Proposal to handle unconfirmed FCP_RSP

Charles Binford, Symbios Logic

Background:

The fact that the delivery of the FCP_RSP packet is unconfirmed in the class 3 AL environment has been discussed over the past few months in committees and on the SCSI and Disk Attach reflectors. (Some also argue that even under class 2, the successful delivery of FCP_RSP cannot be known. This argument hinges on whether or not the ACK implies delivery to the ULP.) A major consideration in the discussion on the potential loss of the FCP_RSP IU (and thus any SCSI Sense data since FCP uses autosense) is whether or not a host actually cares. If the FCP_RSP IU is lost, the typical host will time-out the IO and reissue it. For block devices doing reads and writes this may be a sufficient error recovery scheme. The issue gets interesting, however, when the FCP_RSP IU which is lost contains autosense data WHICH IS NOT ASSOCIATED WITH THE COMMAND. This can occur when the target is attempting to report a Unit Attention or a deferred error.

The intent of this document is not to further the debate concerning the FCP_RSP IU, but rather to propose a solution. To that end the following describes an optional interlock mechanism which may be used by a target to ensure delivery of SCSI Sense data. The additional interlock may be invoked by the target on a per IO basis so as to not impede performance (it is assumed targets would not ask for confirmation for GOOD status). One problem encountered when adding a “handshake” is when to stop: what if the host’s FCP_RSP_ acknowledgment IU gets lost? This proposal addresses that and other potential error scenarios.

1. Two New IUs

- FCP_RSP_REQ_CONFIRM
- FCP_RSP_CONFIRM

1.1 IU Definition

1.1.1 FCP_RSP_REQ_CONFIRM

(target to init) An FCP_RSP_REQ_CONFIRM IU is a normal FCP_RSP IU with the following F_CTL bit changes:

- set Transfer Sequence Initiative
- do not set Last Sequence

1.1.2 FCP_RSP_CONFIRM

(init to target) An FCP_RSP_CONFIRM IU is defined as follows:

- R_CTL bits 31-28: 0000 FC-4 Device_Data
- R_CTL bits 27-24: 0011 Solicited Control
- Type code: 0000 1000 SCSI-FCP
- Payload: 4 bytes, value TBD

2. Interoperability

- Use determined by PRLI parameter.
- Only invoked by target if initiator supports.
- May be invoked by target on a per IO basis (e.g. only when NOT good status)

3. Usage Rules

3.1 Target use of FCP_RSP_REQ_CONFIRM

If the target wishes to request confirmation from the initiator of an FCP_RSP it shall send the FCP_RSP_REQ_CONFIRM IU instead of the normal FCP_RSP.

3.2 Initiator use of FCP_RSP_CONFIRM

When an initiator detects FCP_RSP_REQ_CONFIRM IU it shall send an FCP_RSP_CONFIRM IU.

3.3 Target cleanup of exchange and data

A target which sends an FCP_RSP_REQ_CONFIRM IU shall maintain any associated sense data to allow for a vendor unique number of retries until any of the following:

- an FCP_RSP_CONFIRM IU is received with a payload of (TBD) and FQXID
- an FCP_CMD is received with an OX_ID and S_ID matching that of the yet to be confirmed FCP_RSP_REQ_CONFIRM IU. (Note: this is the case where the FCP_RSP_CONFIRM was lost.)

3.4 Target Error Detection and Recovery

3.4.1 Target detection of lost FCP_RSP_REQ_CONFIRM IU

A target shall assume the FCP_RSP_REQ_CONFIRM IU was not received by the initiator if the FCP_RSP_CONFIRM IU is not received within a target specific time-out.

- the time-out shall be > R_A_TOV

3.4.2 Target retry of FCP_RSP_REQ_CONFIRM IU

The target may retry the FCP_RSP_REQ_CONFIRM IU using the following rules:

- maintain the original FCP_SNS_INFO and FCP_STATUS
- set the RSP_CODE to FCP_RSP_RETRY (value TBD)

3.4.3 Initiator receipt of a retried FCP_RSP_REQ_CONFIRM IU

If an initiator receives an FCP_RSP_REQ_CONFIRM IU with an RSP_CODE set to FCP_RSP_RETRY it shall take one of the following actions:

- if the OX_ID is not currently active, send confirmation (previous confirmation was lost)
- if the FCP_CMD which is active on this OX_ID was sent at time t and (current_time - t) >= R_A_TOV, then treat as normal FCP_RSP_REQ_CONFIRM and send confirmation (previous FCP_RSP_REQ_CONFIRM was lost)
- if the FCP_CMD which is active on this OX_ID was sent at time t and (current_time - t) < R_A_TOV, then ignore (previous confirmation was lost and target attempted

retry at the same time initiator reused OX_ID. The target will see the new FCP_CMD with the given OX_ID and cleanup the previous IO.)

