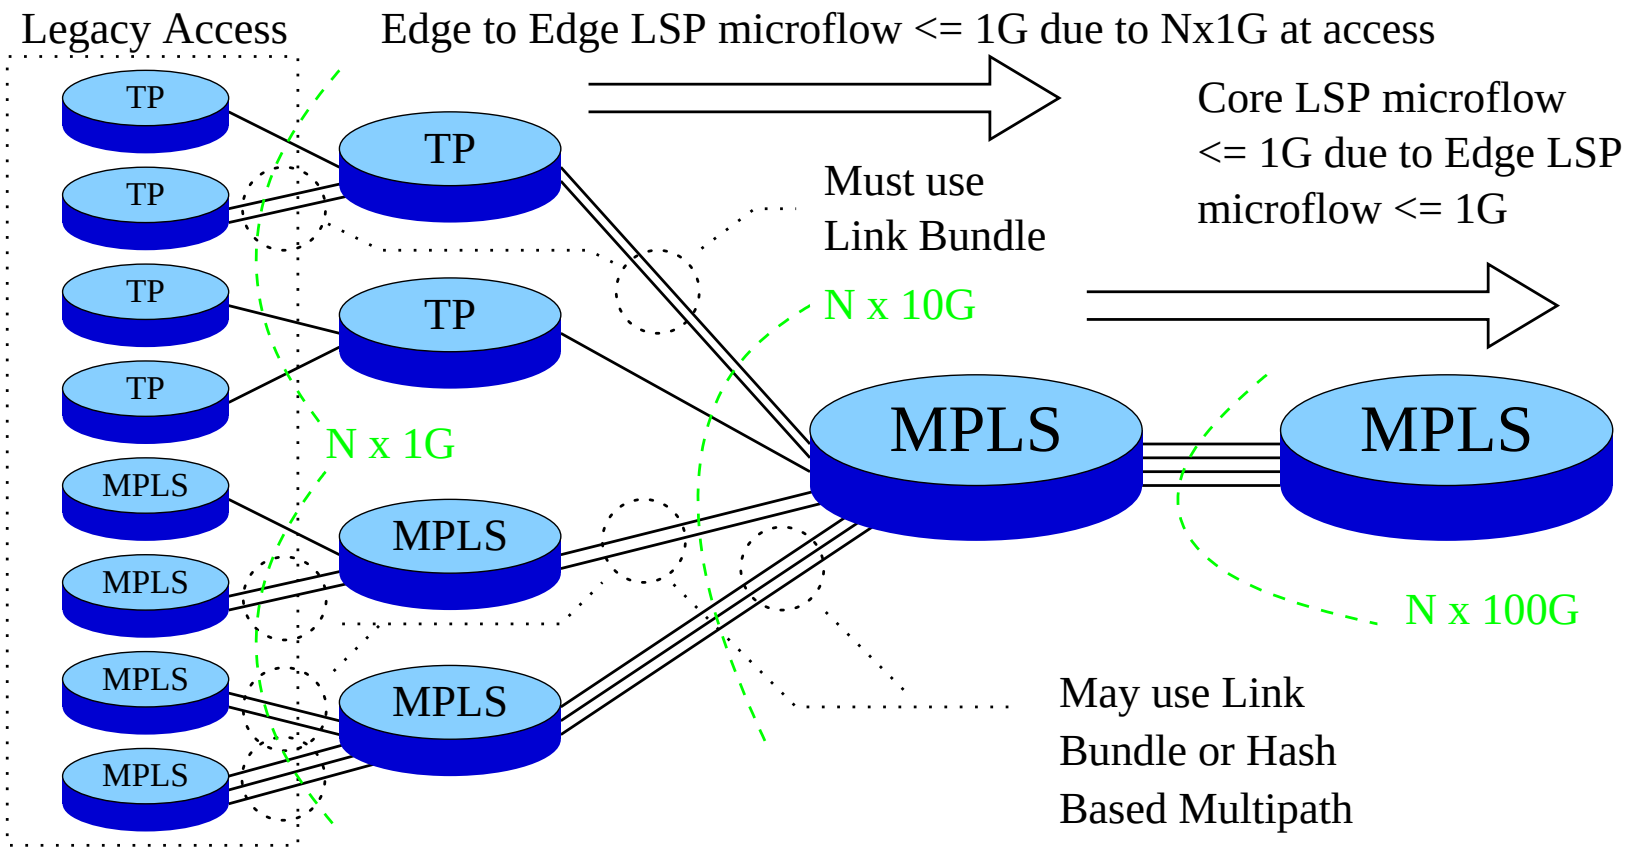

Multipath Extensions - Microflow Estimates

Please don't try to read too much into the above two **examples**.

The point of the two examples is:

1. The estimate of largest microflow within an LSP can be based on limitations imposed by the immediate interfaces in absense of signaling hints (first example).
2. The estimate of largest microflow within an LSP can be based on limitations of the contained LSP if the contained LSP provide signaling hints (second example).

MPLS Multipath Extensions - Link Capabilities

- Can the link support an LSP with strict packet ordering requirements?
  - a simple (non-multipath) link can
  - a traditional (not all-zeros componet) can
  - a traditional LAG cannot
- Does the link supports Entropy Label?
  - if so, a TP LSP can be carried if ELI/EL is used
- How large a microflow can the link support?
  - that value needs to be >= largest LSP microflow
- What fields can the LSR hash on?
  - labels only? IP addresses? Ports? IPv6?
- What hash fields can be disabled per LSP?
  - can looking for an IP header be disabled per LSP?
- What are the limits to the label stack depth usable by the hash or the label stack depth after which an IP header can be found?
  - if too shallow some LSP may have to avoid such hops.

