

**Ladder Diagrams for Error
Recovery For FCP -2 Rev 04
Out-Of-Order Delivery- Annex D**

Carl Zeitler

Compaq Computer Corporation

April 24, 2000


T10/00-137r3

Reference: T11/00-133r1 and r2


Summary of Major Changes

- Put SRR back in for Class 2 recovery
 - D.3, D.5, D.5a, Class 2
- D.5 Class 3 Revert to doing REC and SRR on a closed Exchange
- Generated two possible ways of recovering from lost ACK to FCP_CONF, Class 2--let's vote. See D.?? Possibilities 1 and 2.
- Made decision to Abort the Exchange in multiple error situations. See D.5? and D.5 ??
 - Too many corner cases
 - Fear (reality?) of data integrity exposures
 - Testing is bad enough for single error cases
 - Questionable value
- Added D.6a to show show another Exchange ambiguity case
- Added text to qualify the length of time context must be maintained after FCP_RESP to differentiate between Sequence level and Exchange level recovery
- Added that the use of REC to determine status for recovery in Class 2 is optional.


D.1 Class 3 Error Detection



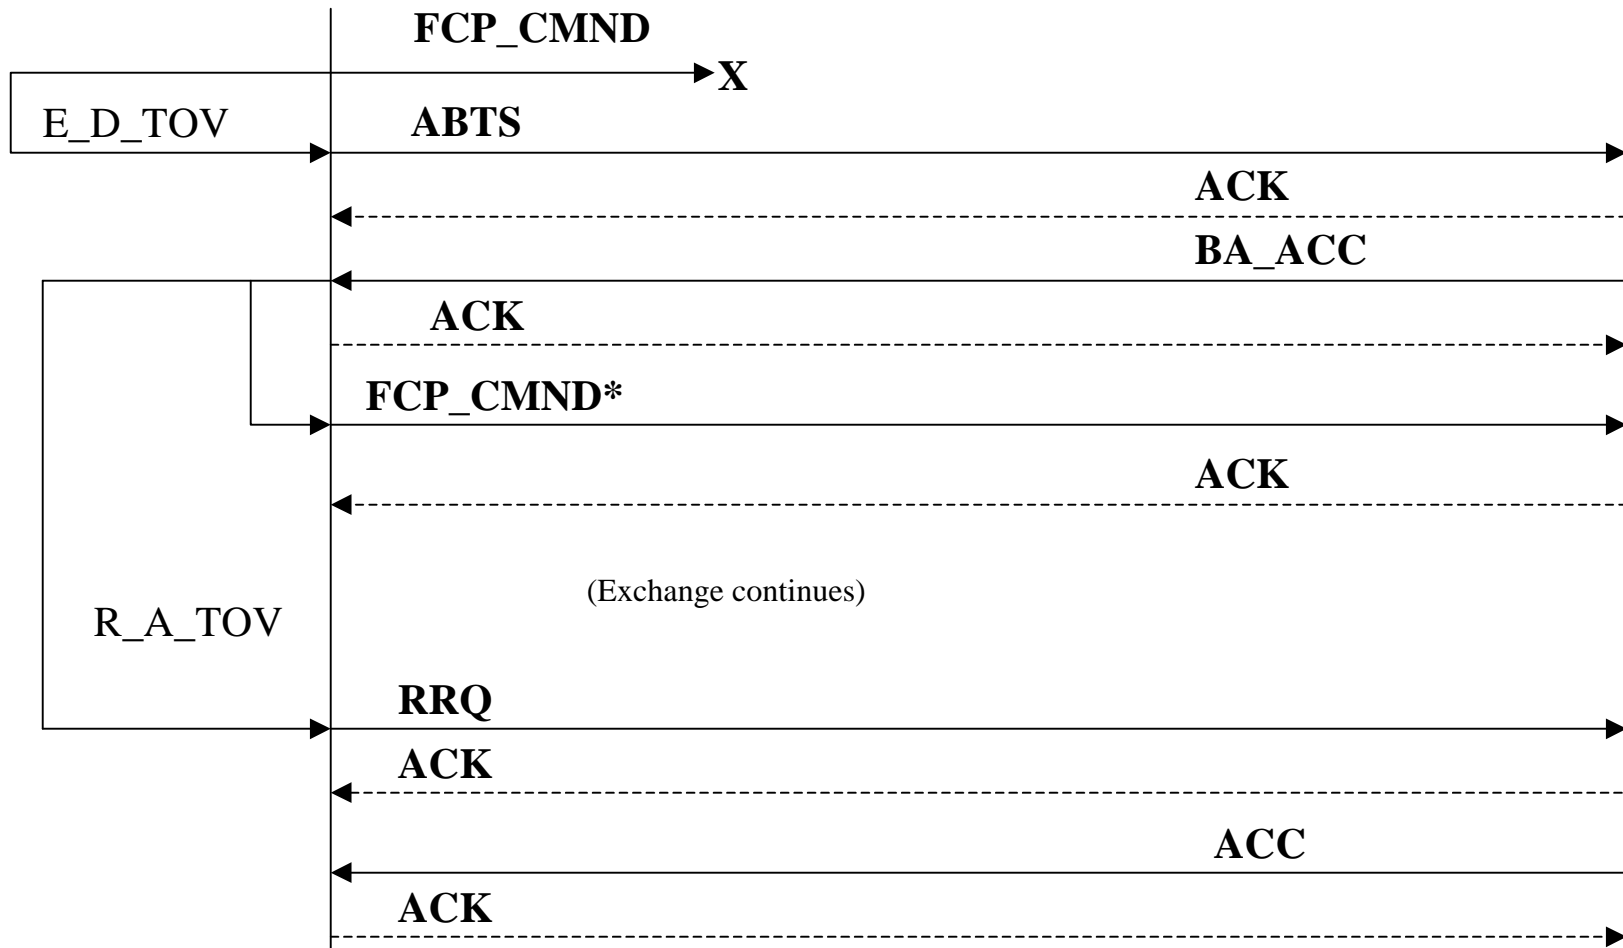
D.1 Class 2 Error Detection



D.2 Class 3 FCP_CMD Lost



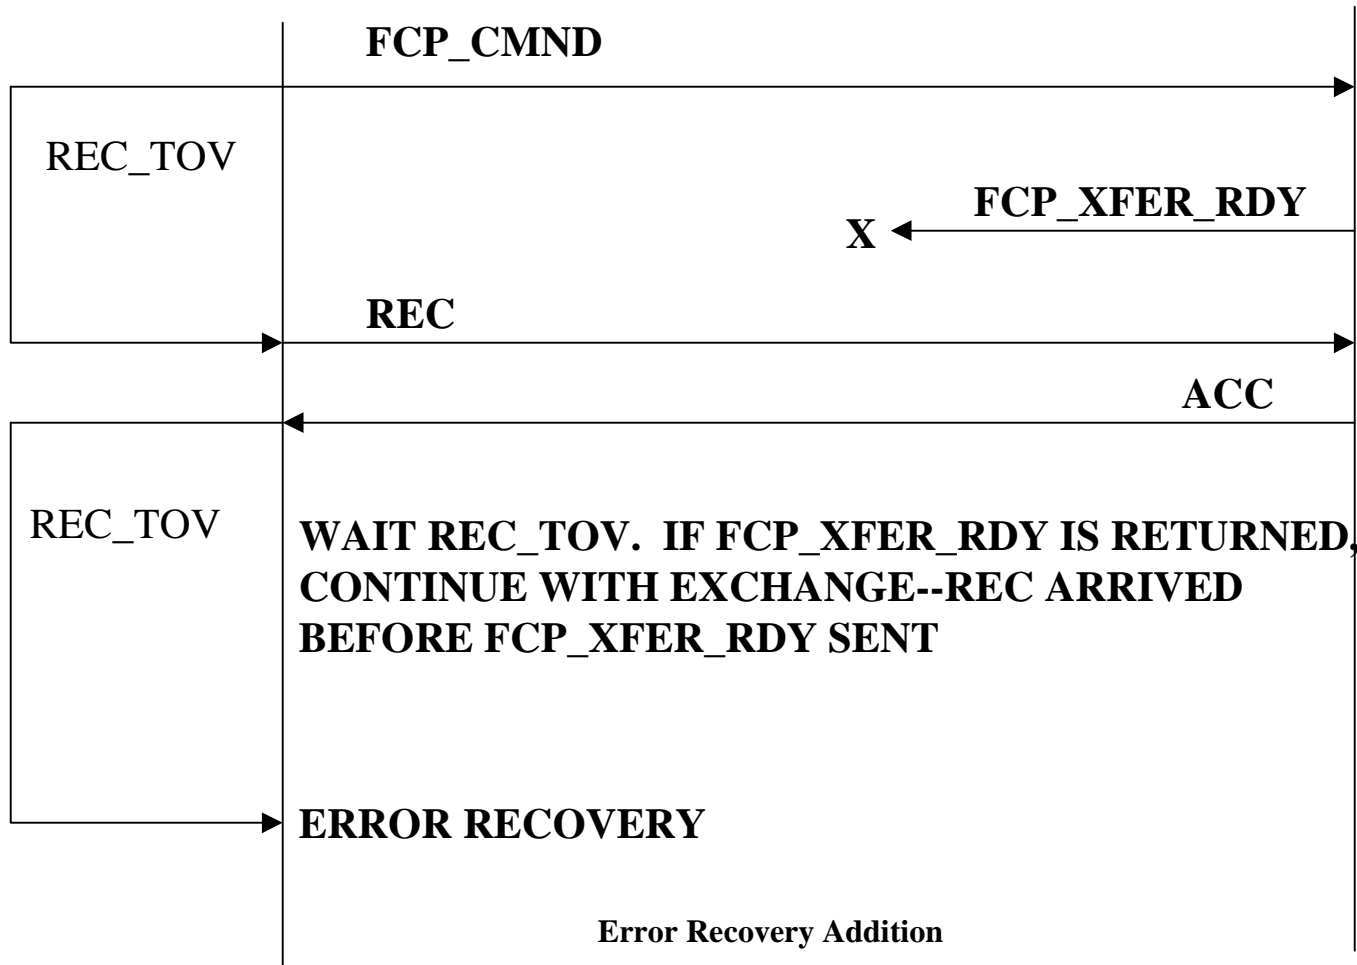
D.2 Class 2 FCP_CMD Lost



Error Recovery

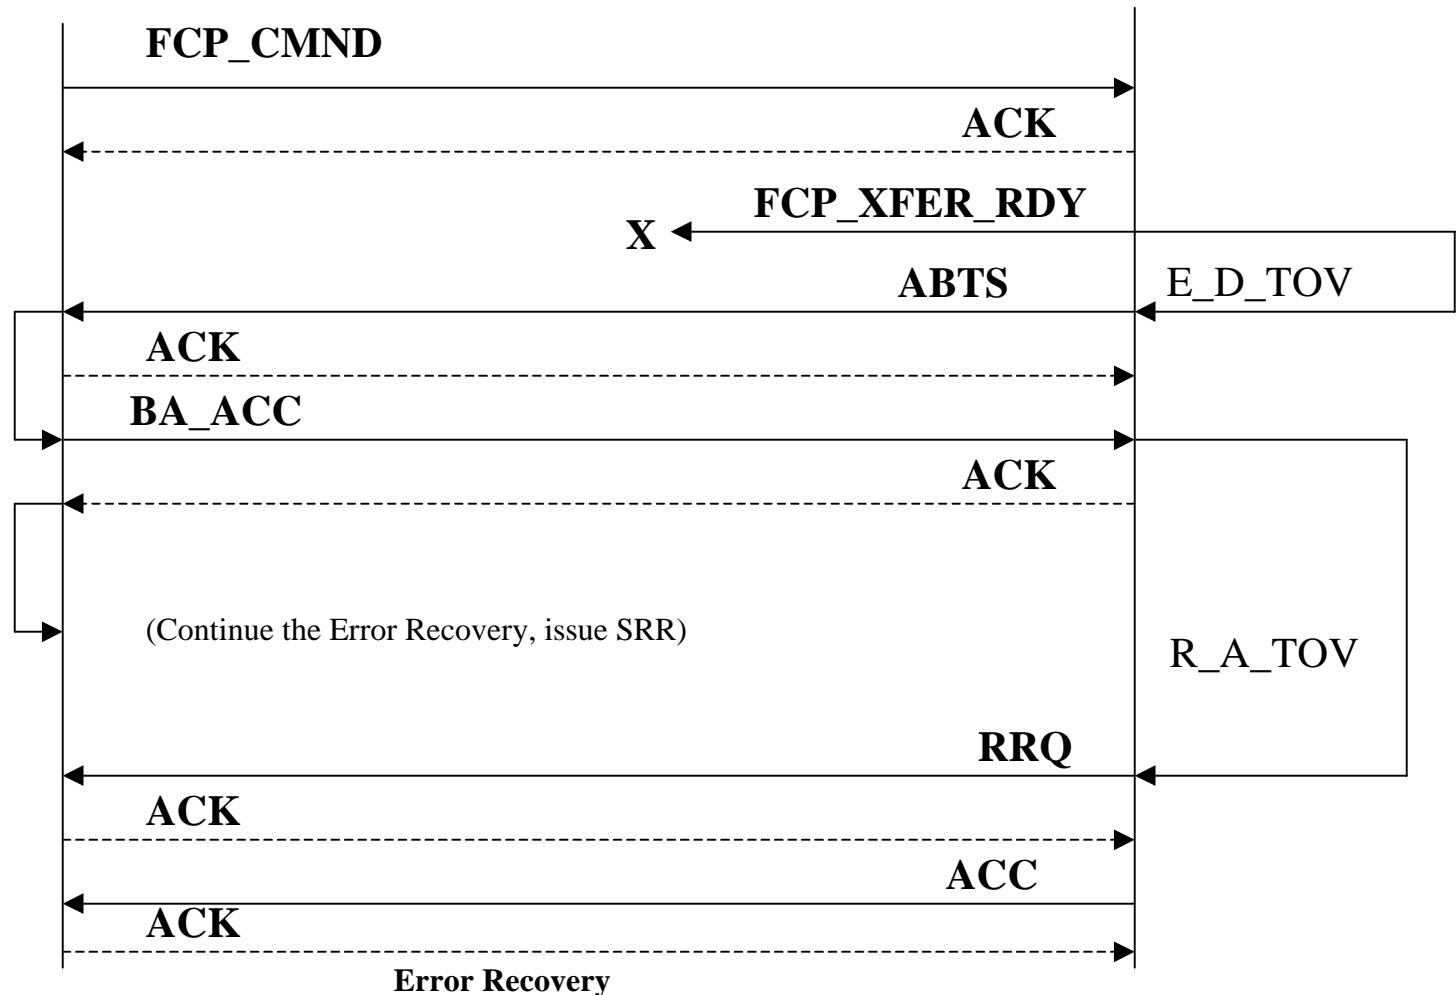
The BA_ACC indicates that the FCP_CMD was never received. (BA-ACC payload : SEQ_ID Validity = invalid, Low SEQ_CNT= 0, High SEQ_CNT= SEQ_CNT of ABTS frame.) Both the Initiator and Target establish Recovery Qualifiers. FCP_CMD* is resent in a new Exchange with a new SEQ_ID. The use of REC to determine status for error recovery is optional.

