

Single VID ECMP Frame Format Considerations

Ben Mack-Crane
(ben.mackcrane@huawei.com)

Goals

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

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

New Frame 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

Existing Frame 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

New Format Requires

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

IEEE RAC Rules...

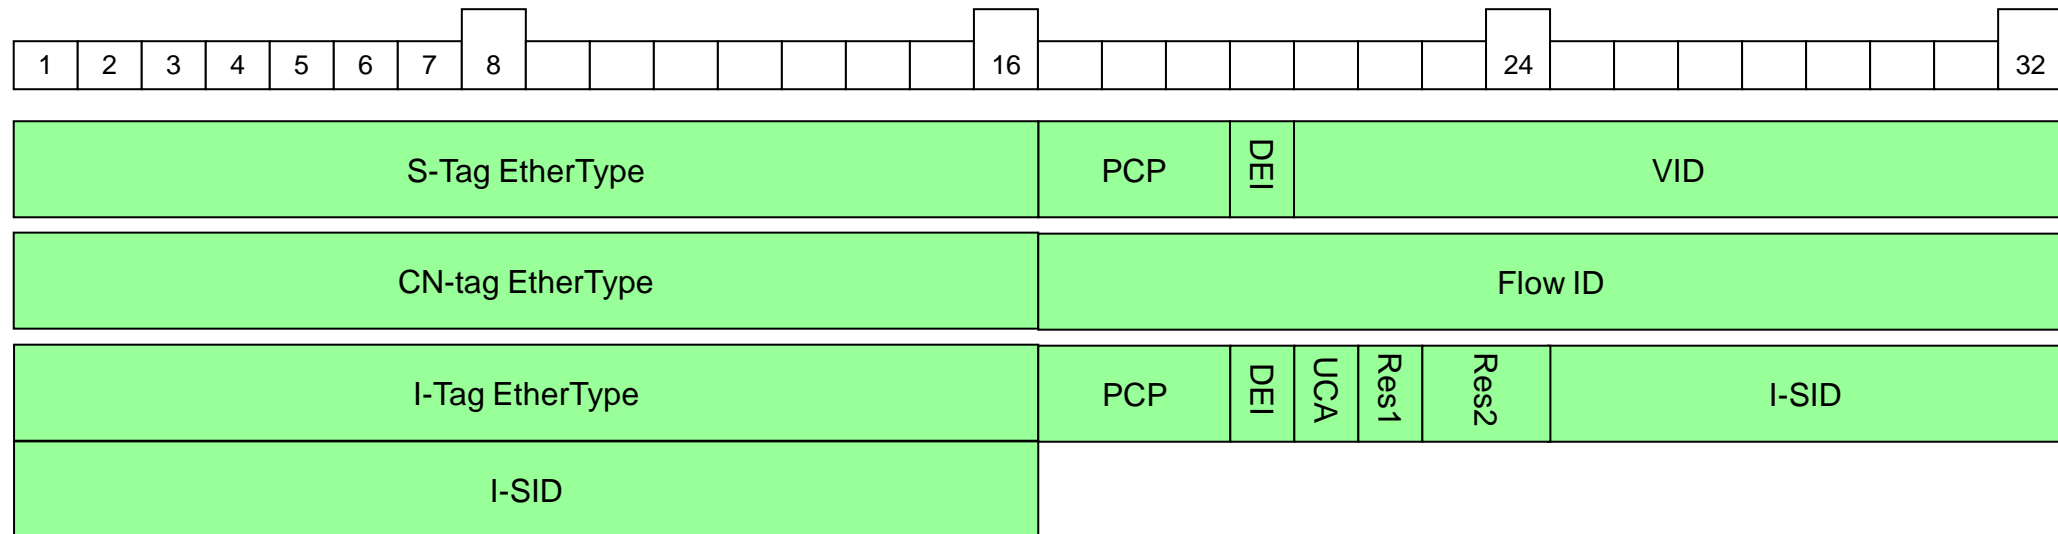
- A number of tags have been defined by 802.1
 - Originally
 - C-Tag
 - S-Tag
 - Recent past
 - I-Tag
 - CN-Tag
 - Potential future
 - Port Extension
 - Single VID ECMP
 - Other new functions (we may not stop working on bridging anytime soon)
 - IEEE RAC rules strongly recommend a subtype field be included so that new EtherTypes are not required for new versions or revisions