MPLS Multipath Extensions - LSP Requirements

- Will use of hash on IP multipath cause reorder problems for the LSP?
  - this would be the case for PW with no CW
- Is IPv4 or IPv6 contained in the LSP?
  - If ELI/EL is not used, then hops must hash on IP
- Is support for Entropy Label required by the LSP?
  - this is the case for MPLS-TP carried as a client in MPLS LSP
  - a hop that does not support ELI/EL with reorder the TP LSP
- What is the largest microflow expected in the LSP?
  - hops with lots of low capacity component links may need to be avoided
- Does the LSP have strict packet ordering requirements but not use ELI/EL?
  - A MPLS PSC LSP can add ELI/EL to support such an LSP
- Does the LSP contain other LSP with strict packet ordering requirements and if so, how deep in the stack?
- How far into the label stack must the hash look for the largest microflow estimate to remain accurate?

Multipath Extensions - Backward Compatibility

- An LSR with legacy forwarding hardware can advertise its limitations so path computation can be more accurate.
  - Lack of Entropy Label support can be indicated.
  - Limited stack depth on hash can be indicated.
- An LSR with some improvements to forwarding hardware can indicate its level of support.
  - Ability to support either traditional or all-zeros component can be advertised.
  - Support for Entropy Label support can be indicated.
- Path computation can avoid problems (ie: MPLS-TP LSP being carried over hash based LAG) with simple software upgrade to the legacy LSR.
- If some legacy LSR do not support a software upgrade supporting multipath extensions, configuration on upgraded LSR can improve the guesses at capabilities of LSR which don't advertise their capabilities.
  - This is bound to be imperfect in some cases.

MPLS Multipath Extensions - TLVs

Two TLVs are proposed. One to advertise link capabilities (IGP-TE):

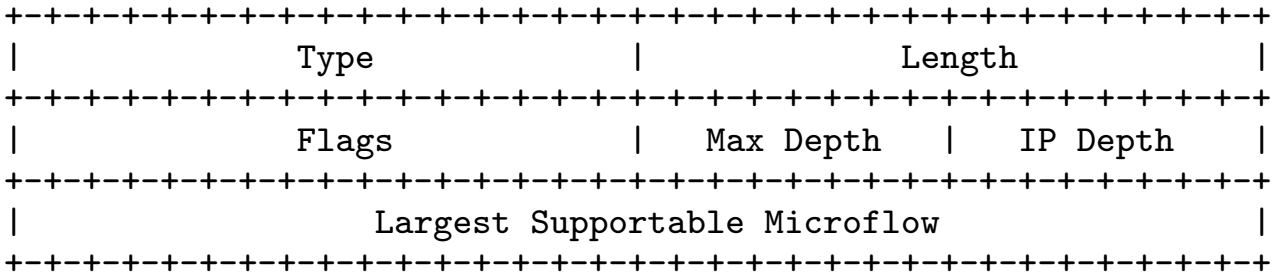

Figure 1: Multipath Capability Sub-TLV

One to indicate LSP characteristics and requirements (RSVP-TE):

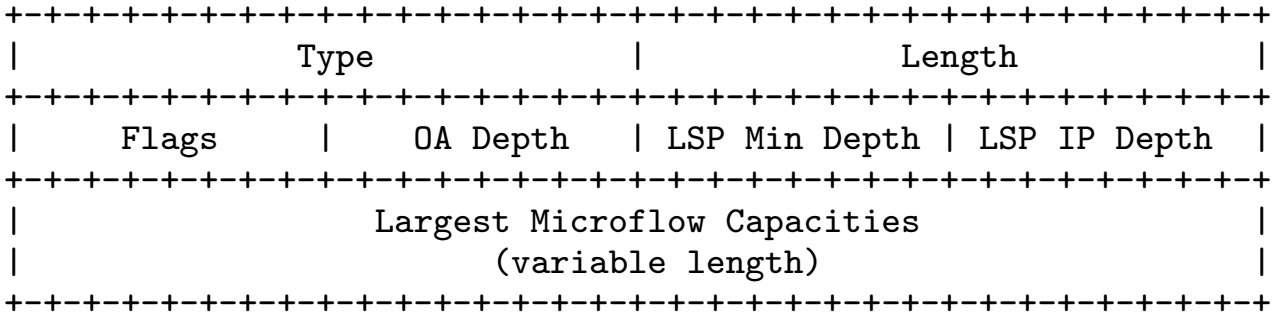

Figure 2: LSP Multipath Attributes TLV

Multipath Extensions - Applying TLVs

- Requirements and characteristics of an LSP may be based on its contained LSP.
- Based on requirements of an LSP, a path is computed.
- The same rules must be applied at each link at connection admission time.

MPLS Multipath Extensions - Conclusion

- Please **Read the draft and comment on it on the MPLS WG mailing list**
- draft-villamizar-mpls-tp-multipath-te-extn-02 is not as short a read as draft-villamizar-mpls-tp-multipath-03 but is still fairly simple.
- Some tweaking of TLV format and semantics is likely, but WG discussion should help.