D.3 Class 3 FCP_XFER_RDY Lost



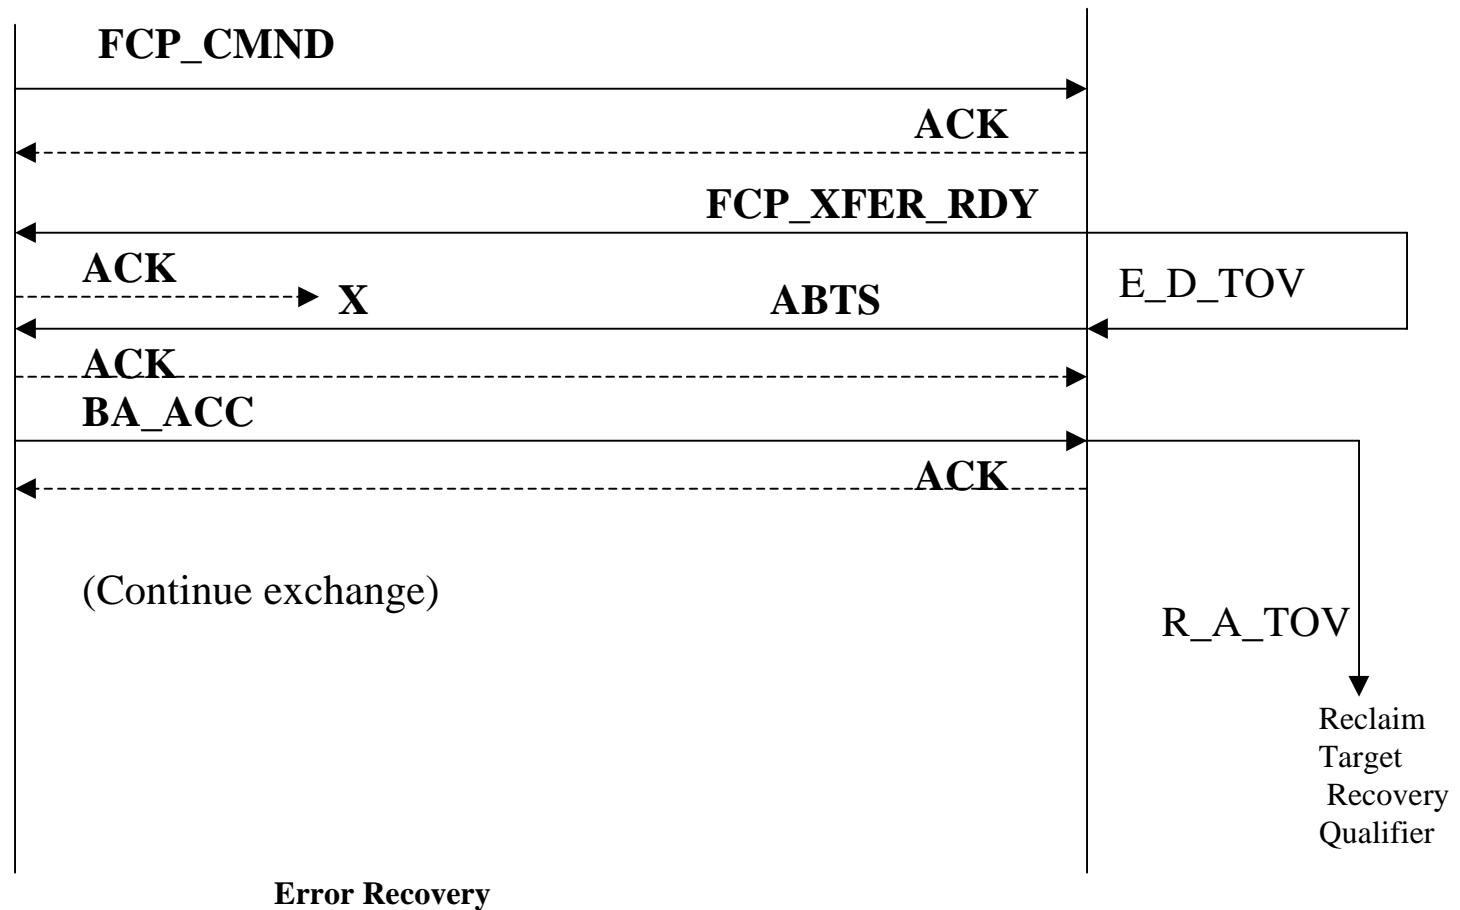
A new SEQ_ID must be used for the retransmission of the FCP_XFER_RDY.

D.3 Class 2 FCP_XFER_RDY Lost



For Class 2 the BA_ACC indicates that the FCP_XFER_RDY was never received by the Target. (BA_ACC Payload is SEQ_ID invalid, low SEQ_CNT=0, high SEQ_CNT =SEQ_CNT in ABTS frame = 1.) Both Target and Initiator must establish Recovery Qualifiers. A new SEQ_ID must be used in the retransmission of FCP_XFER_RDY and the SEQ_CNT value used must be one greater than the value used in the ABTS frame. The use of REC to determine status for error recovery is optional.

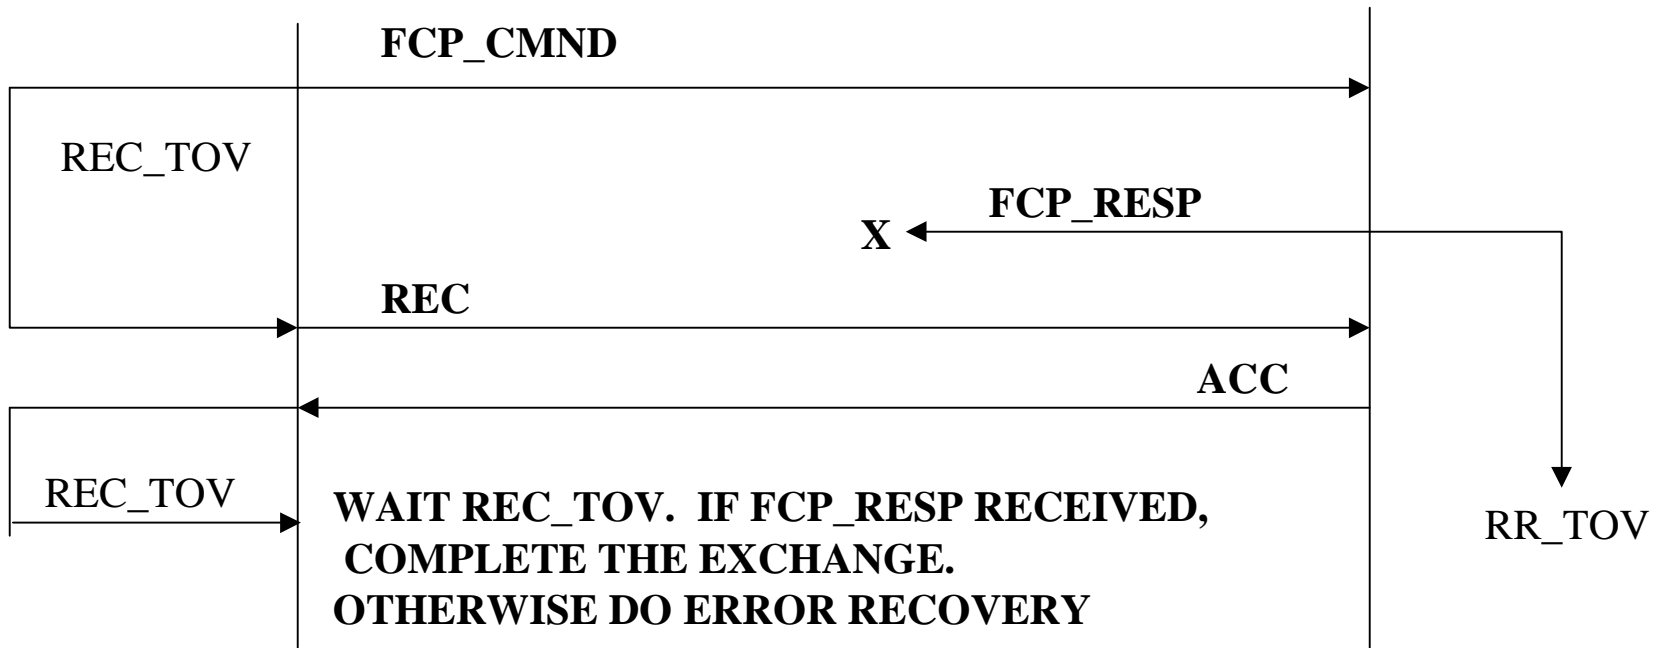
D.4 Class 2 FCP_XFER_RDY Rcvd, ACK Lost



None:

The BA_ACC indicates that the FCP_XFER_RDY was received by the Target, ACC Payload of SEQ_ID valid, SEQ-ID =SEQ_ID of the FCP_XFER_RDY frame, low SEQ_CNT=high SEQ_CNT=SEQ_CNT of the ABTS frame. No error recovery is required. Note: There is no need for the Target to issue the RRQ since no Recovery Qualifier was established by the Initiator in this case. It must still let R_A_TOV expire before reclaiming the resources associated with its Recovery Qualifier. The use of REC to determine status for error recovery is optional.

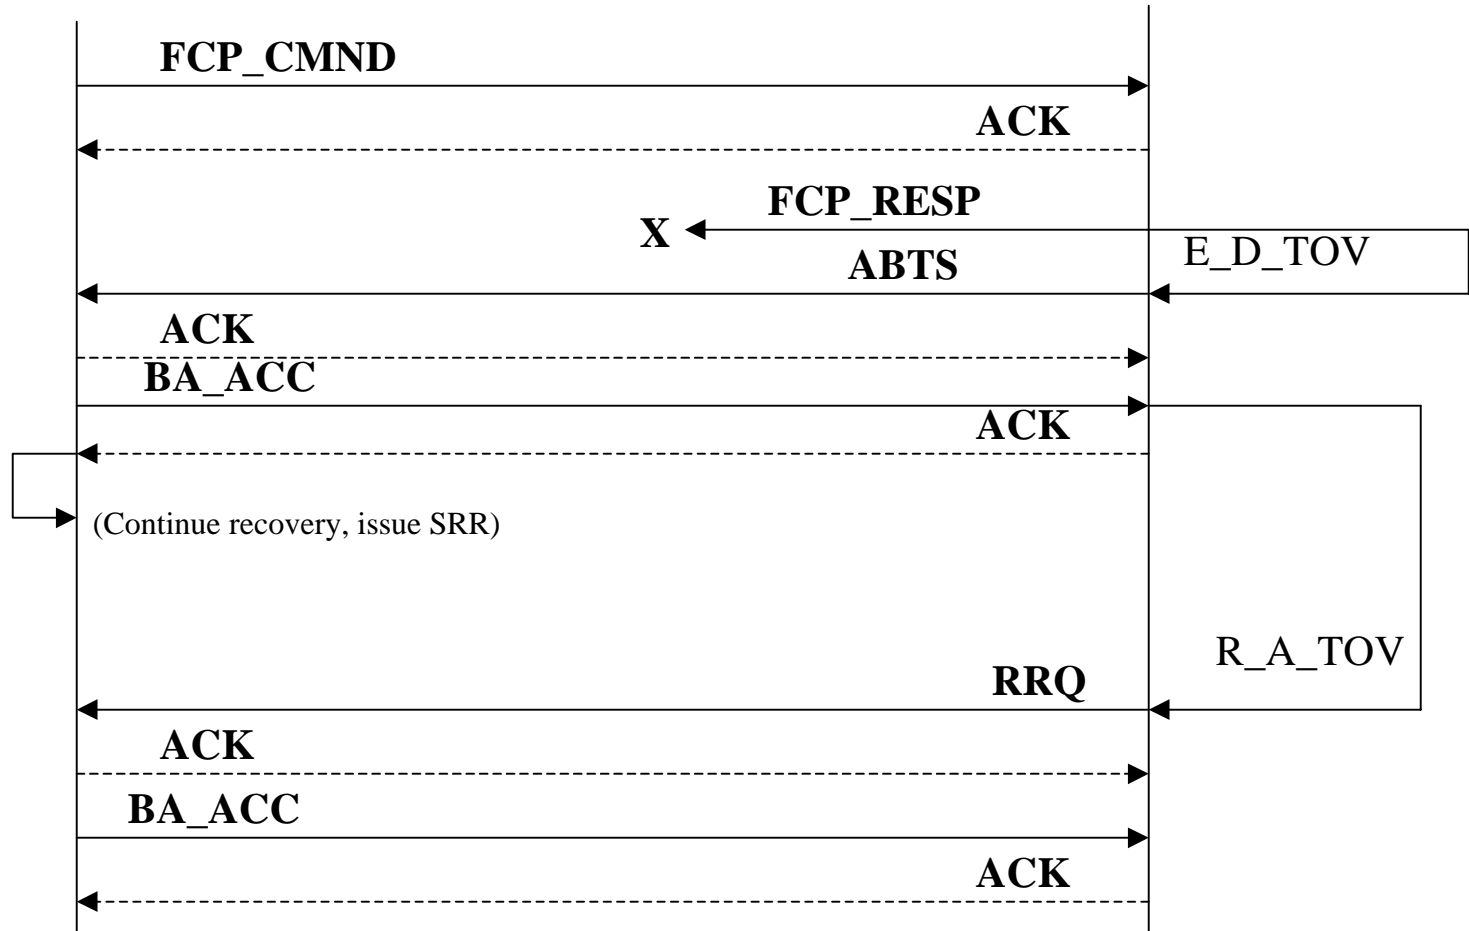
D.5 Class 3 FCP_RESP Lost, No FCP_CONF Req.



For Sequence level recovery, the Target must keep the context of this Exchange until the OX_ID value is reused in a new command, implicitly validating the receipt of FCP_RESP, or for at least RR_TOV, in order to preserve the FCP_RESP information. This long timeout can be avoided by using FCP_CONF. The FCP_RESP retransmission must use a new SEQ_ID.

For Exchange level recovery, the context of the Exchange can be purged by the Target after FCP_RESP is transmitted.

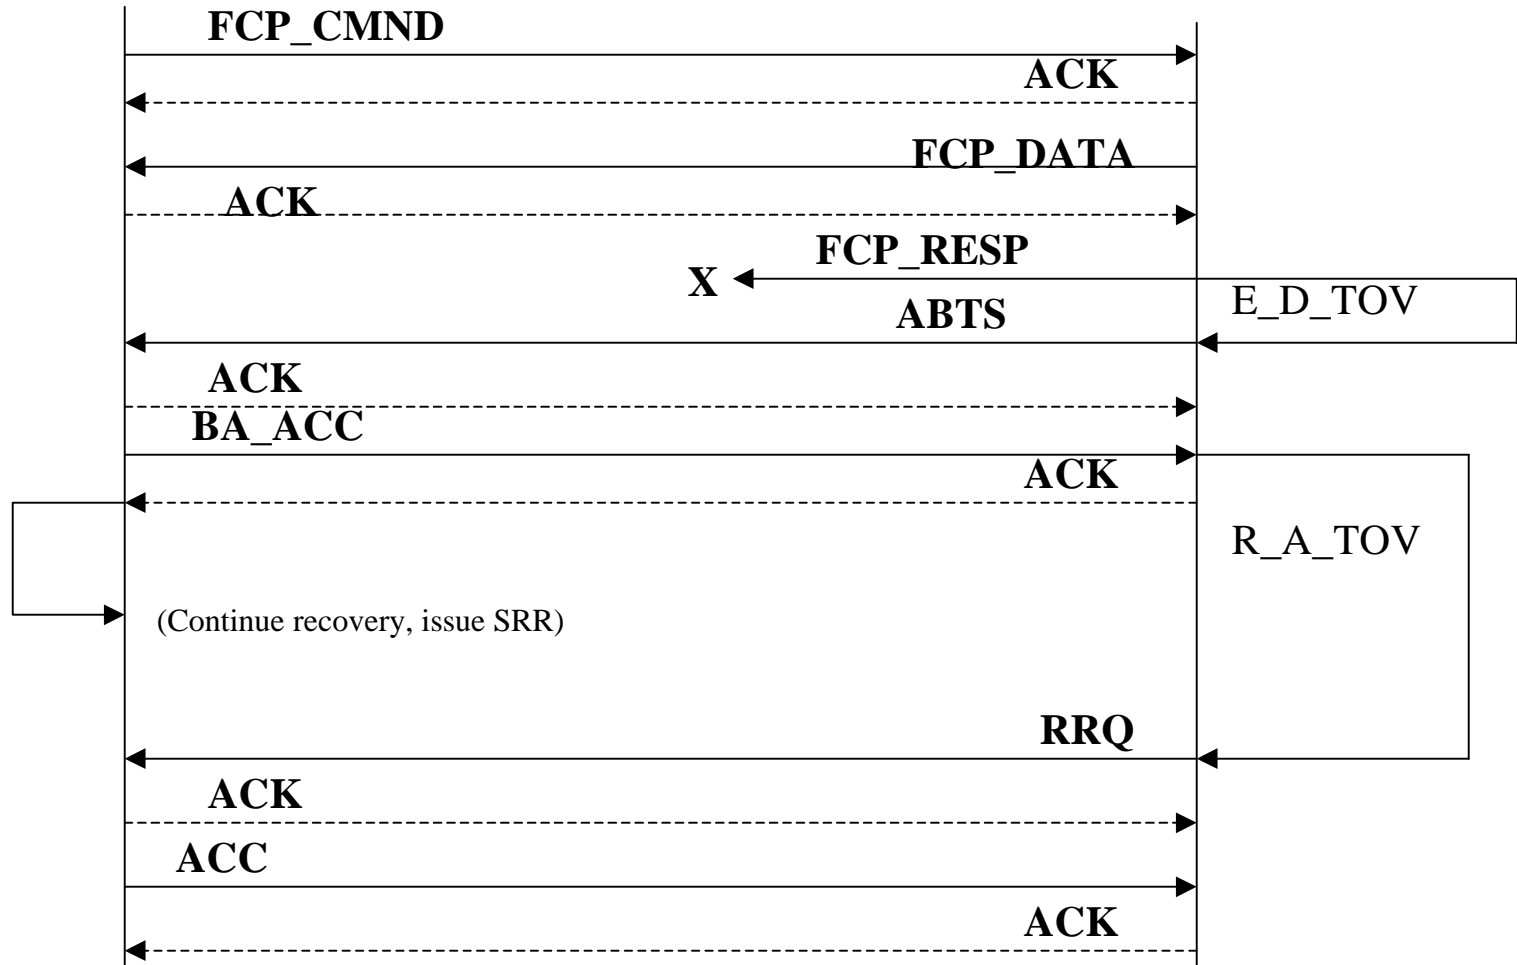
D.5 Class 2 FCP_RESP Lost, No FCP_CONF Req



Error Recovery

BA_ACC indicates that FCP_RESP was never received by the Initiator (Payload is SEQ_ID invalid, low SEQ_CNT=0, high SEQ_CNT=SEQ_CNT in ABTS frame=1.) Both Initiator and Target must establish Recovery Qualifiers. A new SEQ_ID must be used in the retransmission of FCP_RESP and the SEQ_CNT value used must be one greater than the value used in the ABTS frame. The use of REC to determine status for error recovery is optional.