New Tag Format Considerations

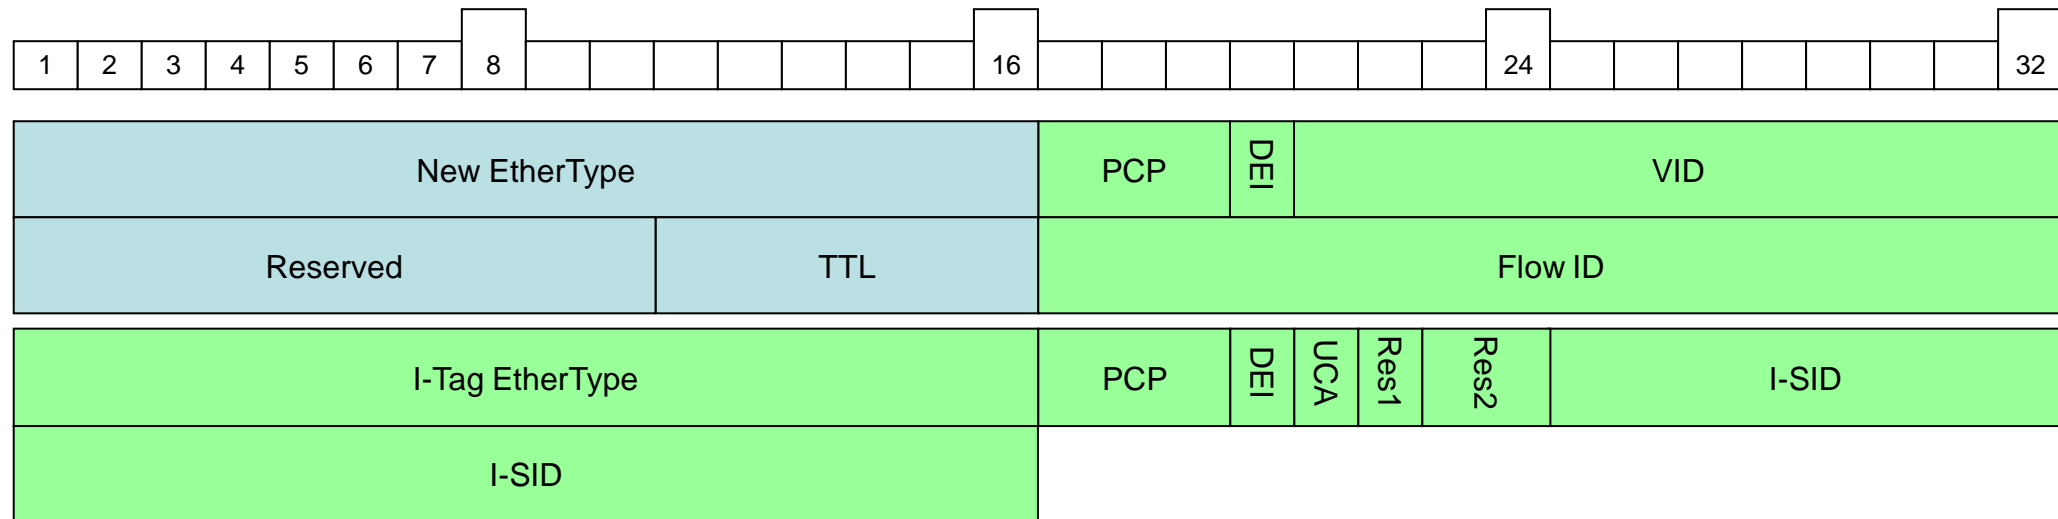
- 1) Subtype/format field for new EtherType
- 2) Explicit vs. implicit Flow ID
- 3) Flow ID field size
- 4) TTL field size
- 5) OAM frame identification
- 6) Reserved bits in tag (potential future evolution)
- 7) Maintain existing field formats
- 8) Overall size

