

## Appendix A. Sync and Steering with Fixed Interval Markers

This appendix presents a simple scheme for synchronization (PDU boundary retrieval). It uses markers that include synchronization information placed at fixed intervals in the TCP stream.

A Marker consists of:

```

Byte /      0      |      1      |      2      |      3      |
   /      /      |      /      |      /      |      /      |
  |7 6 5 4 3 2 1 0|7 6 5 4 3 2 1 0|7 6 5 4 3 2 1 0|7 6 5 4 3 2 1 0|
  +-----+-----+-----+-----+
0| Next-iSCSI-PDU-start pointer - copy #1
  +-----+-----+-----+-----+
4| Next-iSCSI-PDU-start pointer - copy #2
  +-----+-----+-----+-----+

```

The Marker schemes uses payload byte stream counting that includes every byte placed by iSCSI in the TCP stream except for the markers themselves. It excludes also any bytes that TCP counts but are not originated by iSCSI.

The Marker indicates the offset to the next iSCSI PDU header. The Marker is eight bytes in length and contains two 32-bit offset fields that indicate how many bytes to skip in the TCP stream in order to find the next iSCSI PDU header. The marker uses two copies of the pointer so that a marker spanning a TCP packet boundary should leave at least one valid copy in one of the packets.

The inserted value is independent of the marker interval.

The use of markers is negotiable. The initiator and target MAY indicate their readiness to receive and/or send markers during login separately for each connection. The default is NO. In certain environments, a sender not willing to supply markers to a receiver willing to accept markers MAY suffer from a considerable performance degradation.

### A.1 Markers At Fixed Intervals

At fixed intervals in the TCP byte stream, a marker is inserted. During login, each end of the iSCSI session specifies the interval at which it is willing to receive the marker, or disables the marker

altogether. If a receiver indicates that it desires a marker, the sender SHOULD agree (during negotiation) and provide the marker at the desired interval.

The marker interval and the initial marker-less interval are counted in terms of the bytes placed in the TCP stream data by iSCSI.

When reduced to iSCSI terms, markers MUST indicate the offset to a 4-byte word boundary in the stream. The last two bits of each marker word are reserved and are considered 0 for offset computation.

Padding iSCSI PDU payloads to 4-byte word boundaries simplifies marker manipulation.

## A.2 Initial Marker-less Interval

To enable the connection setup including the login phase negotiation, marking (if any) is started only at the first marker interval after the end of the login phase. However, in order to enable the marker inclusion and exclusion mechanism to work without knowledge of the length of login phase the first marker will be placed in the TCP stream as if the Marker-less interval had included markers.

As an example if the marker interval is 512 and the login ended at byte 1003 (first iSCSI placed byte is 0) the first marker will be inserted after byte 1031 in the stream.

## A.3 Negotiation