RoE generic header and the control plane v2

Jouni Korhonen

May 19, 2015

revised May 27

Bit and byte ordering

- Byte ordering
  - RoE shall use network byte order.
  - MSB first..
- Bit numbering for on-wire order
  - From 0 (MSB) to the highest (LSB)
- Example:

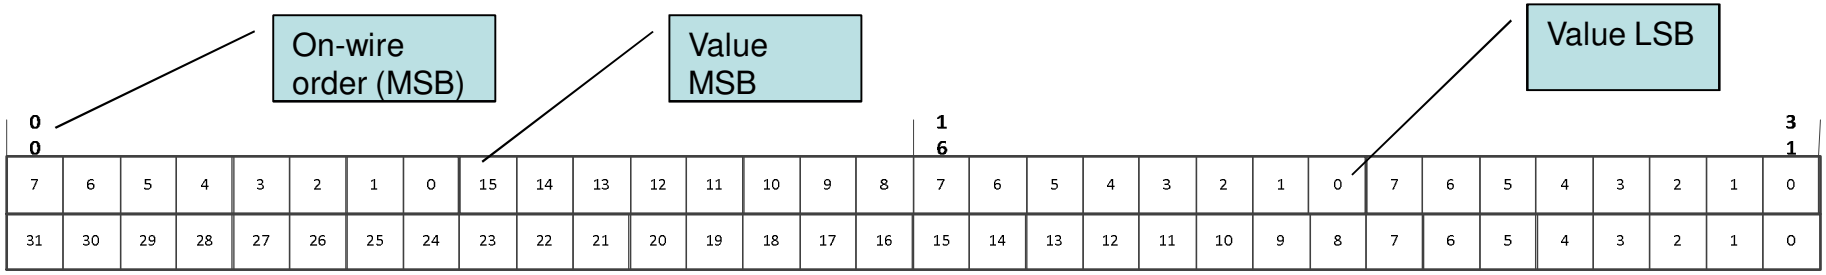

RoE for non-control packets

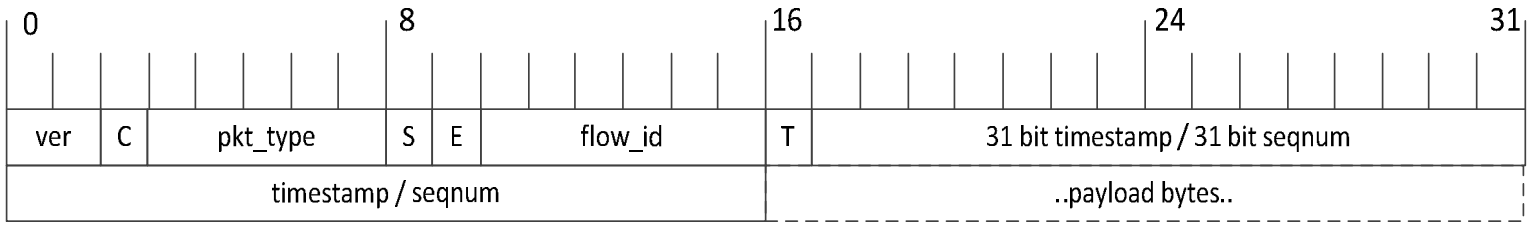

- ver – version field; current 00b
- pkt_type – RoE packet type;
  - 0 reserved for control packet
  - 1-31 other packet types
- flow_id – for multiplexing flows between SA/DA pair:
  - For example flow_id can be AxC number.
- Flags:
  - C = whether additional 32 bith "extended_header_space" follows the timestamp/seqnum; 0=no, 1=yes
  - S & E bits:
    - 10b start of frame (e.g. CPRI hyperframe; interpretation depends on the pkt_type).
    - 01b end of a frame.
    - 00b middle part of a frame.
    - 11b whole frame within the payload.
  - T = timestamp/seqnum selector; 0=seqnum, 1=timestamp

RoE header with extended_header_space

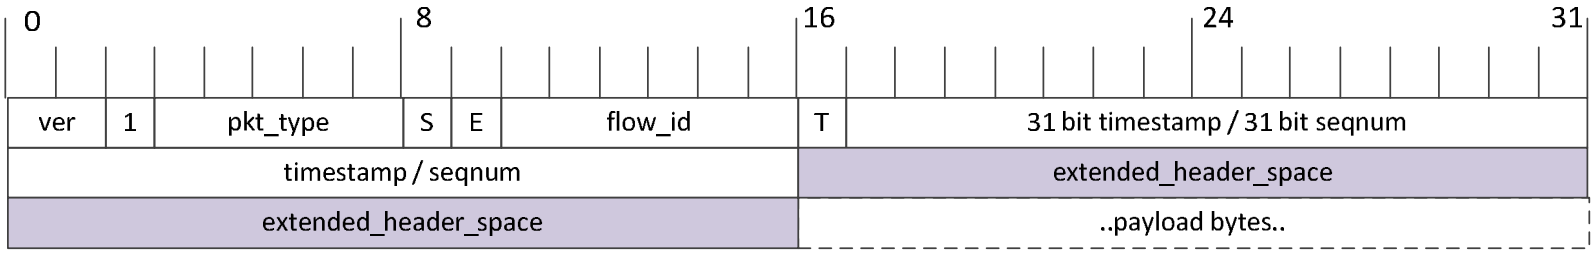

- C-flag = 1
- The content of the "extended_header_space" is opaque to RoE and used by the application.
- Possible uses include:
  - Control data that has to be delivered in timely manner along with the AxC and samples.
  - Data that must always follow the exact same processing path as the payload e.g. not forwarded to the management CPU.

