

Appendix A. Sync and Steering with Constant Overhead Word Stuffing (COWS)


This appendix presents a simple scheme for synchronization (PDU boundary retrieval). The basic mechanism described is inspired by [COBS]. However unlike the mechanism outlined in [COBS] there the PDU extended by 2 or 3 4 byte words regardless of the PDU length. With COWS:

- the iSCSI PDU is "prefixed" by a 8 byte COWS header (CH) including a 4 byte word aligned COWS framing pattern (CFP) and a 4 byte COWS link chain element (CLCE)
- any appearance of the pattern within the PDU is replaced by a forward or (for long payloads) a backward link to the next appearance of the pattern or to the end (for forward links) or an end of chain in indicator formatted as a CLCE.
- the iSCSI PDU may be followed by trailer consisting of a single CLCE

All those elements form a COWS Extended PDU (CEPDU)

COWS uses a framing pattern defined by the sender and a special version of COWS that does not require pattern replacement but requires the sender to guarantee that a CH will always appear at the beginning of a TCP segment may be used by specialized software stacks and/or hardware adapters and its use may be negotiated (PDU Alignment Option - PAO).

The CH format is:



CFP is a pattern selected by the sender and communicated to the receiver during the login phase. A CFP has to be specified by the sender for each direction.

Except for the PAO every occurrence of CFP within the payload is replaced by a CLCE.

The CLCE is composed of:

- N a bit is selected to be the complement of the corresponding bit in the Framing pattern
- Link - number of non-link 4 byte words to the next/previous link
- LT - link type coded as follows:
 - 00 - Forward Last Link - no more links follow - the Link field is the number of words to the end of PDU
 - 01 - Forward More Links - the link field is the number of words to the next link
 - 10 - Backward Last Link - the link field is 0
 - 11 - Backward More Links - the link field is the number of words from the preceding field

The LT field in the CH MUST be 00 or 01 and all CLCE within the iSCSI header (if any) MUST have an LT field of 00 or 01 (the iSCSI header is encoded ONLY with forward links).

The backward linking mechanism MAY be used by the sender to avoid storing very long data payloads before sending them and MUST be processed by the receiver.

If using backward linking the sender MUST include in the ECPDU a tailing CLCE and the tailing CLCE is the first CLCE in the "back-tracking chain" and MUST be linked to by the last CLCE in the "forward-tracking chain".

COWS ECPDU can follow one of the following outlines:

a) Only Forward Pointing

CH (mandatory)

.

.

Forward Pointing CLCE (optional)

.

.

.

Forward Pointing CLCE (optional)

.

.
Last Payload Word

b) Forward-and-Backward Pointing

CH (mandatory)

.
.
Forward Pointing CLCE (optional - may point to trailer if
last forward)

.
.
.
Backward Pointing CLCE (optional - may have link of 0)

.
.
Last Payload Word
Backward Pointing CLCE (mandatory - may have a link of 0)

c) PDU alignment option

CH (mandatory)

.
.
.
Last Payload Word (also end of TCP segment)

For cases a) and b) the payload on the wire is guaranteed not to contain a CFP a word aligned pattern anywhere but in CH. For case c) the CFP is supposed to appear only aligned to TCP segment boundaries and be implemented with specialized software stacks and hardware. For this case the Link value, LT and the Reserved bits may be used as a further validity check (TBD???) .

In all cases the iSCSI PDUs are not constrained to a one or a limited number of TCP segments.

A.1 Negotiation

- A.2 Sent PDU processing
 - pseudo-language description...
- A.3 Received PDU processing
 - pseudo-language description
- A.4 Search for framing processing