ForCES Packet Parallelization

IETF – 85 Atlanta

Evangelos Haleplidis (ehalep@ece.upatras.gr)
Joel Halpern (jmh@joelhalpern.com)

Problem Drivers

- Actual hardware can sometimes perform multiple actions in parallel
- There is currently no way for a CE to tell an FE when and how it can make use of this capability
  - A really smart FE could guess, but…
- We therefore propose two new LFB classes to allow the CE to manipulate parallelism
  - Many devices will not support these. Some will.

Packet Parallelization Types

- Flood
  - Copies of a packet is sent to multiple LFBs
- Split
  - A packet is split into multiple equal size chunks[1] (CE-specified) and sent to multiple LFBs.

[1]Except the last one

Introduced LFBs

- Splitter
  - Splits the path of a packet and sends to multiple LFBs.
  - 1 singleton input port.
  - 1 group output port.
- Merger
  - Receives packets or chunks and merge them into 1.
  - 1 group input port.
  - 1 singleton output port.

Idea status

- This needs more work ad review
- This is currently not within the WG chart
  - But is a natural fit for this WG
- We suggest this be included as a work item in the new charter for the WG.

Use cases

Split Type:

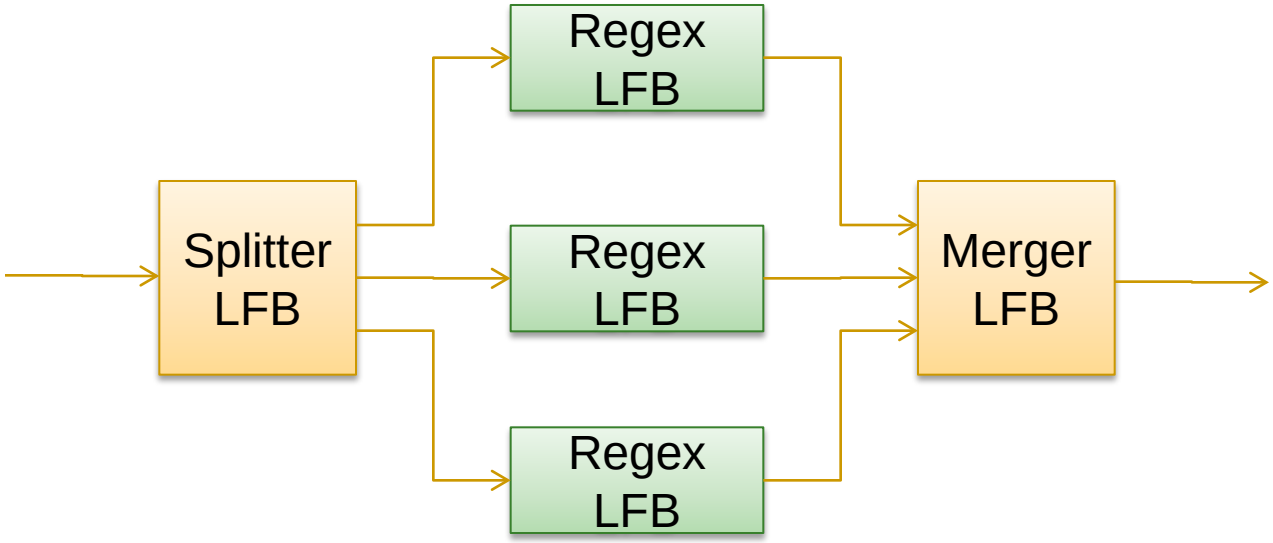

Flood Type:

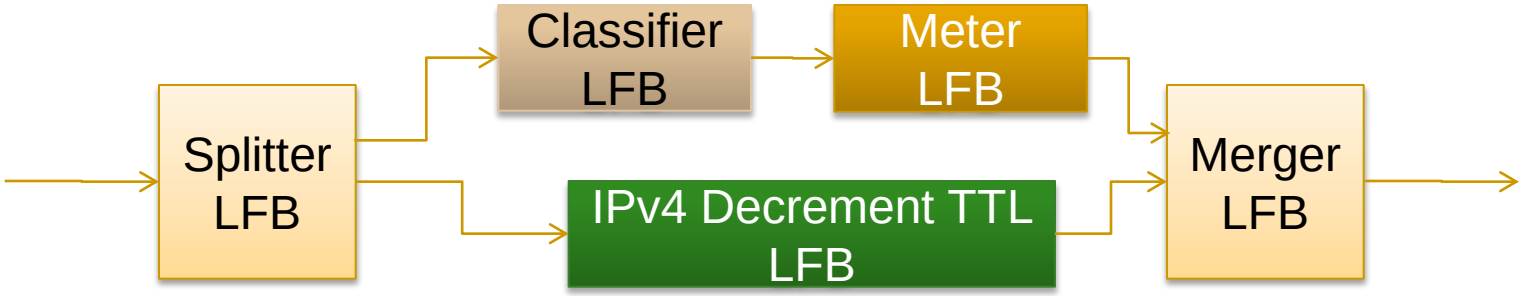

Use cases – complex

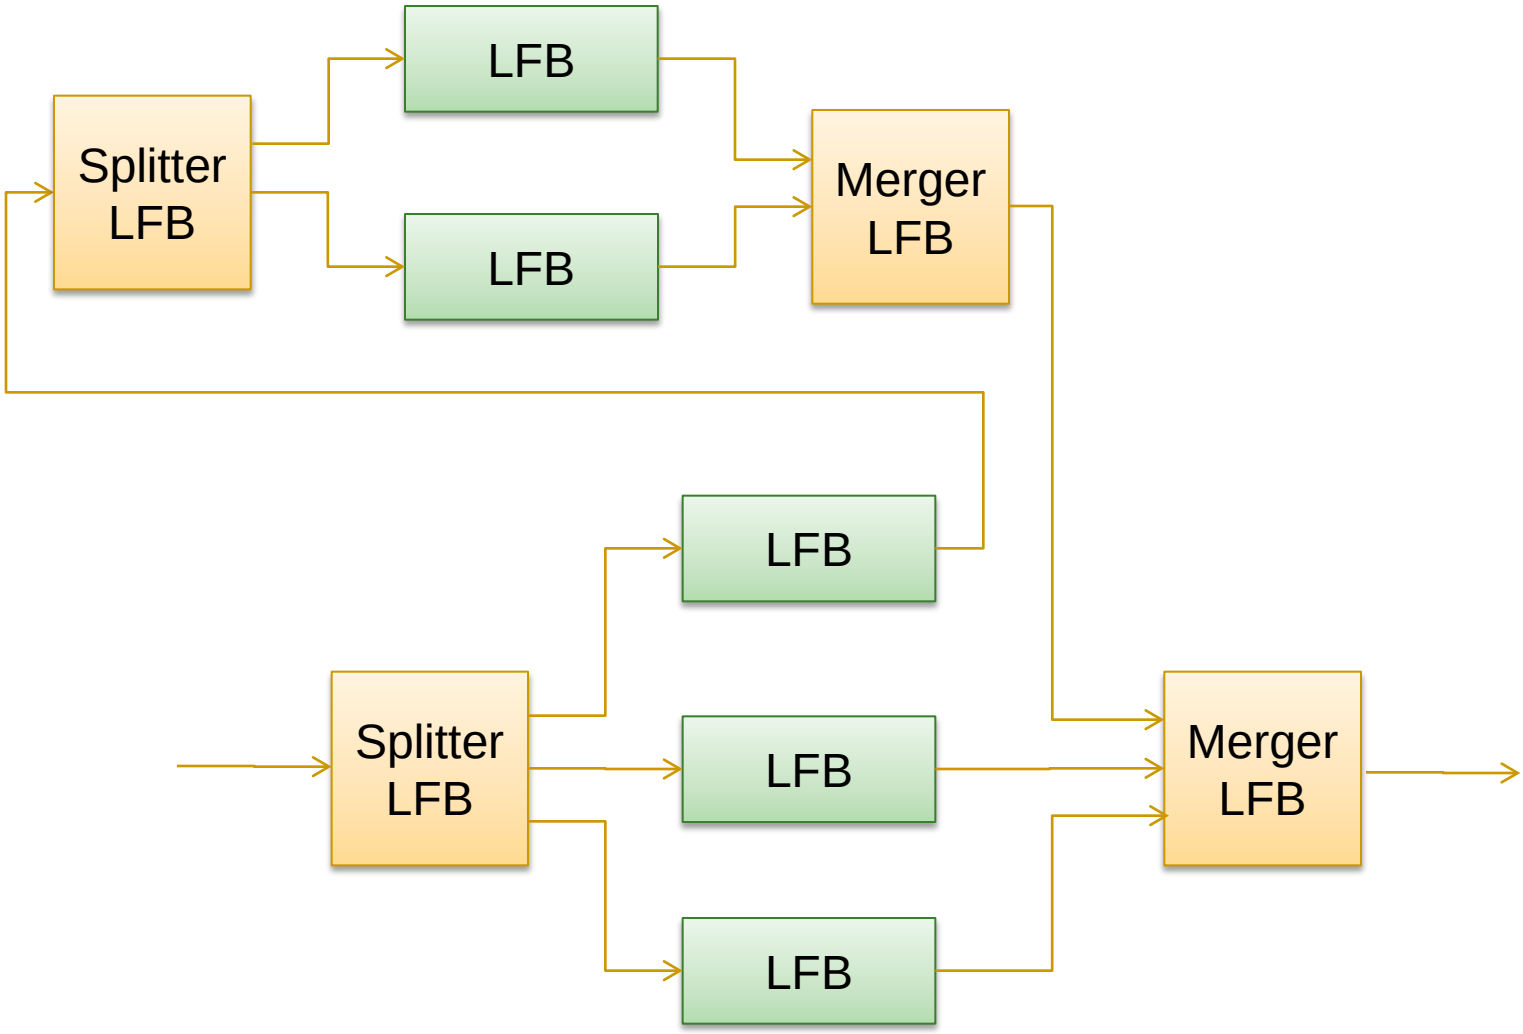

Split/Merge metadata

- Splitting/Merging are implementation issues – Document specifies operational parameters to control splitting/merging.
- Metadata created by splitter LFB to be received by Merger LFB – Opaque to LFBs in parallel paths:
  - ParallelType – Specifies flood or split.
  - Correlator – Identify packets or chunks belonging to the same original packet.
  - ParallelNum – Packet/Chunk number.
  - ParallelPartsCount – Total Number of Packets/Chunks.

Parallel Path mechanics

- In case of another splitter/merger in the path, the merge metadata MUST be tunneled through.
- A chunk/packet MAY be dropped in the path but the merge metadata MUST reach the Merger LFB. (opposite text remained by error in draft – will fix)
- Metadata produced in parallel paths MAY be aggregated with the merger LFB and sent on.
- In case of same metadata produced with different values, the first received MUST be kept.