About timestamp & seqnum

- Timestamp
  - 31 bits in size; units in nanoseconds
  - Represented as a _presentation_time_ at the received (and calculated by the sender based on its clock)
  - Can present time ~1s in future
  - Carries lower 31 bits of the calculated presentation time (~2s on wire but ~1s window)

- Sequence number (seqnum)
  - 31 bits in size; after 2^31-1 wraps to 0
  - Increments by a constant value known by both sender and receiver (configured or negotiated)

RoE control packet and payload TLVs

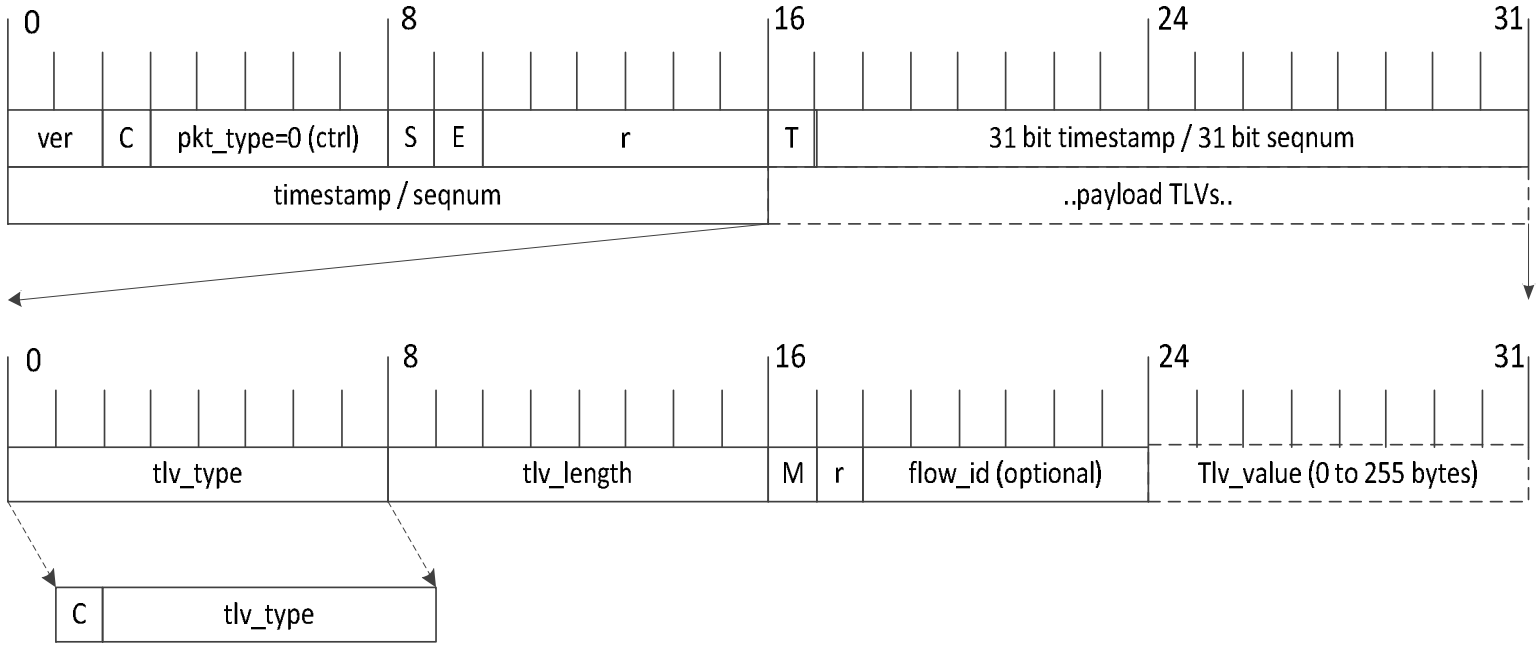

- In a RoE control packet the header level flow_id is insignificant
- S & E flags can be used to spread the control message over multiple RoE packets.
- TLVs
  - tlv_type (0-127 ignored if not understood by the receiver, 128-255 cause error if not understood by the receiver)
  - tlv_length excludes tlv_type, tlv_length and M/T/AxC_id
  - M = the TLV continuation flag;
  - flow_id – which flow this TLV concerns; optional and depends on the tlv_type

Potential TLVs

- Configuration:
  - Sample width, number of samples per RoE packet (i.e. how many Tc to interleave into one packet), ..
  - (obviously a state machine needed here for the link configuration time)
- For CPRI mapper use:
  - Link setup/negotiation TLVs -> C&M speeds, "CPRI" link speed, mapper method (and Na, K, S, M, etc)
  - Container for VSDs
  - Container for Slow/Fast C&M
  - Container for ctrl_AxC
- Generic container TLVs:
  - Carry some "alien" protocols like 1588.
- For path measurements/debug purposes
  - E.g. figure out intermediate node residence times.
- And so on.. A registry needed to be maintained somewhere.

Additional discussion

- Current RoE Control packet is designed to curry TLV and has implicitly only one packet type.

- Question: Do we need packet sub-type in the header format or is a “sub-type TLV” adequate?

Proposal

- Approve the RoE header format for data traffic.
- Approve the RoE header format for control traffic.
- Approve the RoE control packet's TLV format.