Previous Tag Formats Considered

B-tag, CN-tag, I-tag



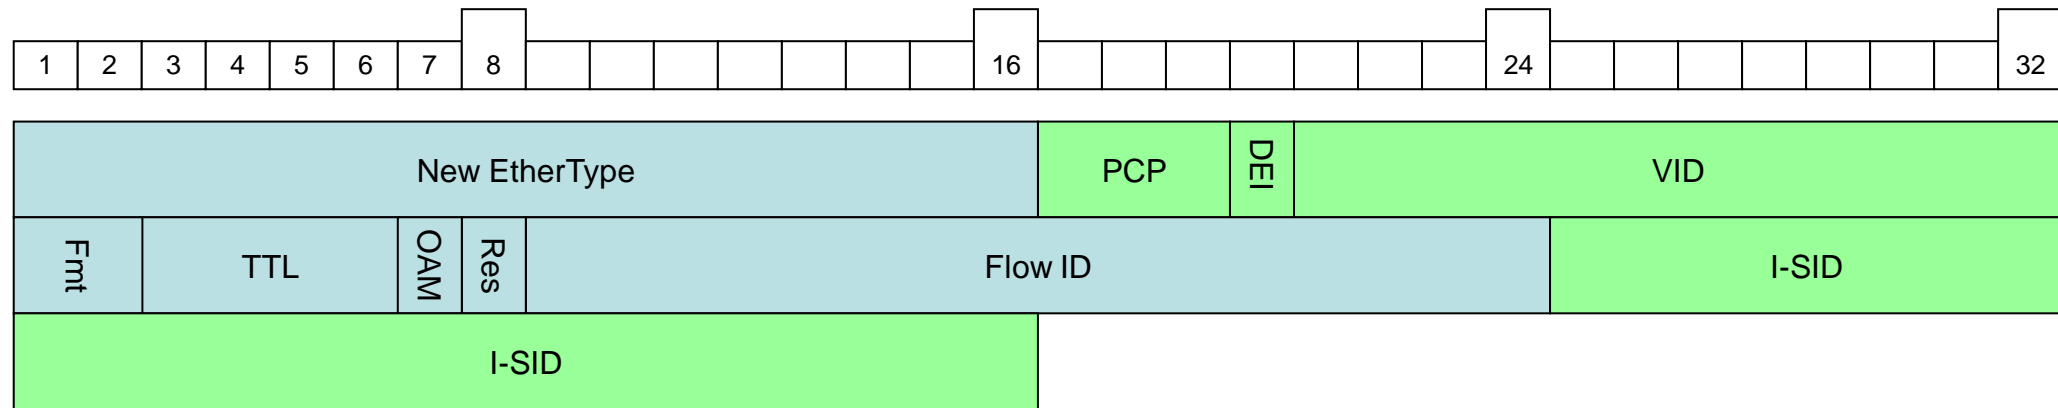
New Tag combining VID, FlowID



Notes:

- 1) No subtype or format field to support extension or revision (could be taken from reserved field?)
- 2) Explicit Flow ID (preserves independent semantics, simplifies processing, no new fault modes, etc.)
- 3) 16 bit Flow ID
- 4) 6 bit TTL
- 5) No OAM indicator (could be taken from reserved field?)
- 6) 10 reserved bits
- 7) I-tag remains unchanged, location of existing fields in B-tag and CN-tag unchanged
- 8) 14 bytes including new tag and I-tag