D.5a Class 2 FCP_RESP Lost after a Write Cmnd, No FCP_CONF Req.



Error Recovery

BA_ACC indicates that FCP_RESP was never received by the Initiator (Payload is SEQ_ID valid, SEQ_ID=SEQ_ID of FCP_DATA Sequence, low SEQ_CNT=0, high SEQ_CNT=SEQ_CNT in ABTS frame=1.) Both Initiator and Target must establish Recovery Qualifiers. A new SEQ_ID must be used in the retransmission of FCP_RESP and the SEQ_CNT value used must be one greater than the value used in the ABTS frame. The use of REC to determine status for error recovery is optional.

D.6 Class 2 FCP_RESP Rcvd, ACK Lost

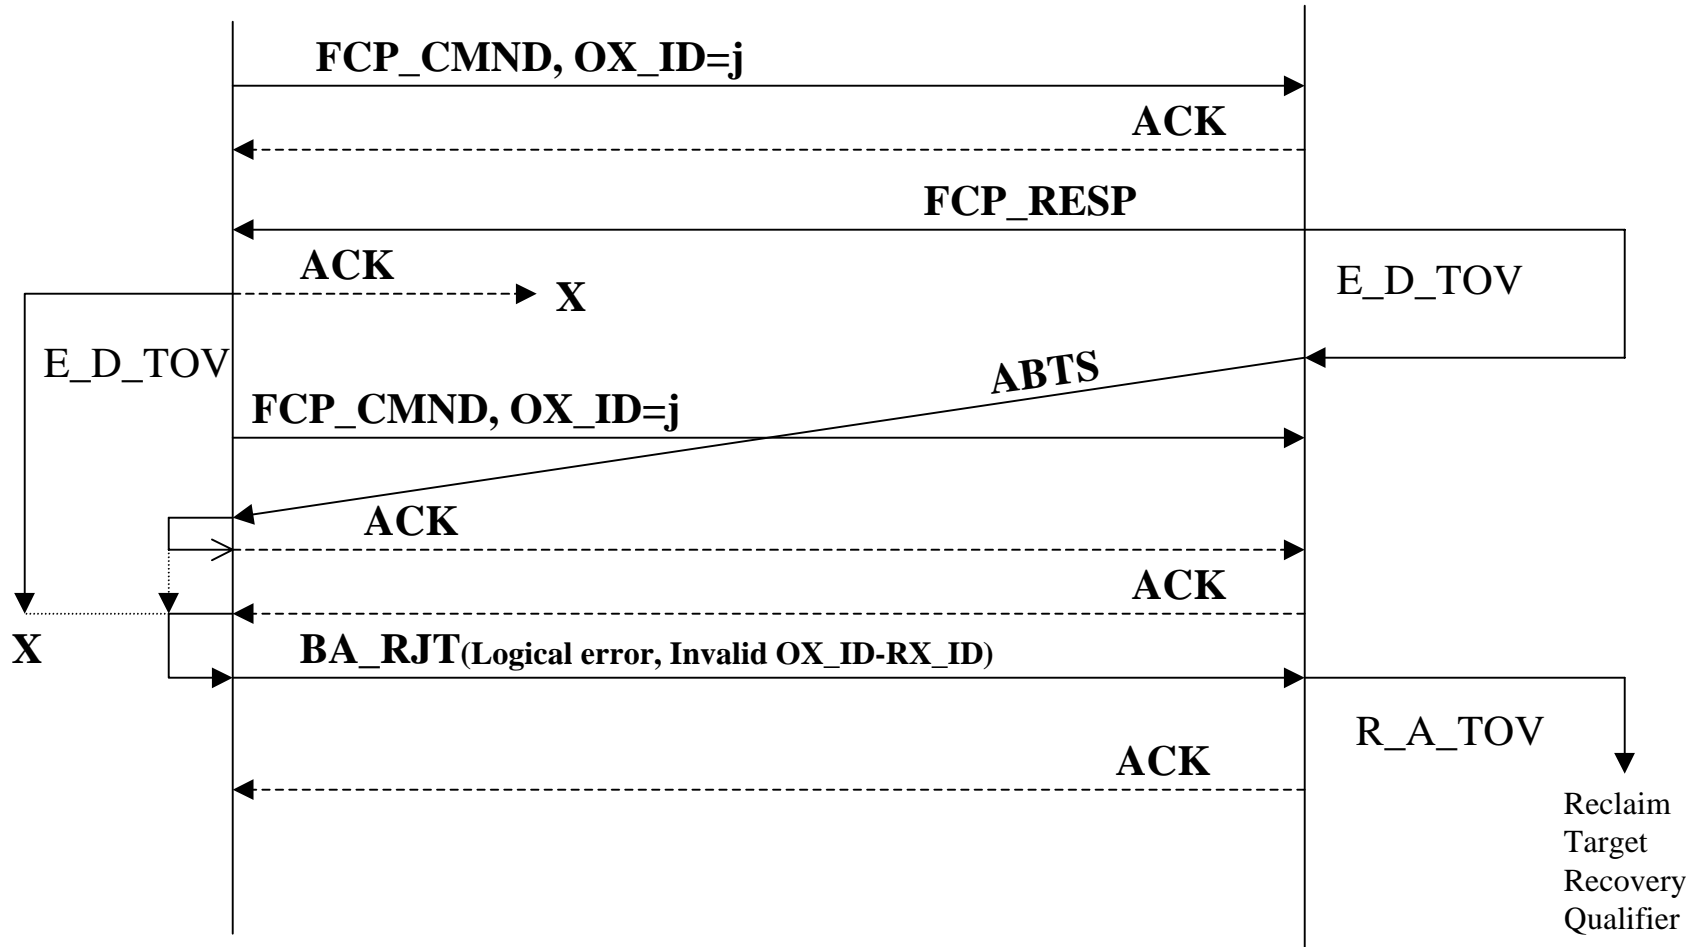


None:

The BA_RJT for the ABTS indicates the Exchange is unknown and therefore complete and no error recovery is required. The Target must establish a Recovery Qualifier. The associated resources cannot be reused for a period of R_A_TOV. (See D.5?? where the same OX_ID is used before the expiration of R_A_TOV)

Note: The Target need not issue RRQ as no Recovery Qualifier was established by the initiator.

D.6a Class 2 FCP_RESP Rcvd, ACK Lost, Case 2



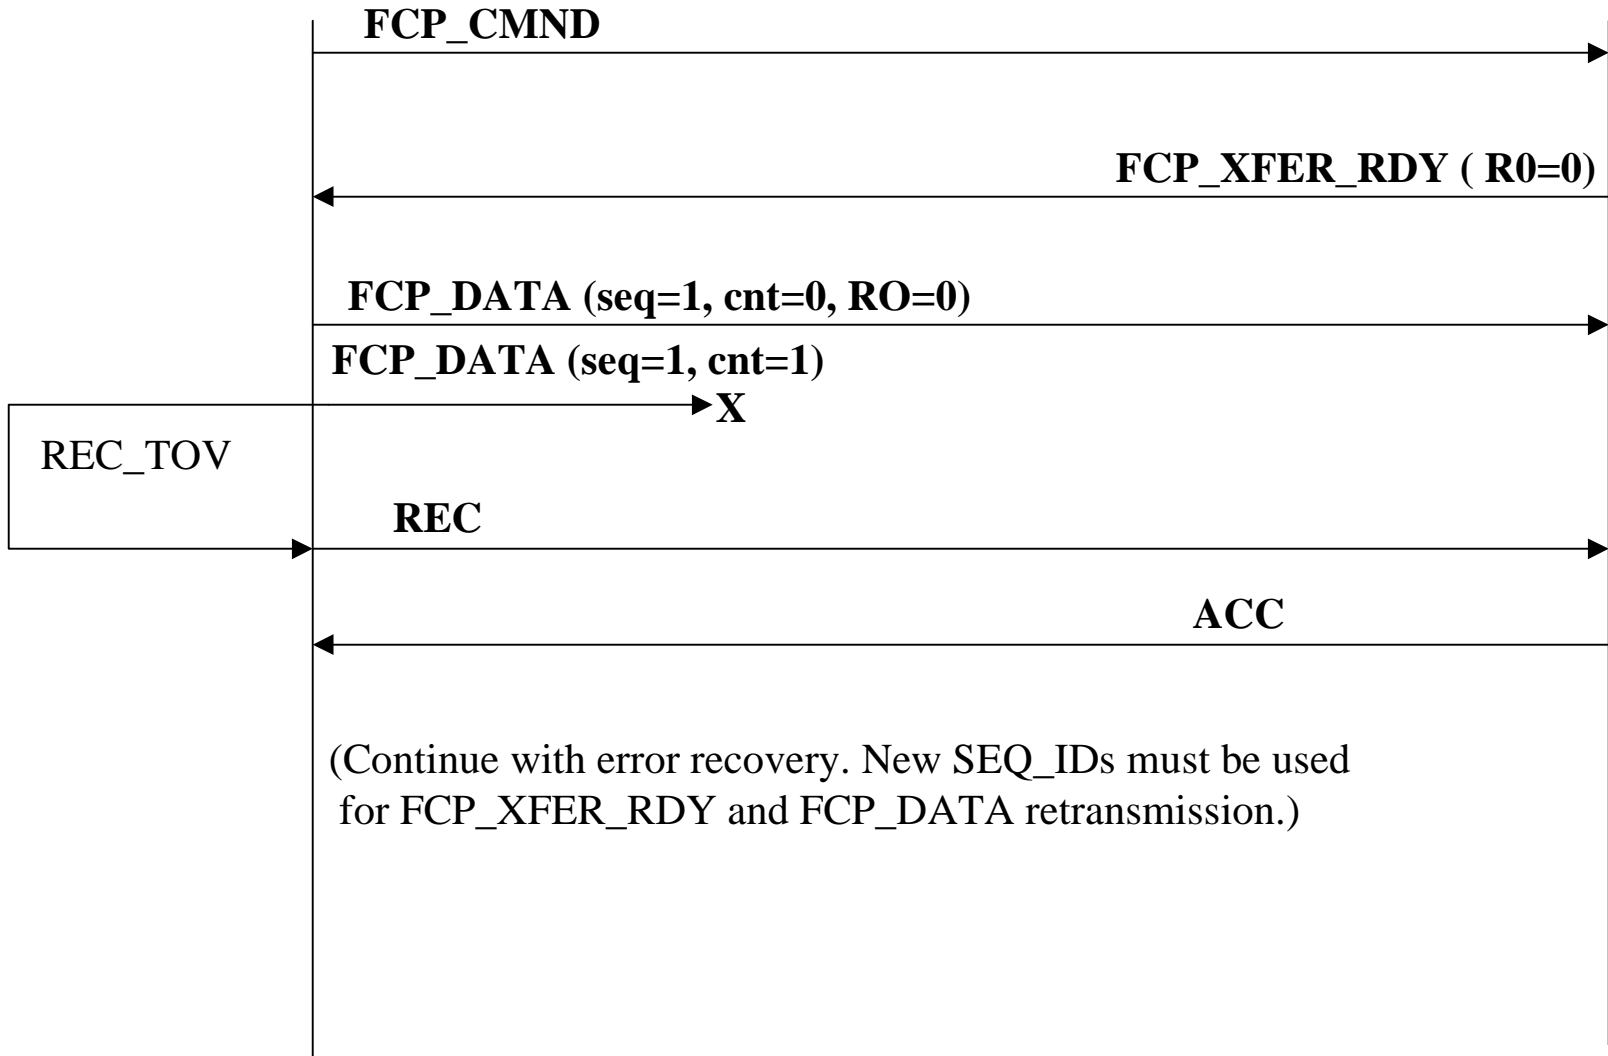
Error Recovery

None:

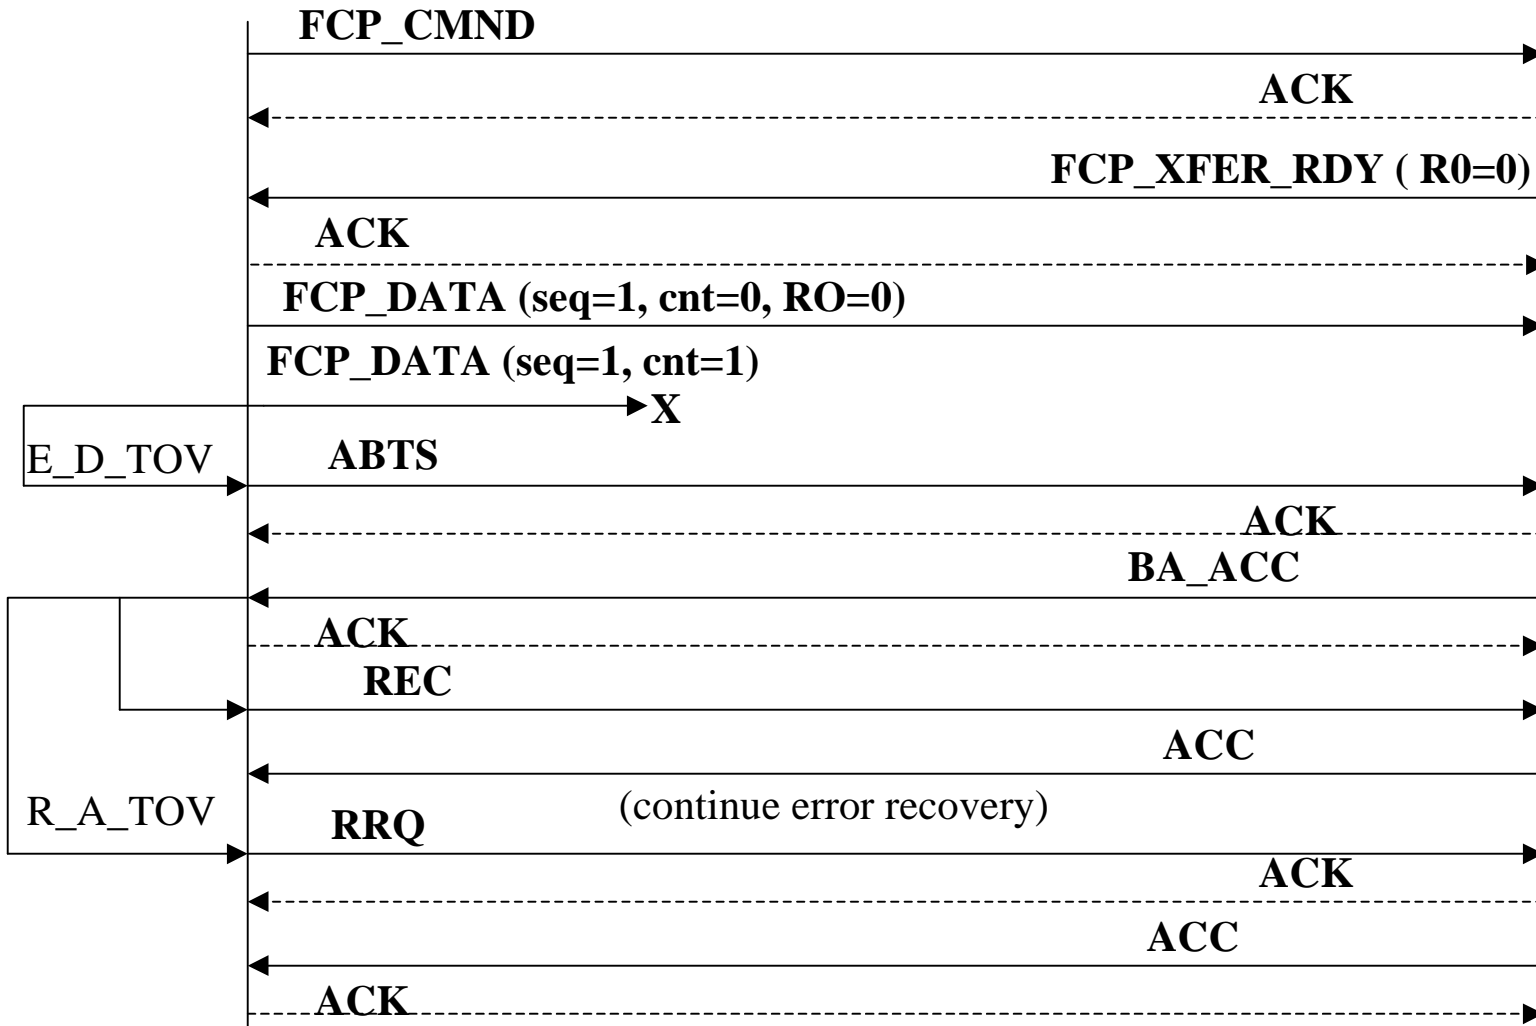
The BA_RJT for the ABTS indicates the Exchange is unknown and therefore complete and no error recovery is required. The Target must establish a Recovery Qualifier. The associated resources cannot be reused for a period of R_A_TOV. Note that no action is taken on ABTS until the ACK to the outstanding Sequence has been received, allowing the analysis to take into consideration the RX_ID, to eliminate ambiguity. (See D. 5?? For another case.)

