

Single VID ECMP Frame Format Considerations

Ben Mack-Crane
(tmakcrane@huawei.com)

Goals

- Make the data path processing efficient
- Maintain existing field format and function
- Keep the frame overhead small

These goals may lead in different directions...
tradeoffs may be necessary

New Information

- Flow ID
 - Enables deep packet hash (e.g., 5-tuple hash) at edge only
 - Reduces frame processing at intermediate hops
 - Maintains independence from client at internal interfaces
 - Needs to be big enough to support sufficient entropy for ECMP
- TTL
 - Protects network from temporary loops during reconfiguration
 - Enables more rapid recovery in some cases (no need to wait for Agreement Protocol)
 - Protects network from control software variations (implementation bugs) that can cause forwarding loops
 - Need to be big enough to cover ECMP network domain diameter

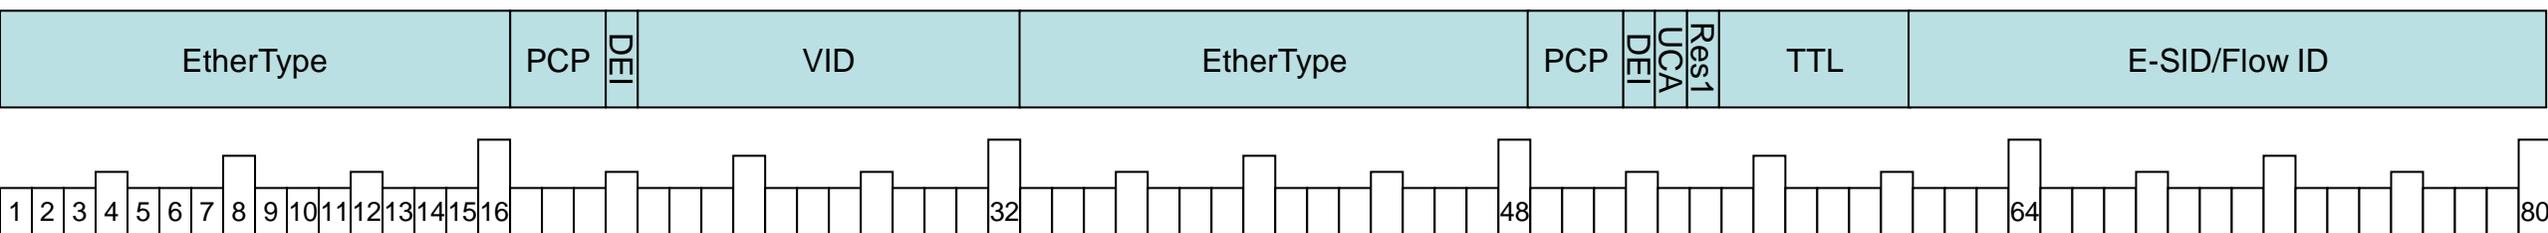
Old Information

- I-SID
 - Backbone service instance identifier
 - Needs to be big enough to support maximum number of services
 - Currently 24 bits
- BSI PCP and DEI
 - Copy of priority and discard eligibility for service
 - Enables service values to be different from backbone values
 - 4 bits
- UCA bit
 - Indicates distinct addressing carried in I-Tag
 - Used for OAM to allow standard MEPs/MIPs to work for BSIs

Required Information

- EtherType
 - Tags use fixed format
 - EtherType is format indicator

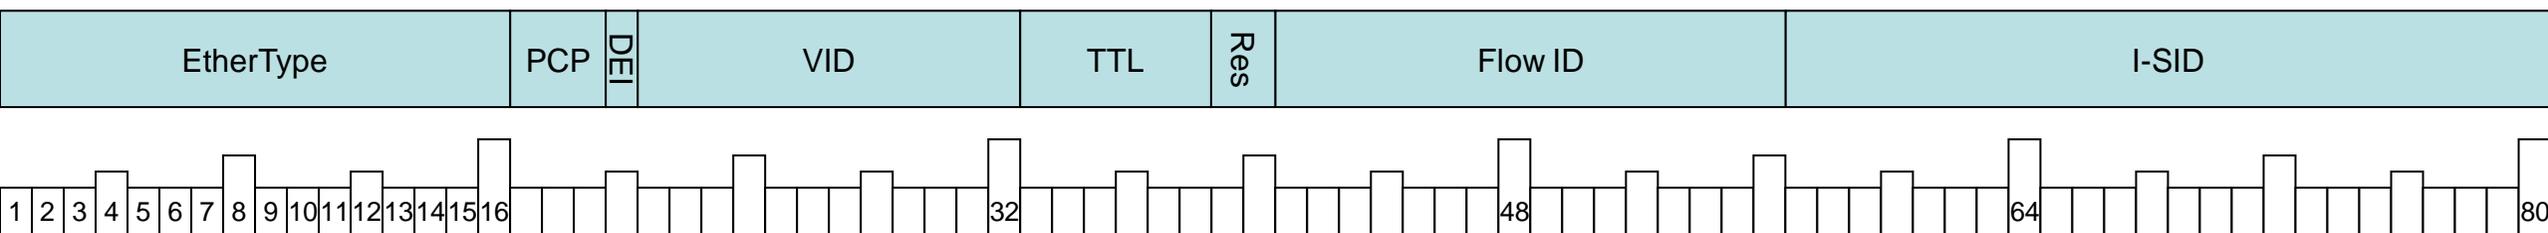
Example frame formats – Replace I-Tag with New Tag (E-SID)



Notes:

- 1) Standard B-Tag
- 2) New tag replacing I-Tag
- 3) I-SID replaced by 20 bit E-SID
- 4) Flow ID modulated over E-SID field
- 5) 6 bit TTL
- 6) Retains UCA to support OAM (are the semantics changed?)
- 7) 1 reserved bit
- 8) Uses two EtherTypes (one for B-VLAN and one for BSI/ECMP)
- 9) Concern that modulation requires hash calculation before forwarding lookup at egress (serializes otherwise parallel functions)

Example frame formats – New Tag with Both VID and I-SID



Notes:

- 1) Single tag does not require reading/checking a second EtherType (consistent with 802.1 practice)
- 2) Saves 2 octets in frame (does not require a second EtherType value)
- 3) First part of tag is identical to B-Tag/S-Tag
- 4) 24 bit I-SID
- 5) No separate service and backbone PCP/DEI (S-Tag can carry service values if needed)
- 6) Flow ID and I-SID independent (no field modulation)
- 7) 16 bit Flow ID
- 8) 6 bit TTL
- 9) 2 reserved bits (need to use one for OAM?)

Other Considerations...

- A number of new tags may be defined in 802.1
 - Recent past
 - I-Tag
 - CN-Tag
 - Potential future
 - Port Extension
 - Single VID ECMP
 - Other new functions (we may not stop working on bridging anytime soon)
 - Should we add a format indicator within (the next) 802.1 tag so new EtherTypes are not required each time?