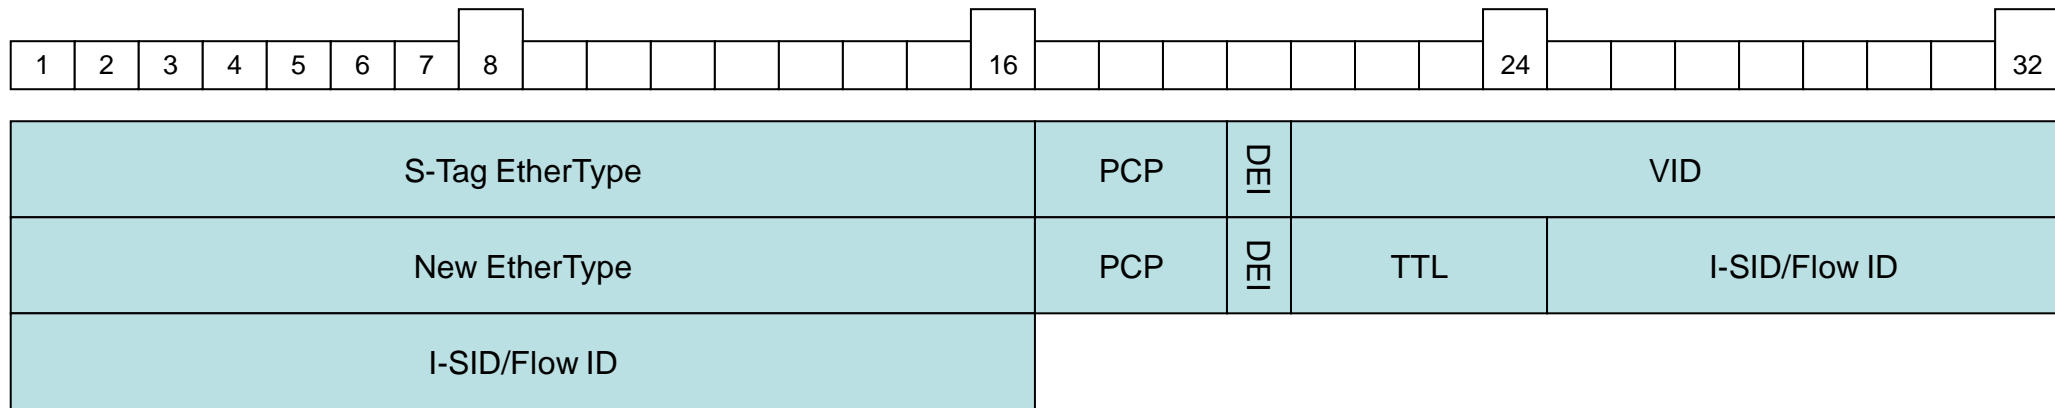
New Tag combining VID, FlowID, I-SID



Notes:

- 1) 2 bit subtype/format field (can reserve a value for extending subtype/format field)
- 2) Explicit Flow ID (preserves independent semantics, simplifies processing, no new fault modes, etc.)
- 3) 16 bit Flow ID
- 4) 4 bit TTL
- 5) OAM indicator bit
- 6) 1 reserved bit
- 7) I-SID remains 24 bits, first part of tag is identical to B-Tag/S-Tag
 - Easy to convert to B-Tag or B-Tag+I-Tag at domain boundary, if needed (just overwrite some fields)
- 8) 10 bytes including B-VID and I-SID