Note: The Target need not issue RRQ as no Recovery Qualifier was established by the initiator.

D.7 Class 3 Lost Write Data, Last Frame of Seq.



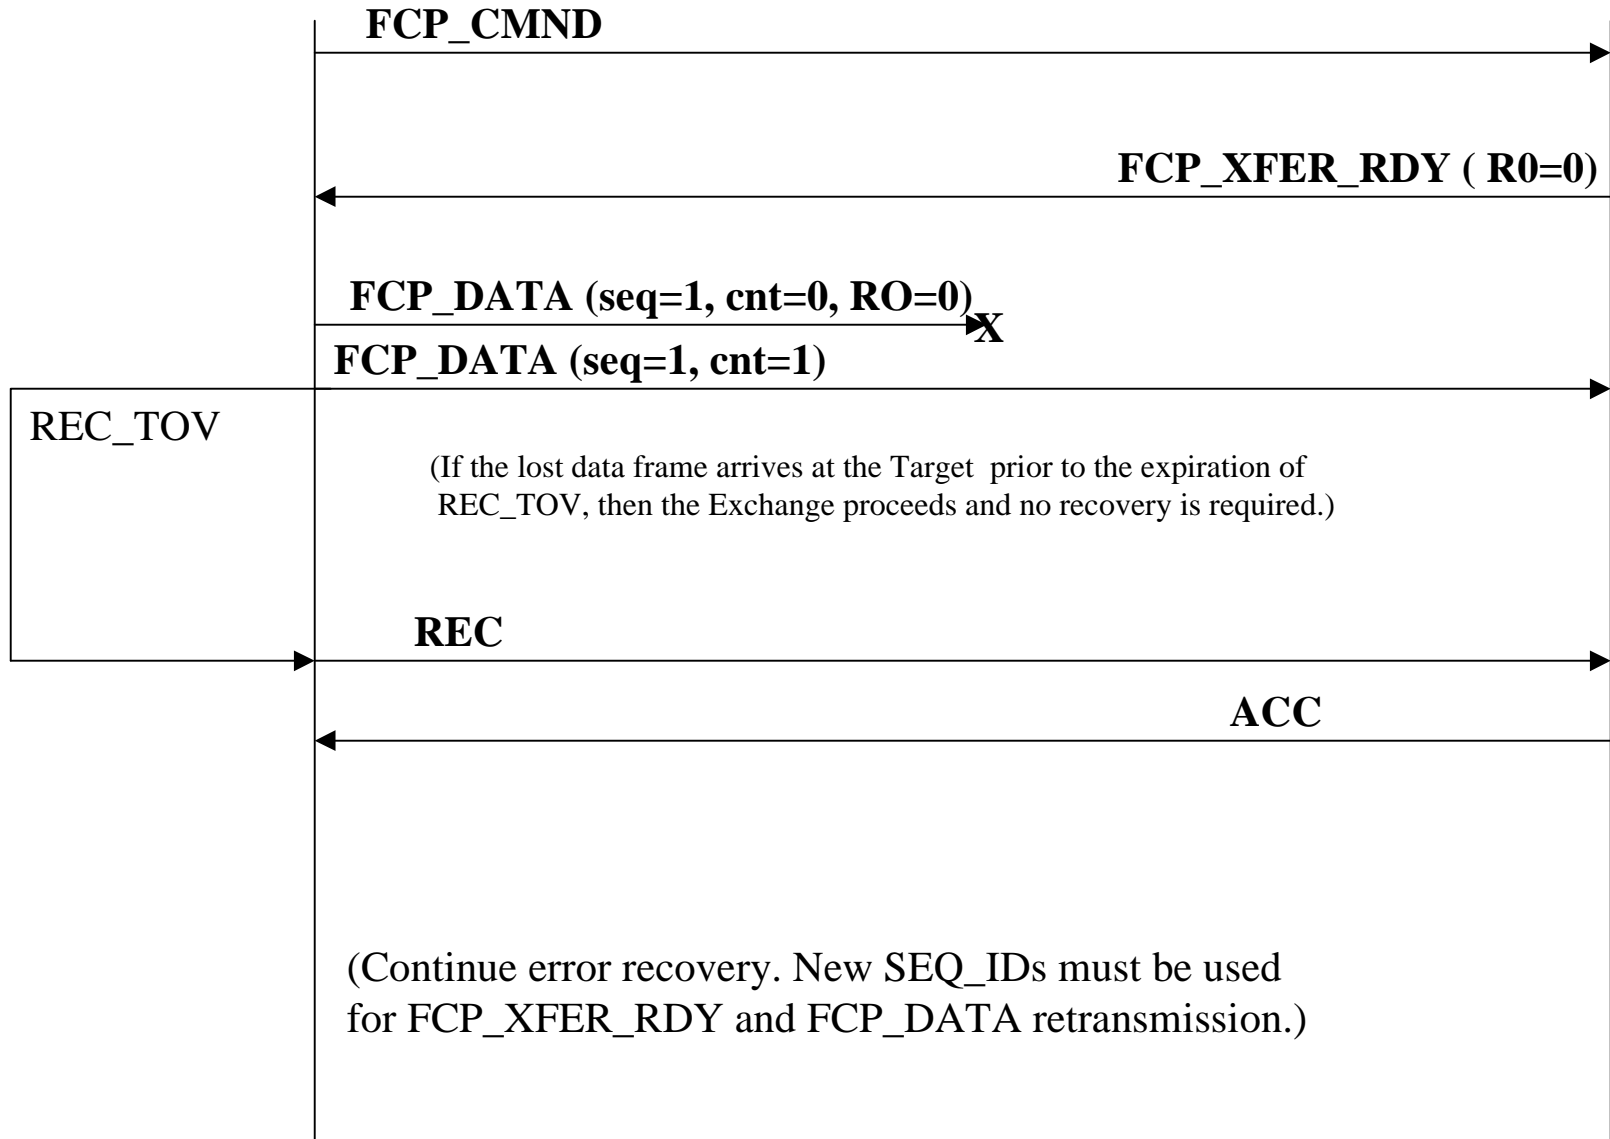
D.7 Class 2 Lost Write Data, Last Frame of Seq.




Error Recovery Addition

BA_ACC indicates by its payload, that FCP_DATA Sequence was not completely received (Payload: SEQ_ID valid, SEQ_ID value = SEQ_ID of FCP_CMND, low SEQ_CNT of 0 not equal to high SEQ_CNT of 2). New SEQ_IDs shall be used for the retransmission of FCP_XFER_RDY and FCP_DATA. For Class 2 the starting SEQ_CNT value used with FCP_DATA must one greater than the value used in ABTS. The ACKs for REC/ACC are not shown.

D.8 C1 3, Lost Write Data, Not Last Fr. of Seq.

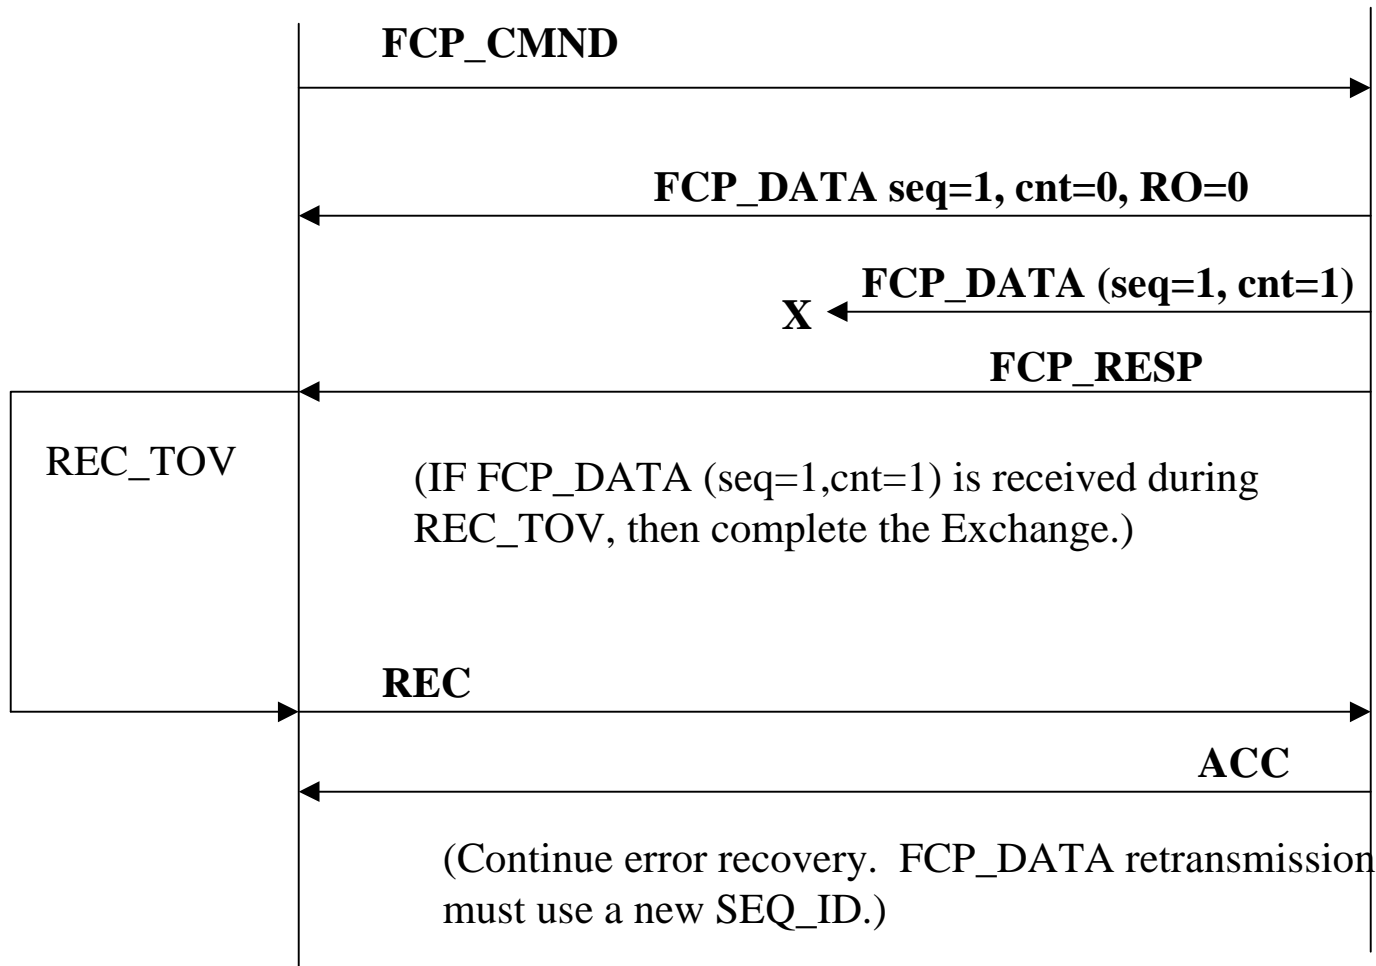


D.8 Class 2 Lost Write Data, Not Last Frame of Seq.