Backup – 1: Splitter LFB

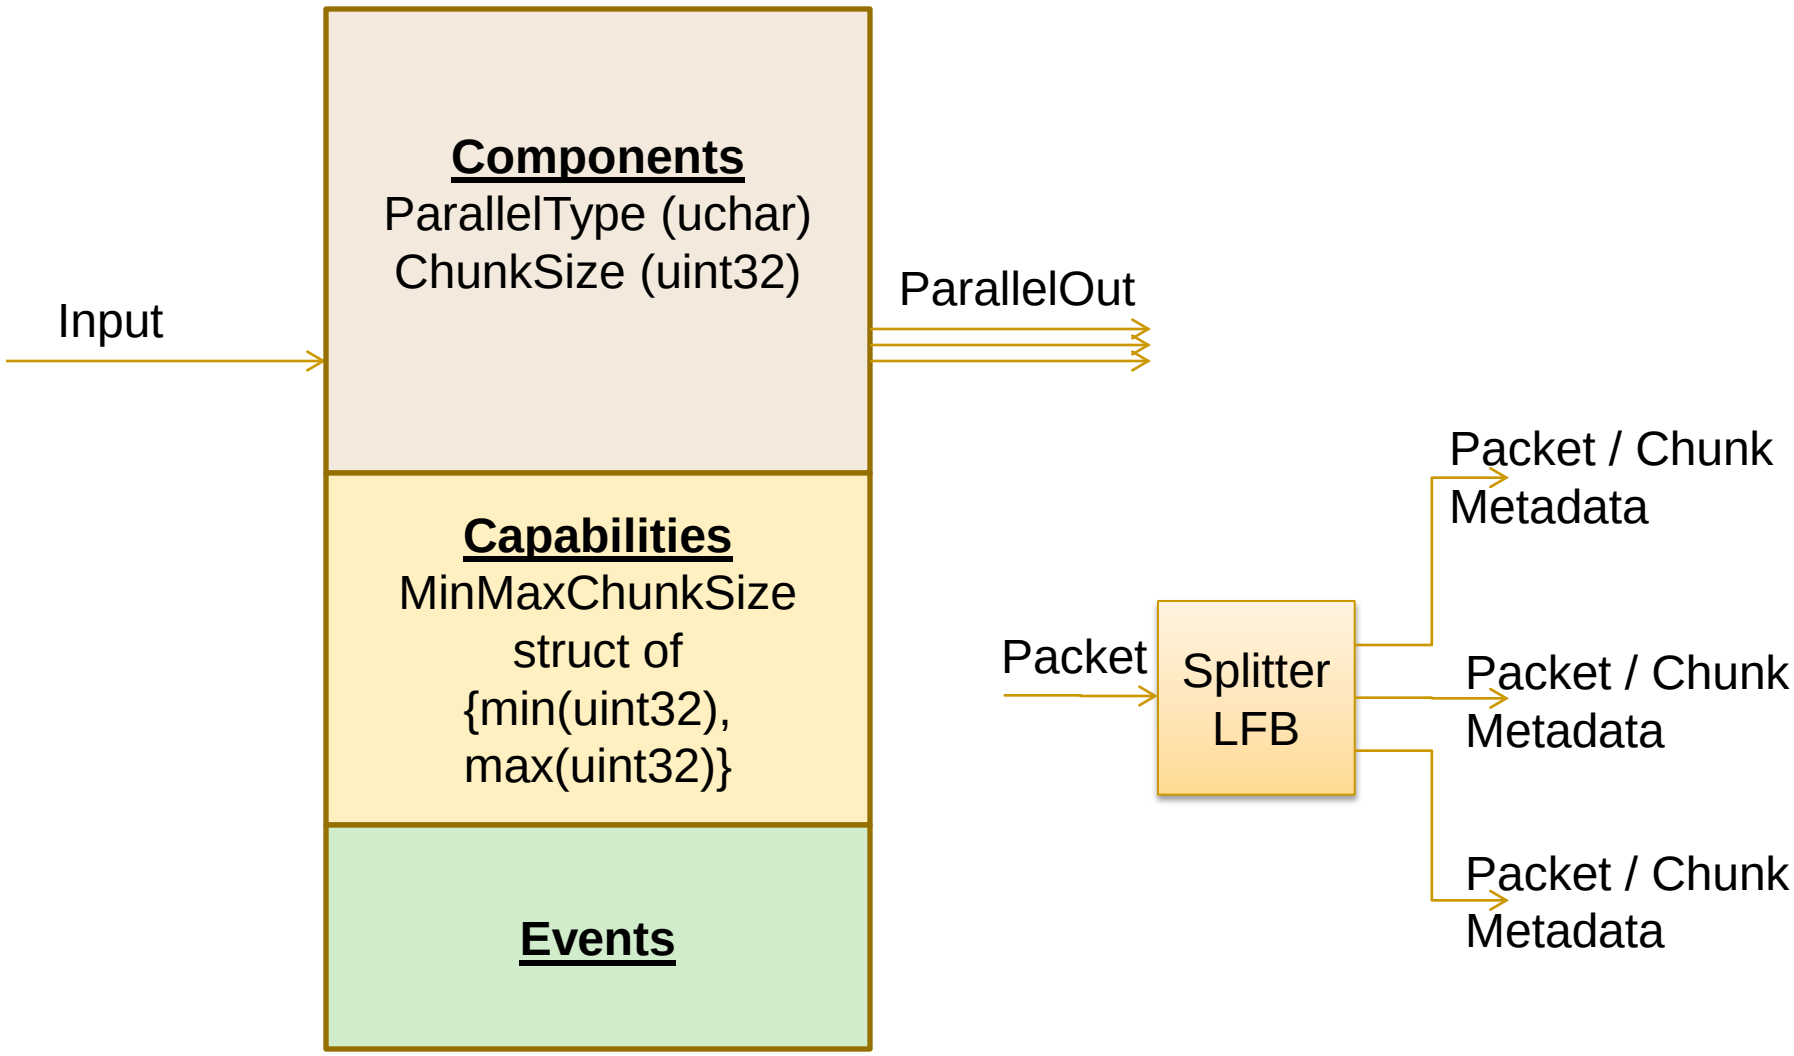

Backup – 2: Splitter LFB (2)

