

# Packet Formats of Non-Renegable Selective Acknowledgments (NR-SACKs) for MPTCP

## I. MODIFIED MULTIPATH CAPABLE (MP\_CAPABLE) OPTION

Before sending/receiving NR-SACKs, two end hosts must negotiate NR-SACK usage during the connection initiation phase. A proposed modified MP\_CAPABLE option is shown in Figure 1. Two bits — 'N' and 'n' — are used. During the three-way handshake, N bits of the two SYNs (SYN and SYN/ACK) indicates "NR-SACK capability of the SYN's sender". The decision of using NR-SACKs in data transfer is confirmed by the setting of N bit in the third packet (the ACK). N bit in the ACK packet =  $N_{SYN} AND N_{SYN/ACK}$ , which means NR-SACK is used only if both endpoints are NR-SACK capable.

In a packet, the n bit has meaning only if N = 1, otherwise the n bit MUST be ignored. n = 1 indicates the size of one NR-SACK block is 6 bytes, and n = 0 means the size of one NR-SACK block is 8 bytes. The reason for using variant NR-SACK block size is explained in section 3.2. The decision of the size of one NR-SACK block in data transfer is confirmed by the setting of n bit in the ACK packet. n bit in the ACK =  $n_{SYN} AND n_{SYN/ACK}$ , which means the size of one NR-SACK block is 6 bytes only if both endpoints set n = 1 in their SYNs, else the size is 8 bytes.

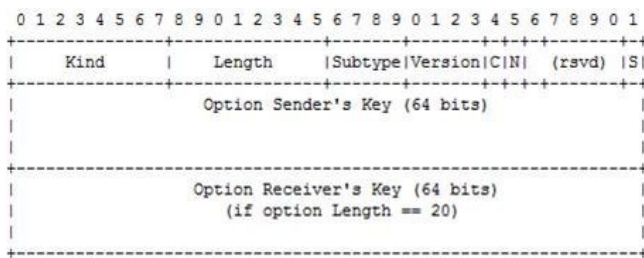


Figure 1. Modified MP\_CAPABLE Option

## II. MODIFIED DATA SEQUENCE SIGNAL (DSS) OPTION INCLUDING NR-SACK

Before talking about the proposed DSS option, consider how many NR-SACK blocks can be present in the TCP

option field. During unidirectional MPTCP data transfers, the NR-SACKs are carried by pure acks (acks without application data). The maximum size of the TCP option field is 40 bytes. A timestamp option occupies 12 bytes (with padding) leaving 28 bytes. Assuming no SACK information, a DATA ACK needs 8 or 12 bytes (depending on flag 'a'), thus only up to 20 bytes can be used for NR-SACKs. To decrease the number of bytes needed to represent one NR-SACK block, the left and right edge values of a reported NR-SACK block can be defined relative to the DATA ACK value. For example, if the MPTCP receiver receives out-of-order data with DSNs from  $DSN_{start}$  to  $DSN_{end}$ , the left and right edge values of the reported NR-SACK block are  $DSN_{start} - DATAACK$  and  $DSN_{end} + 1 - DATAACK$ , respectively. With 6-byte NR-SACK block, up to 3 blocks can be present and out-of-order bytes within  $2^{24}$  (16MB) from the DATA ACK can be reported. When an MPTCP receive buffer size is  $\leq 16MB$ , 6 bytes is sufficient. However, when an MPTCP receive buffer size is  $\geq 16MB$ , 6 bytes may not be enough. In this situation, the size of one NR-SACK block can be negotiated to be 8 bytes during connection establishment. Only 2 NR-SACK blocks will fit if the size of one NR-SACK block is 8 bytes.

The proposed modified DSS options with NR-SACKs are shown in Figure 2 (each NR-SACK is 6 bytes) and 3 (each NR-SACK is 8 bytes). A 2-bit unsigned integer — 'C' — is used to indicate the number of presented NR-SACK blocks. When the size of one NR-SACK block is 6 bytes and 1 or 3 NR-SACK blocks are present, two bytes paddings are used for alignment. The NR-SACKs can be present only when DATA ACK is present, and NR-SACKs yield the TCP option space to all TCP and other MPTCP options. As specified for SACKs in TCP, NR-SACKs always report the block containing the most recently received data, because this approach provides a MPTCP sender with the most up-to-date information about the state of a MPTCP receive buffer.

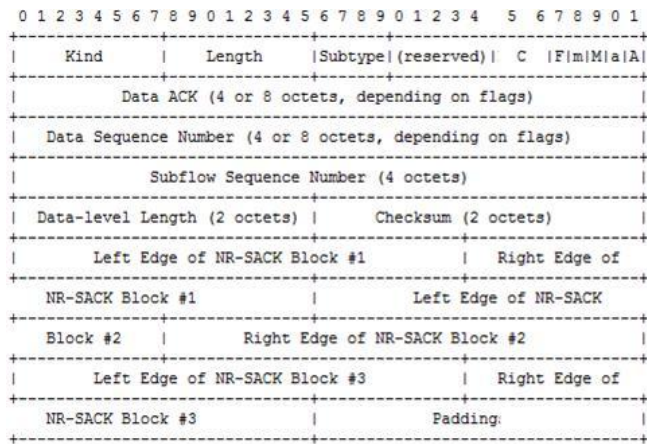


Figure 2. Modified DSS Option (each NR-SACK is 6 bytes)

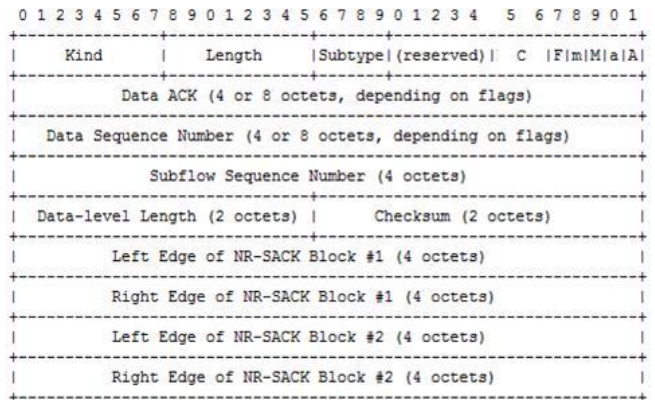


Figure 3. Modified DSS Option (each NR-SACK is 8 bytes)