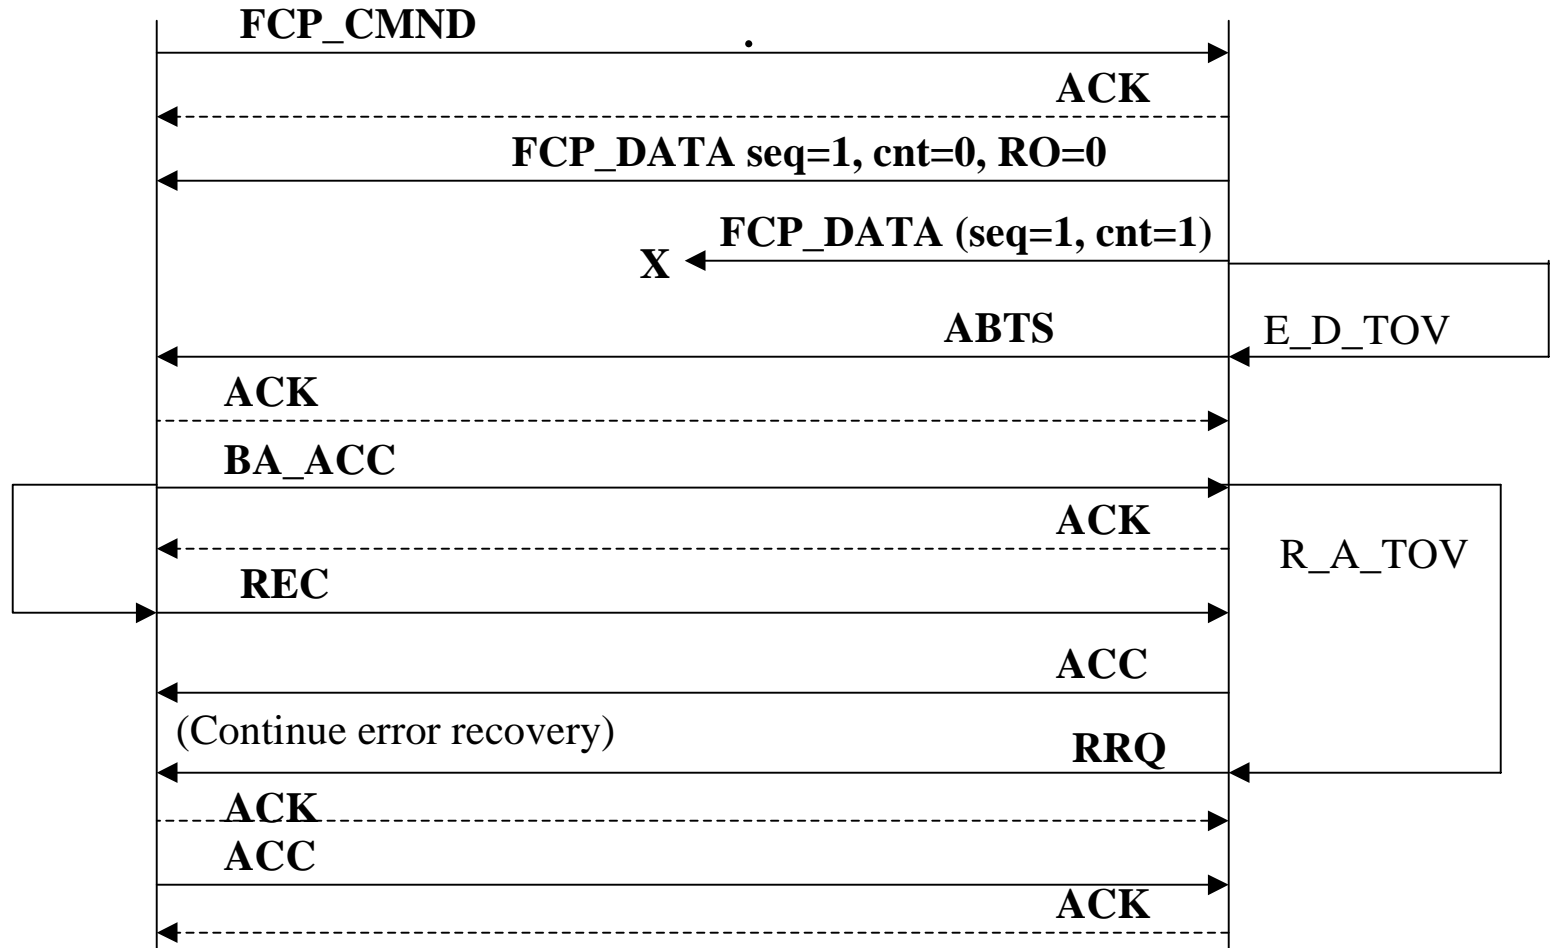


BA_ACC indicates by its payload, that the FCP_DATA Sequence was not completely received (Payload: SEQ_ID valid, SEQ_ID value = SEQ_ID of FCP_CMND, low SEQ_CNT of 0 not equal to high SEQ_CNT of 2). New SEQ_IDs shall be used for retransmission of FCP_XFER_RDY and FCP_DATA. For Class 2 the Sequence count value used with the retransmission of FCP_DATA must be one greater than the value used in ABTS. Note that if all data frames arrive at the Target prior to the expiration of E_D_TOV (out-of-order), then there is no error and no recovery is necessary. ACKs for REC/ACC are not shown.

D.9 Class 3 Lost Read Data, Last Frame of Seq.



D.9 Class 2 Lost Read Data, Last Frame of Seq



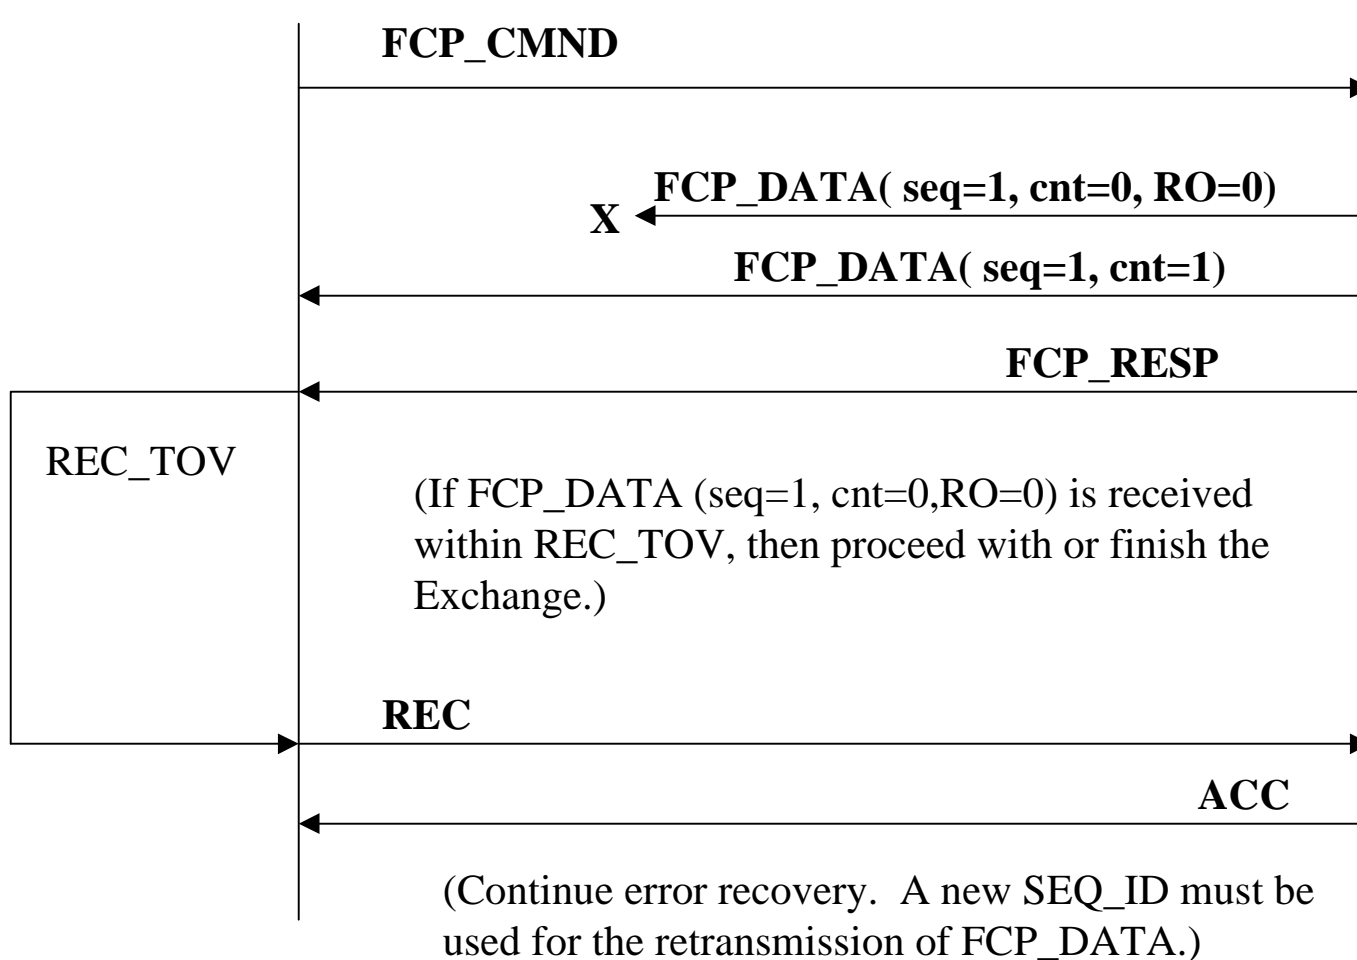
Error Recovery Addition

BA_ACC indicates by its payload, that the FCP_DATA Sequence was not completely received (Payload: SEQ_ID invalid, low SEQ_CNT of 0 not equal to high SEQ_CNT of the ABTS frame of 2).


A new Sequence ID shall be used for retransmission of FCP_DATA.

For Class 2, the Sequence count value used with the retransmission of FCP_DATA shall be one greater than the value used in ABTS. The ACKs for REC/ACC are not shown.

D.10 Class 3 Lost Read Data, Not Last Frame of Seq

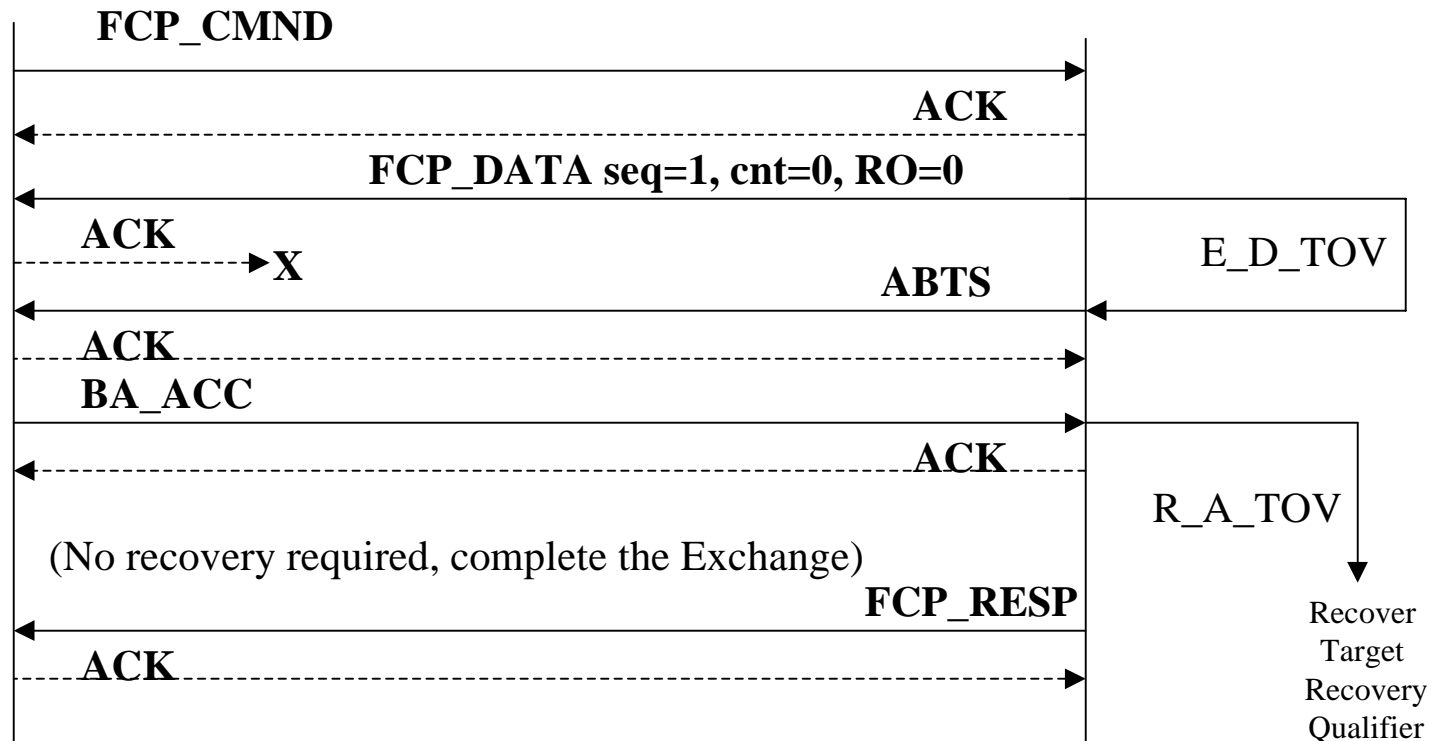


D.10 Class 2 Lost Read Data, Not Last Frame of Seq



BA_ACC indicates by its payload, that the FCP_DATA Sequence was not completely received (Payload: SEQ_ID invalid, low SEQ_CNT of 0 not equal to high SEQ_CNT of the ABTS frame of 2). A new Sequence ID shall be used for retransmission of FCP_DATA. For Class 2, the SEQ_CNT used with the retransmission of FCP_DATA shall be one greater than the value used in ABTS. Note that if all data frames arrive at the initiator before E_D_TOV expires, then no recovery is required; a frame or frames arrived out-of-order. The ACKs for REC/ACC are not shown.

D.11 Class 2 ACK Lost on Read



Error Recovery

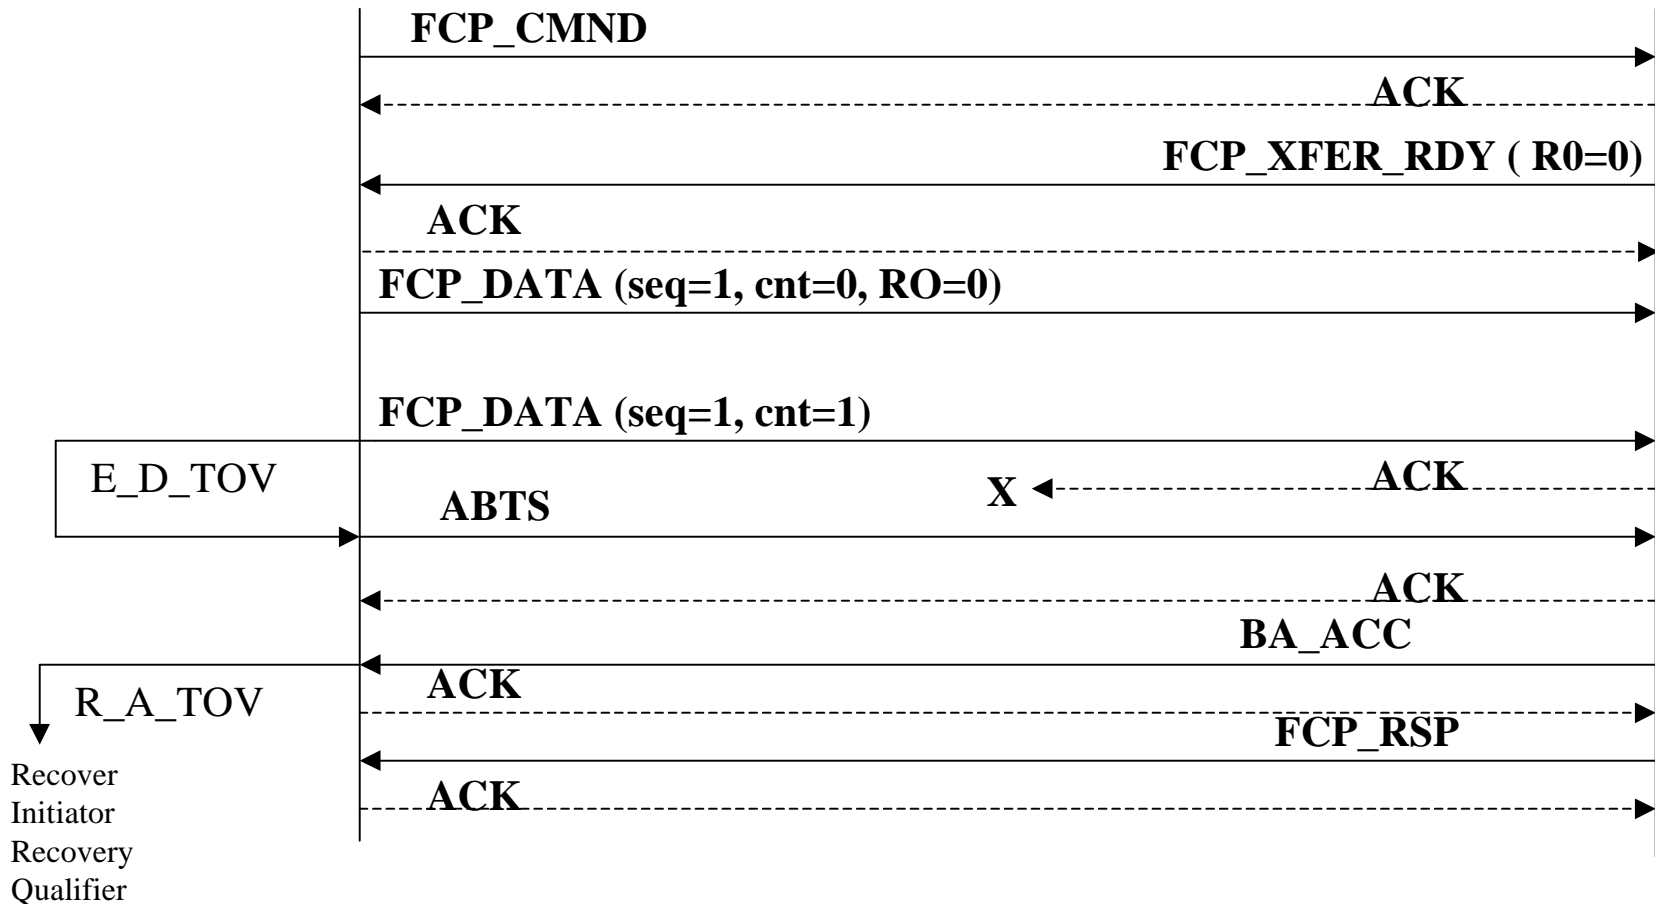
None:

The initiator has received the FCP_DATA Sequence. No error recovery is required.