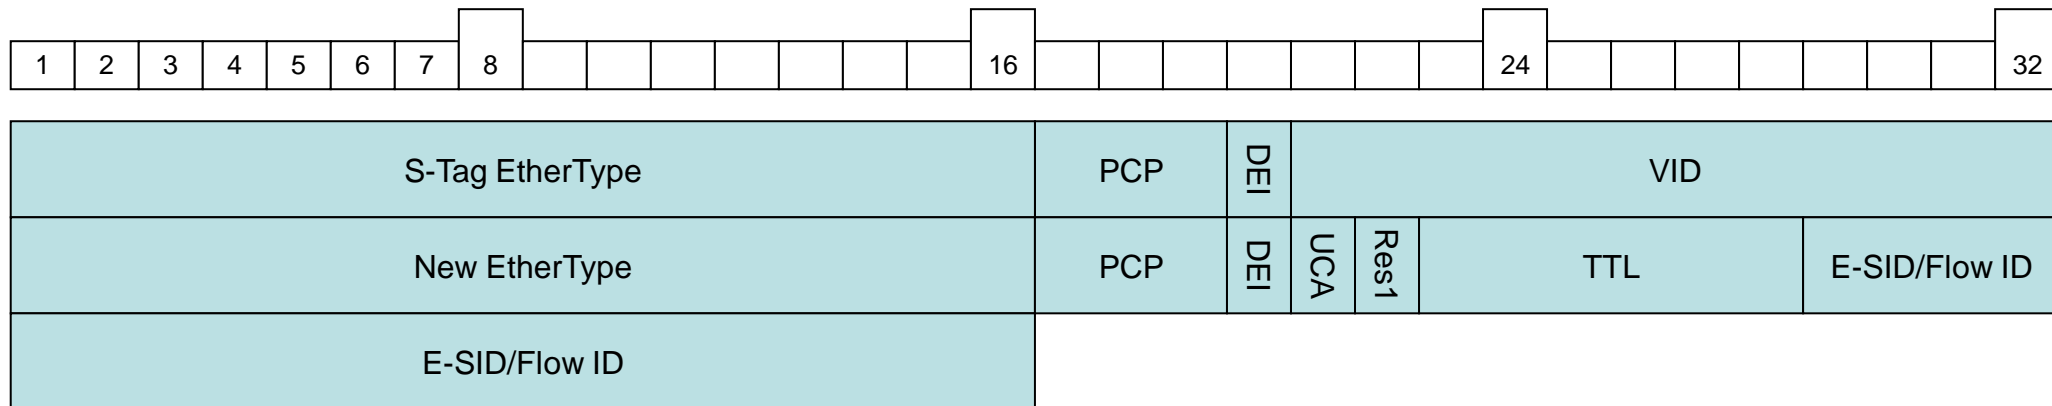
ECMP-only Tag with I-SID/Flow ID



Notes:

- 1) No subtype or format field to support extension or revision
- 2) Implicit Flow ID (Flow ID modulated over I-SID field)
 - Impacts egress processing efficiency
 - Introduces new failure mode
 - Hash mismatch can cause frame discard or mis-delivery
 - New CFM required?
 - Cannot mix hash algorithms in a single VLAN
- 3) 24 bit Flow ID (modulated with I-SID)
- 4) 4 bit TTL
- 5) Second new Ethertype for OAM
- 6) No reserved bits
- 7) I-SID remains 24 bits , first part of frame is B-Tag/S-Tag

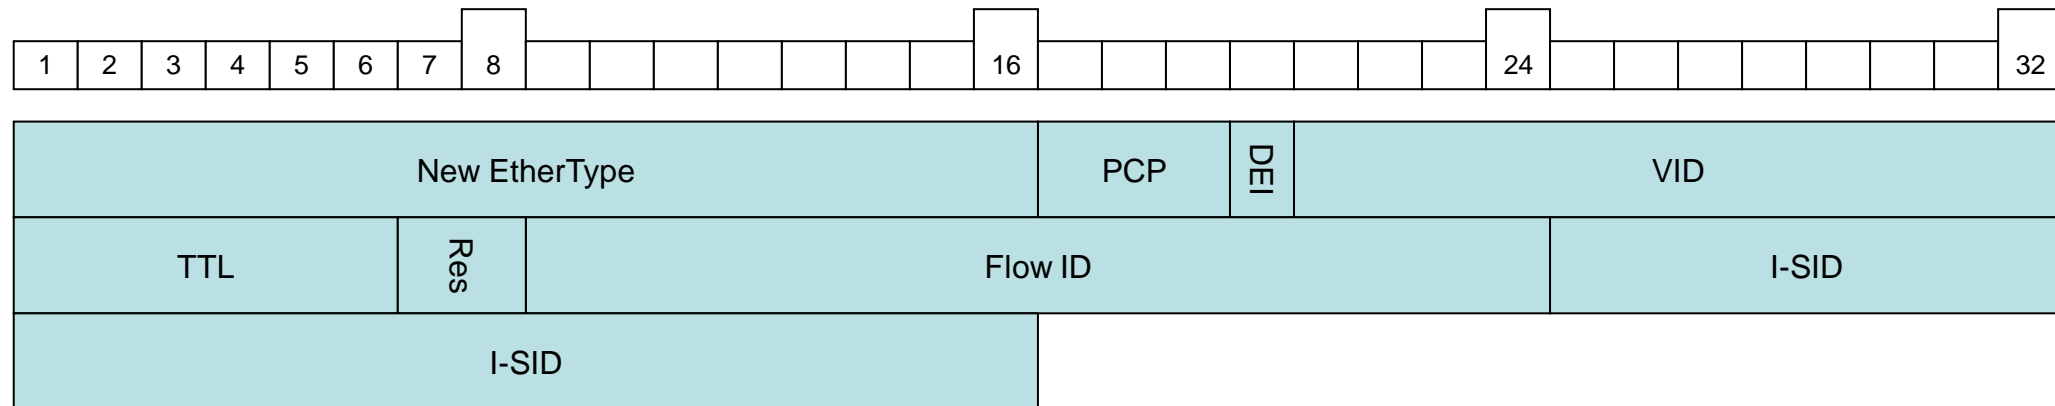
Replace I-Tag with New Tag (E-SID)



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)

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?)