- Packet is received.
  - ParallelType == 0
    - Copies of packet is sent through parallel out through all output instances along with metadata.
  - ParallelType == 1
    - Packet is split into chunks of size==ChunkSize and each chunk is sent through one of instance of output instance's in a round-robin fashion.
    - Last chunk size's may be size<ChunkSize.

Backup – 3: Merger LFB

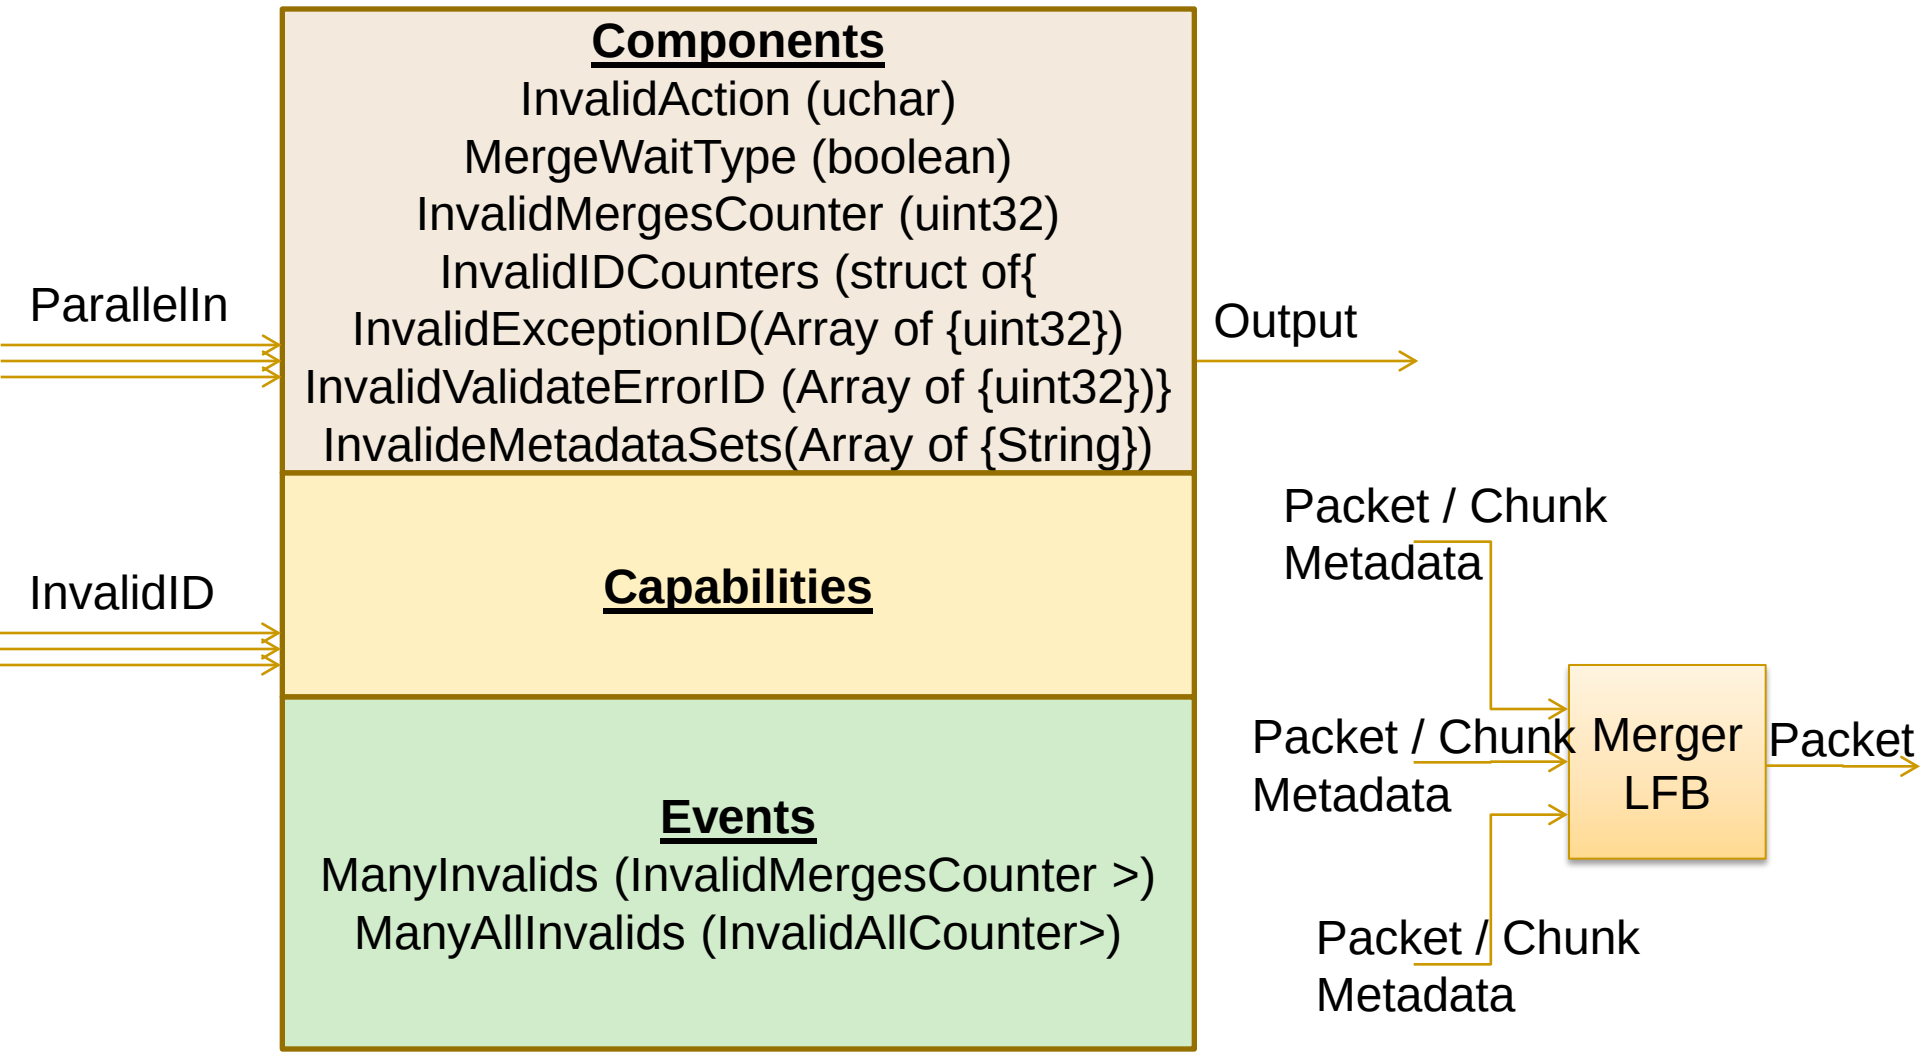

Backup – 4: Merger LFB (2)

- Receives packet/chunk via group input ParallelIn along with merging metadata.
- If packet/chunk was invalid it MUST receive the merging metadata and MAY receive an ExceptionID or ValidateErrorID and MAY receive the packet/chunk as well.
- If MergeWaitType==false the Merger LFB will strt merging upon receiving the first packet/chunk.
- If Invalid Action==0 it drops all packet/chunk. If 1 it will continue with the merge.

Backup – 5: Merger LFB (3)

- The merger LFB for statistics keep counters for the following:
  - InvalidMergesCounter – Merges with at least one Invalid
  - InvalidAllCounter – Merges with all invalid.
  - InvalideMetadataSets (optional) – Stores metadata sets along with the error id as a string.
- Includes two events:
  - InvalidMergesCounter greater than value.
  - InvalidAllCounter greater than value.