Note: The BA_ACC indicates the FCP_DATA sequence was received (Payload is SEQ_ID valid, SEQ_ID value =SEQ_ID value of FCP_DATA Sequence, low SEQ_CNT=high SEQ_CNT= SEQ_CNT of ABTS frame).

Note: The Target must establish its Recovery Qualifier. The resources associated with the Recovery Qualifier can be reclaimed after R_A_TOV. The issuance of RRQ is optional as no Recovery Qualifier was established by the Initiator in this case. FCP_RESP can be received anytime after the transmission of FCP_CMD due to out-of-order delivery.

D.12a Class 2 ACK Lost on Write

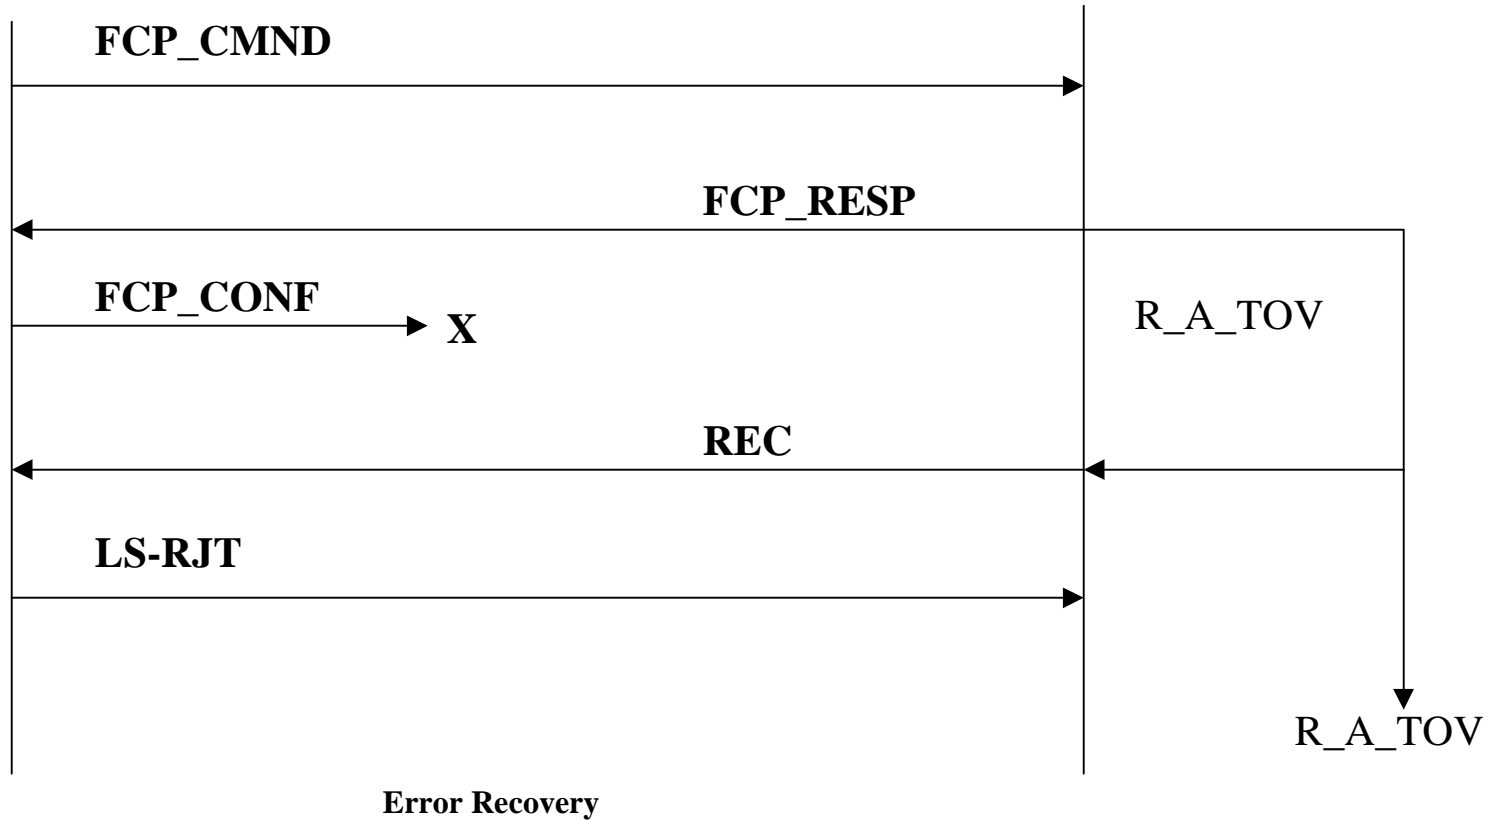


Error Recovery

None: The Target received the FCP_DATA sequence. No error recovery is required.

Note: The BA_ACC indicates the data sequence was received (Payload is SEQ_ID valid, SEQ_ID value =SEQ_ID value of the FCP_DATA Sequence, low SEQ_CNT=high SEQ_CNT=SEQ_CNT of ABTS frame). The Target and Initiator continue the Exchange. The Initiator must establish its Recovery Qualifier. The resources associated with the Recovery Qualifier can be reclaimed after R_A_TOV. The issuance of the RRQ is optional as no Recovery Qualifier was established by the Target. FCP_RESP can be received at any time after the last FCP_DATA frame has been transmitted.

D.?1 Class 3 FCP_CONF Lost

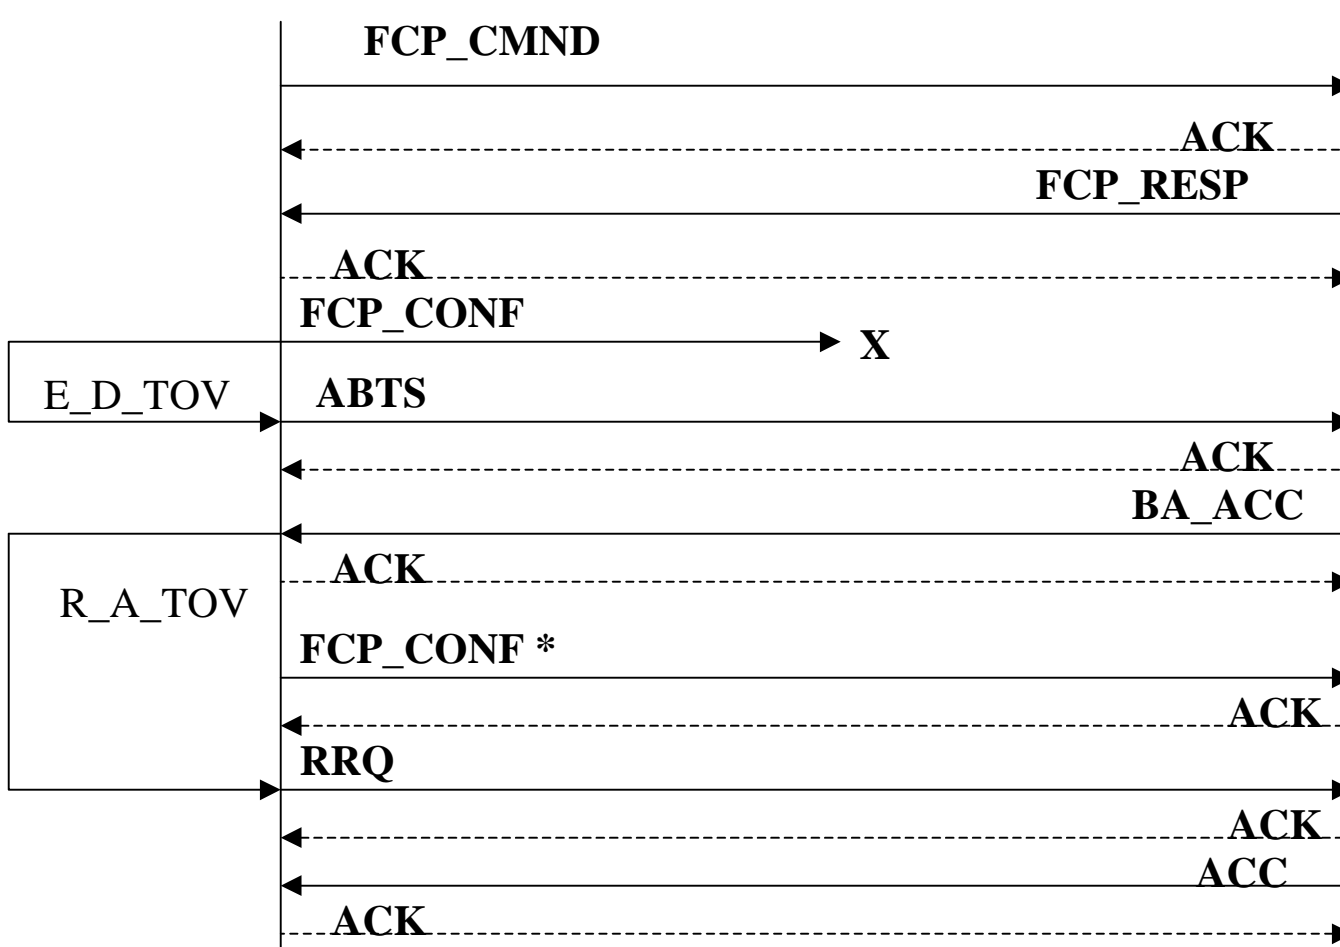


None.

LS-RJT implicitly indicates that the Initiator received FCP_RESP and sent FCP_CONF, since no context exists for the Exchange.

The context for the Exchange in the Target must be preserved for another R_A_TOV to prevent possible aliasing. If FCP_CONF is received after REC is transmitted, it is accepted and discarded and the context for the Exchange can be purged.

D.?? Class 2 FCP_CONF Lost(Possibility 1)

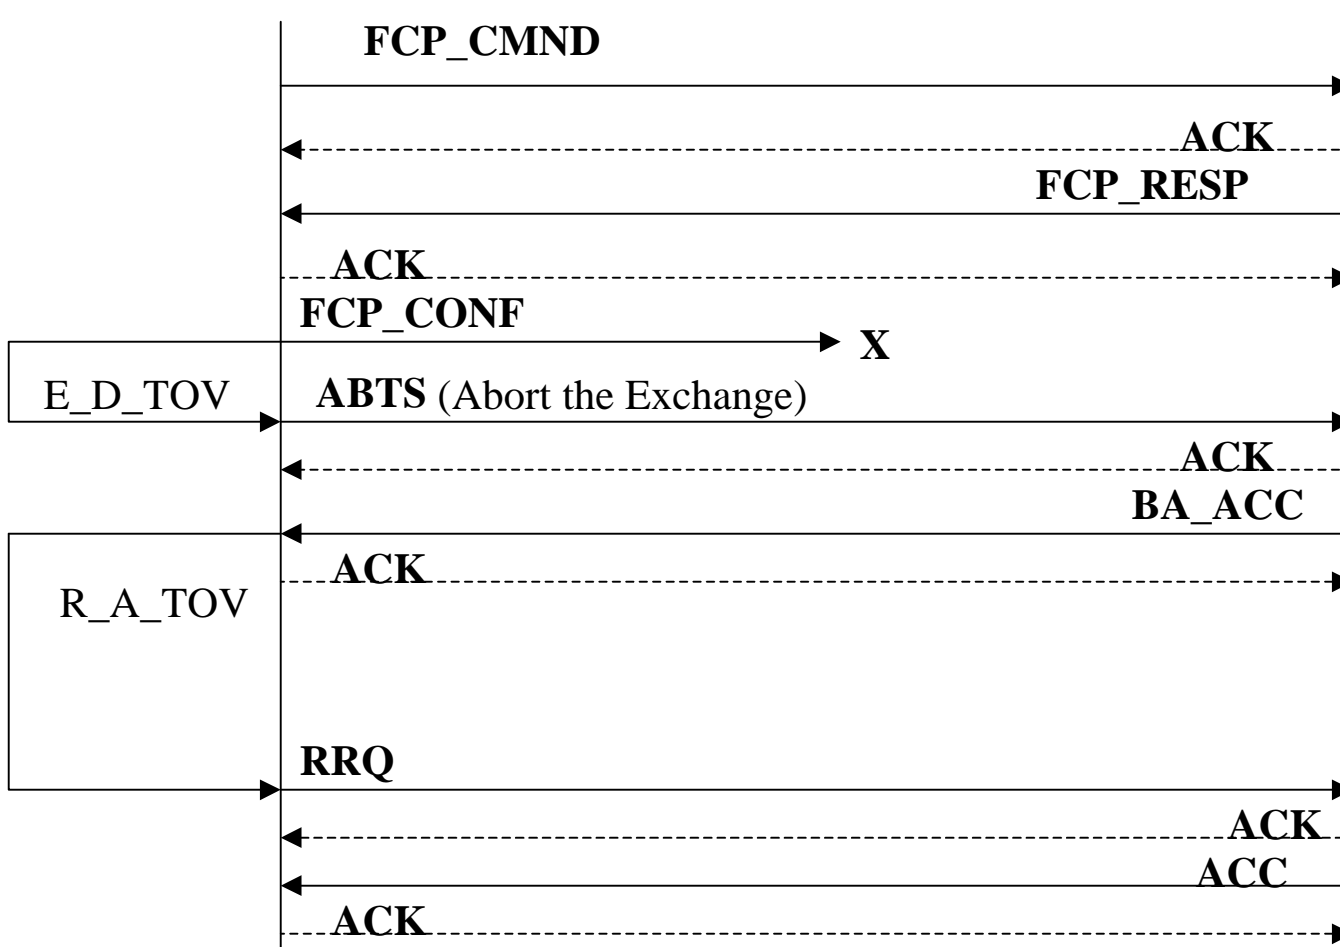


Error Recovery

BA_ACC payload indicates that FCP_CONF was not received (low SEQ_CNT=0, not equal to high SEQ_CNT =1, SEQ_ID valid, SEQ_ID value =SEQ_ID of FCP_CMND)

* Second FCP_CONF must be sent with a different SEQ_ID and the SEQ_CNT value must be one greater than the value used in the ABTS frame.


D.?? Class 2 FCP_CONF Lost(Possibility 2)



Error Recovery

Since the Initiator has received FCP_RESP, the Initiator can Abort the Exchange.

D.??? Class 2 ACK Lost on FCP_CONF

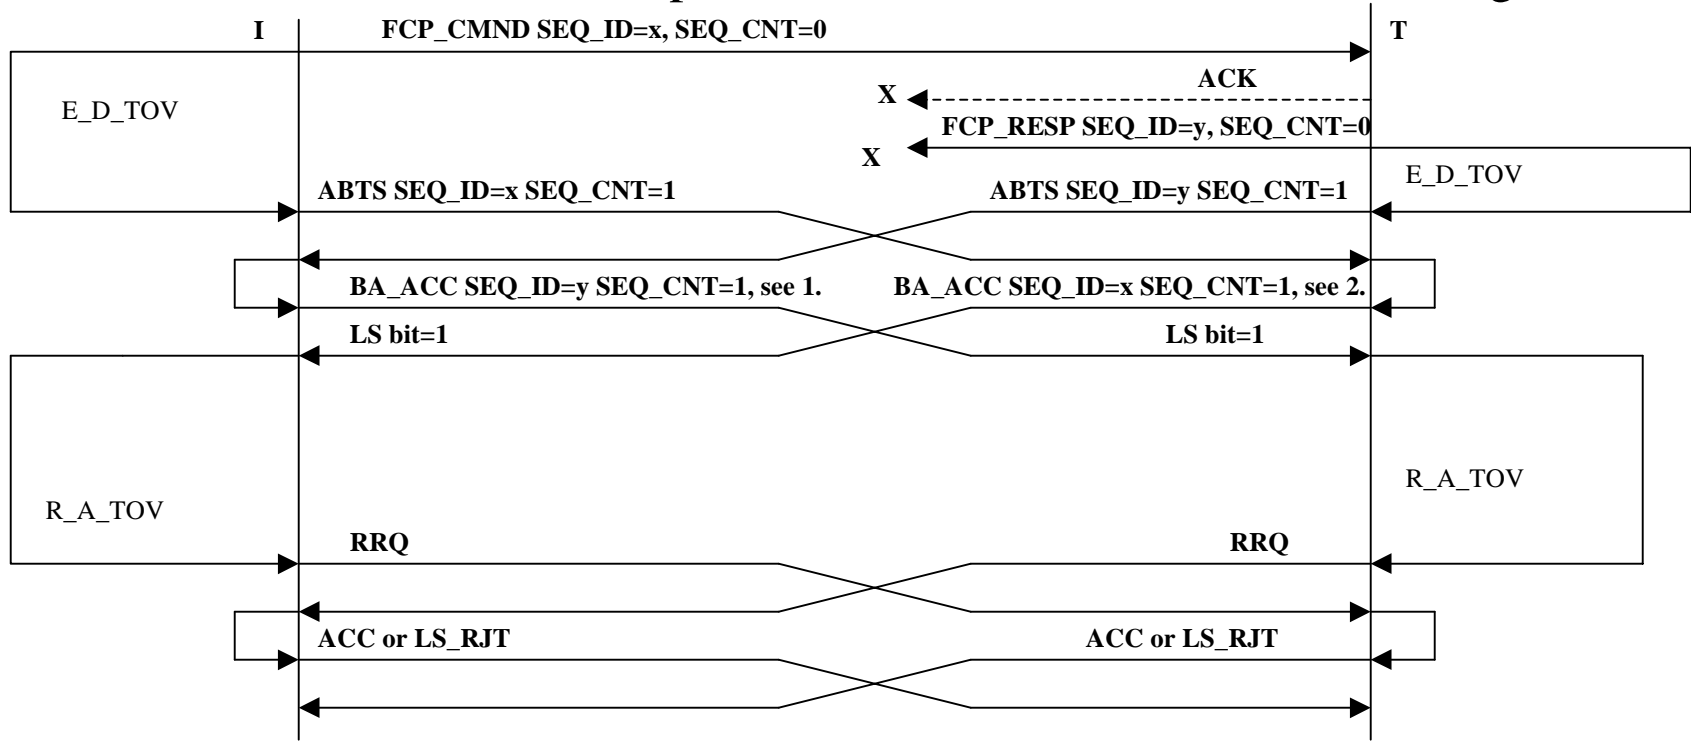


Reclaim
Initiator
Recovery
Qualifier

None:

BA_RJT is the response to the ABTS since no context exists for this Exchange and the ABTS was not issued on the first Sequence of a new Exchange. The Initiator must establish a Recovery Qualifier on receipt of the BA_RJT. The resources associated with the Recovery Qualifier can be reclaimed when R_A_TOV expires. Note that the issuance of RRQ is optional as no Recovery Qualifier was established by the Target.

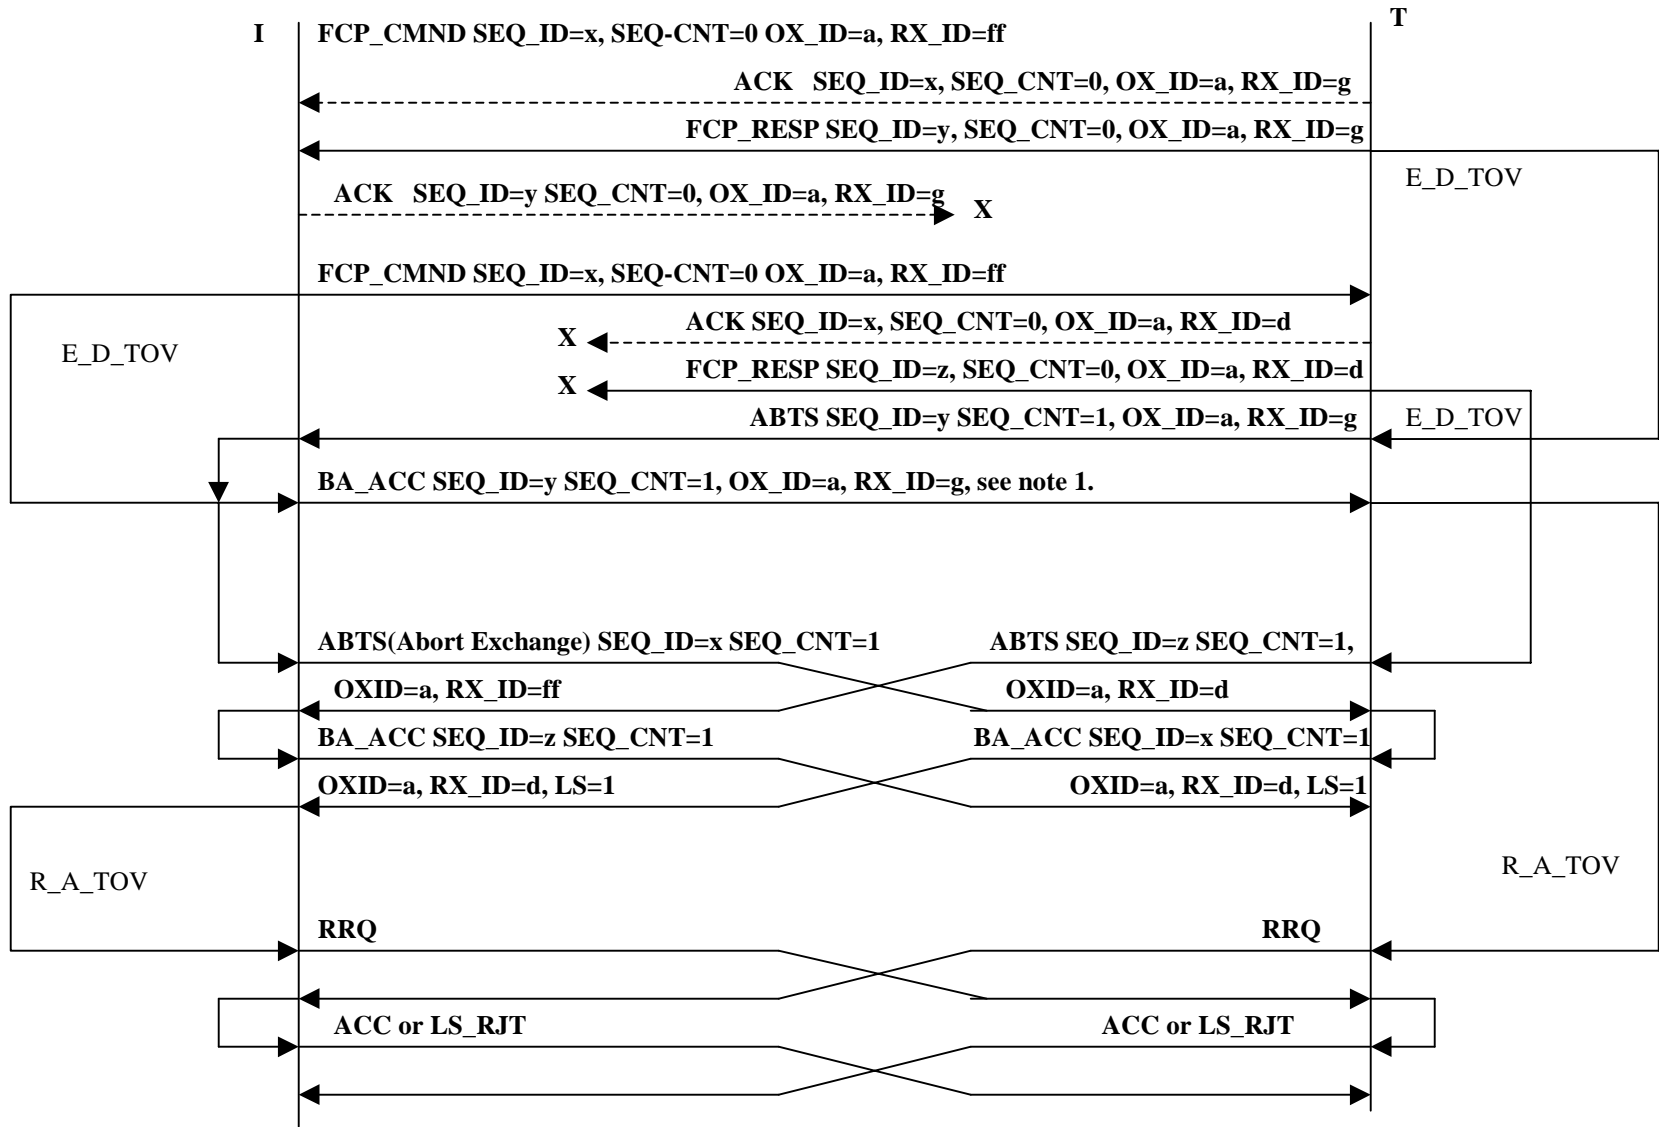
D.5? Class 2 Multiple Error Condition, Abort the Exchange



Error Recovery

Error recovery shall not be attempted if multiple errors, or the appearance of multiple errors, have occurred in an Exchange. A multiple error arises when a Recovery Qualifier has been established and then either an ABTS is received or the recovery action indicates the need to send an ABTS. The Exchange shall be aborted by issuing ABTS(Abort Sequence) if called for in the recovery process or by setting the Last_Sequence bit to one in BA_ACC with payload of SEQ_ID Validity = invalid, Low SEQ_CNT=hex '0000', High SEQ_CNT= hex 'FFFF', or both.

D.5?? Mutliple Error Condition, Exchange Ambiguities, Abort the Exchange

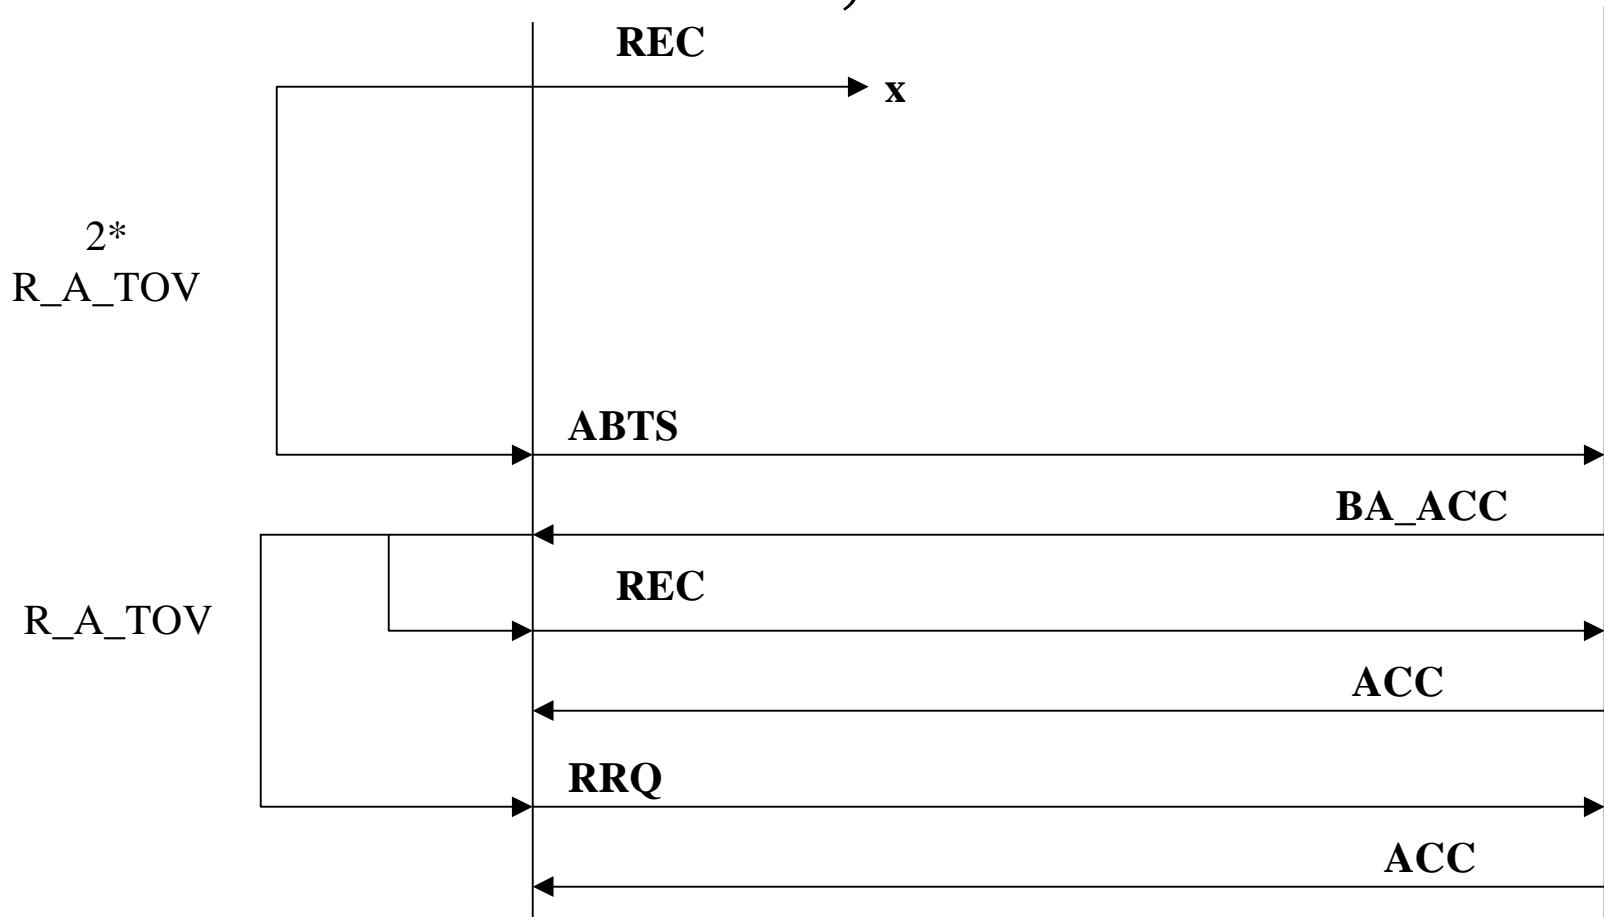


Error Recovery

Reuse of Exchange IDs can produce ambiguities. In the event of multiple errors, the current Exchange shall be aborted.

1. The Initiator shall not act on the ABTS until either the ACK to FCP_CMND is received or E_D_TOV expires.


D.13 C1 3, REC Lost



Since the REC was never received by the Target, the BA_ACC payload is SEQ_ID invalid, low SEQ_CNT =0, high SEQ_CNT= SEQ_CNT of ABTS=1. Recovery qualifiers are established on both sides. The second REC must be issued in a new Exchange.

Change E_D_TOV in the test to $2^*R_A_TOV$ to agree with the text in 12.6.2


D.13a C1 3, REC Response Lost



Since the Target has already responded with ACC then no context exists for the Exchange. The Target would view the ABTS as being issued on a new Exchange, establish a Recovery Qualifier and respond with BA_ACC (Payload is SEQ_ID invalid, low SEQ_CNT=high SEQ_CNT= SEQ_CNT of ABTS). Since REC does not change any state, it can be reissued unconditionally. The second REC must be issued in a new Exchange.

Change E_D_TOV in the test to $2 * R_A_TOV$ to agree with the text in 12.6.2

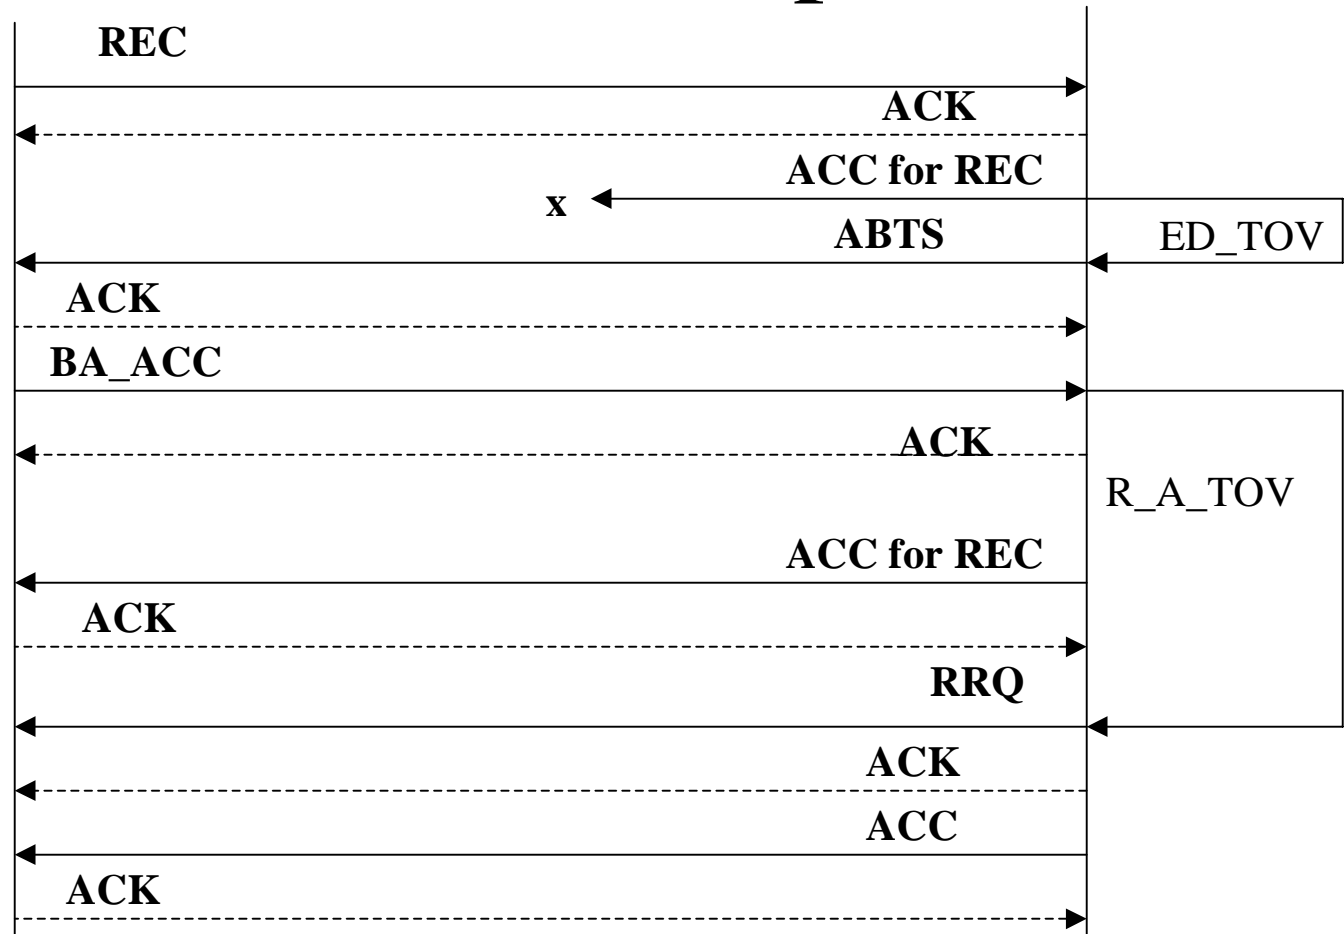
D.13b Class 2, REC Lost



The BA_ACC payload indicates that the REC was never received by the Target (Payload is SEQ_ID invalid, low SEQ_CNT=0, high SEQ_CNT=SEQ_CNT in ABTS frame) Recovery Qualifiers are established on each side.

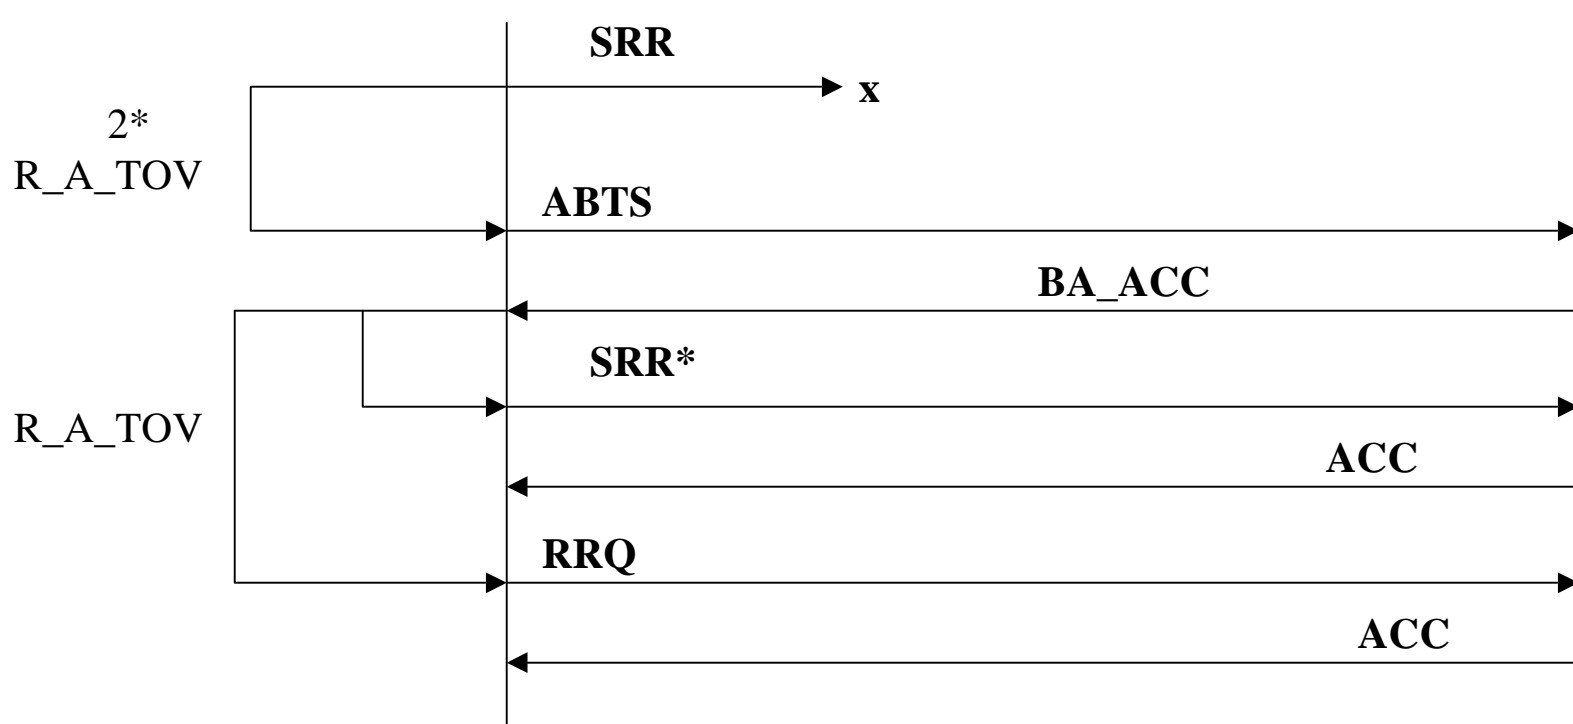
The second REC is issued using a new Exchange.

D.13c Class 2, REC Response Lost



Note: The Target determines that the ACC was never received by the Initiator; BA_ACC payload is SEQ_ID invalid, Low SEQ_CNT=0, High SEQ_CNT=SEQ_CNT of ABTS. Target reissues the ACC.

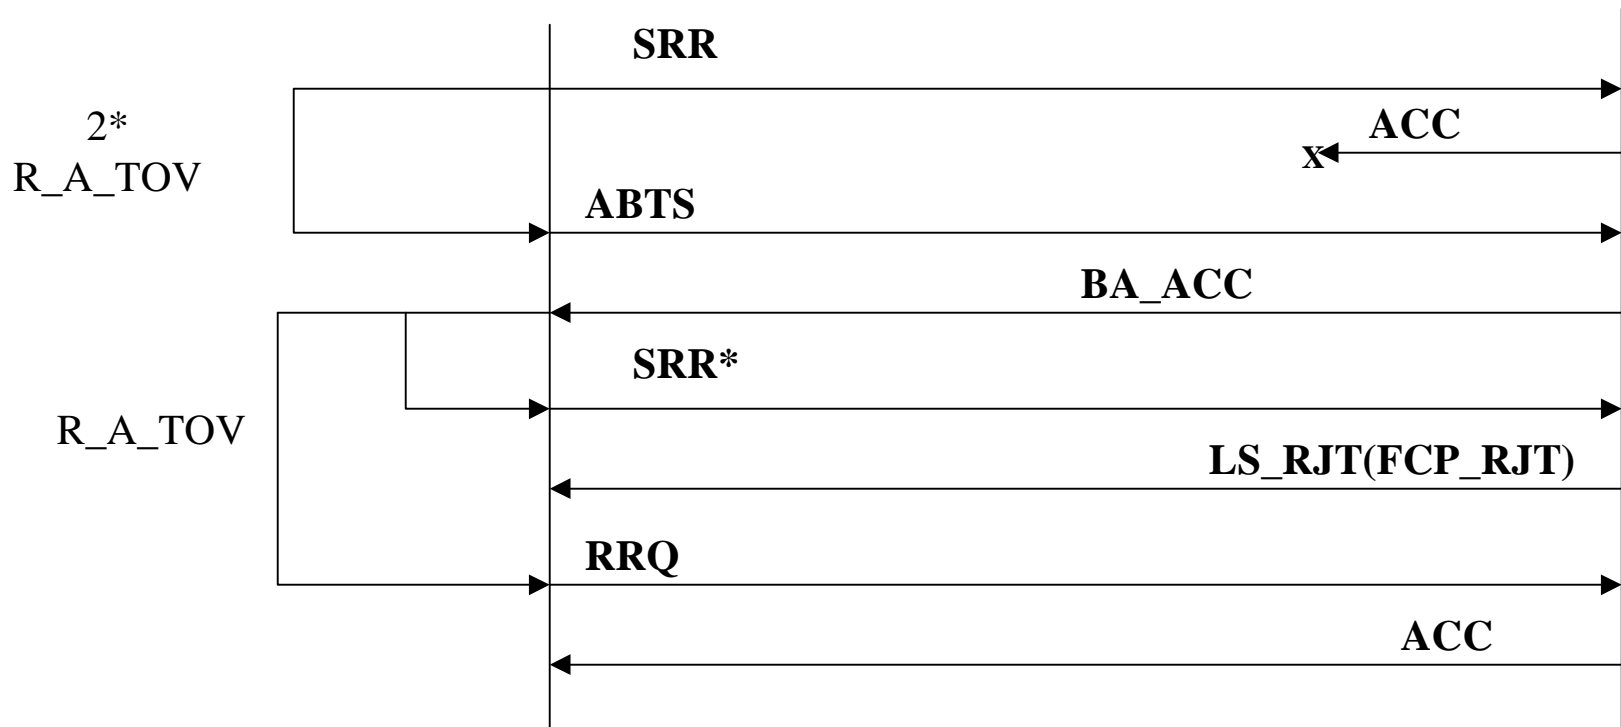
D.14a Class 3, SRR Lost



Note: BA_ACC Payload: SEQ_ID Validity = invalid, low SEQ_CNT=0, high SEQ_CNT =SEQ_CNT of ABTS frame. Recovery Qualifiers are established on both sides. SRR* is issued in a new Exchange. The Target restarts the original Exchange per the SRR* Payload.

Change E_D_TOV in the test to $2 * R_A_TOV$ to agree with the text in 12.6.3


D.14b Class 3, SRR Response Lost



Note: The Exchange for the SRR completed. Since to the Target this looks like an ABTS on a new Exchange, the BA_ACC Payload is SEQ_ID Validity = invalid, low SEQ_CNT=0, high SEQ_CNT = SEQ_CNT of ABTS. FCP_RJT is the response to the SRR* since the the original Exchange has been restarted by the Target per the Payload of the SRR. The original Exchange is in process, or it has completed and no context for it (OX_ID-RX_ID) remains.

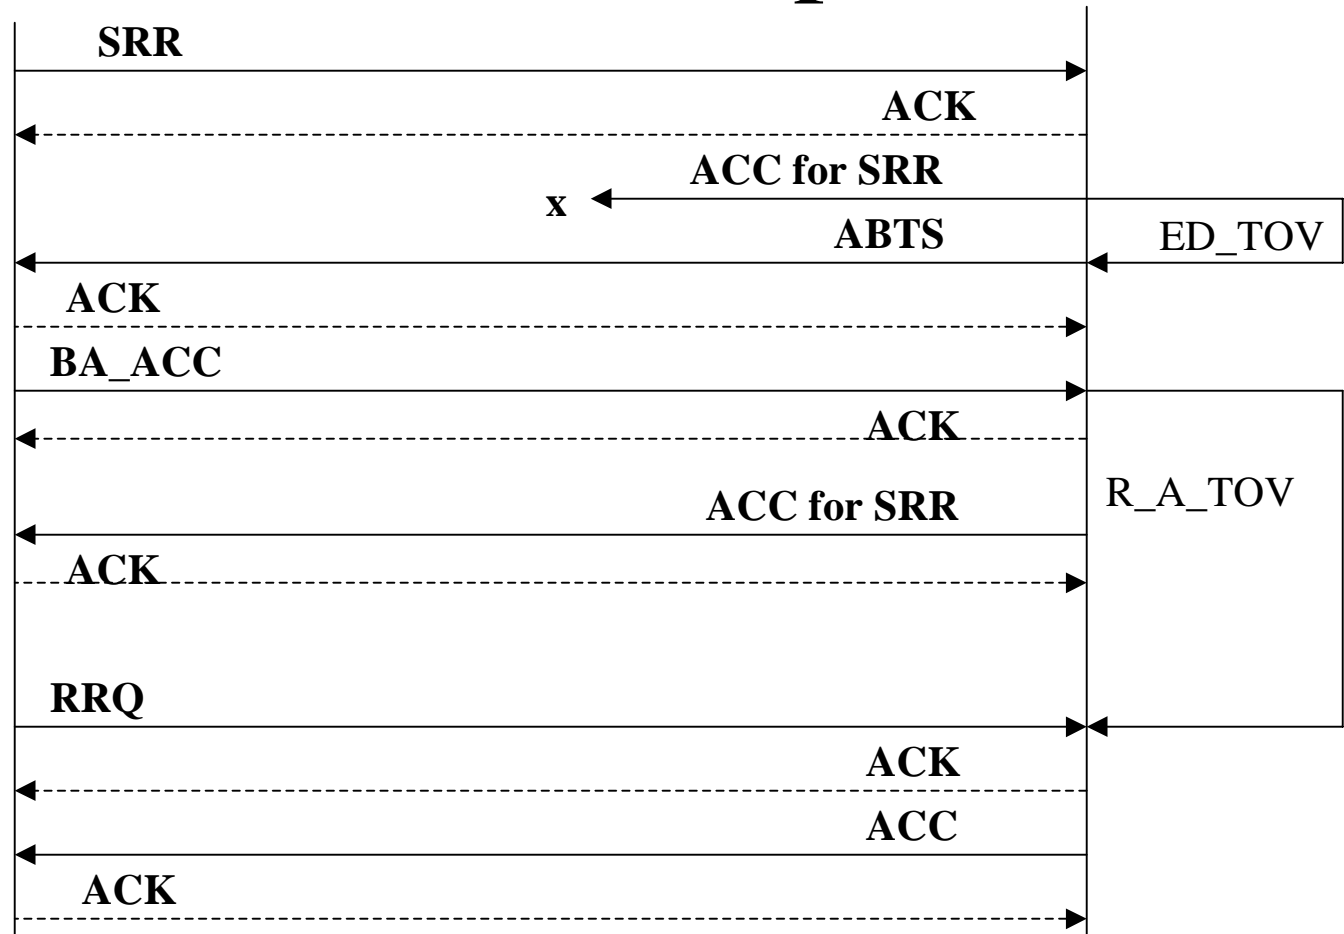
Change E_D_TOV in the test to $2^*R_A_TOV$ to agree with the text in 12.6.3

D.14c Class 2, SRR Lost



Since this is an ABTS on a new Exchange, Recovery Qualifiers must be established. BA-ACC indicates Invalid SEQ_ID, low SEQ_CNT= 0 and high SEQ_CNT = SEQ_CNT of the ABTS. The second SRR is issued using a new Exchange.

D.14d Class 2, SRR Response Lost



Note: The BA_ACC payload indicates SEQ_ID invalid, low SEQ_CNT=0 and high SEQ_CNT=SEQ_CNT of the ABTS, which indicates that the ACC for SRR was not received and will be discarded if it is received. Recovery Qualifiers are established on both sides. The ACC for SRR is issued with a new SEQ_ID and a SEQ_CNT one greater than used in the ABTS.