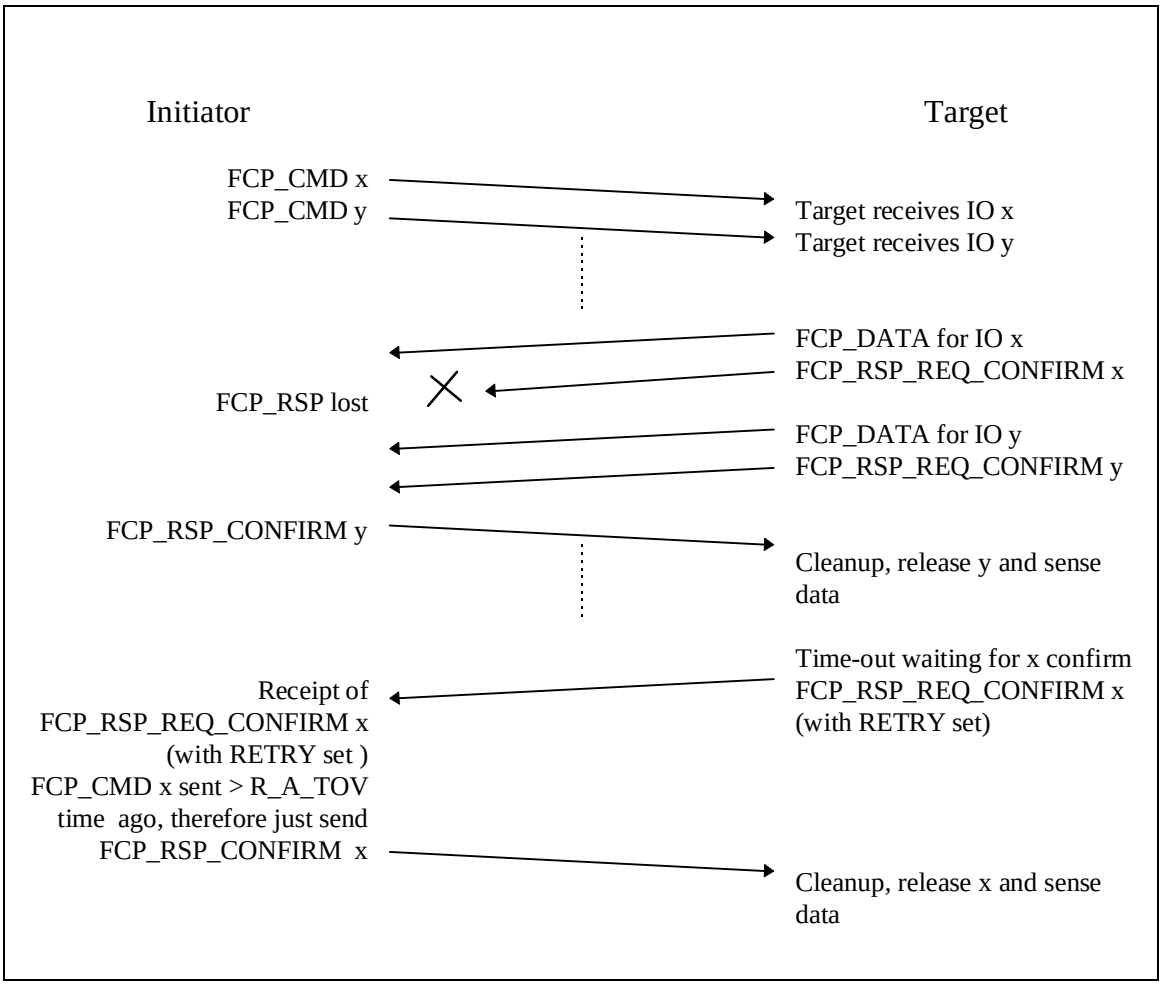

Figure 1: Example Flow

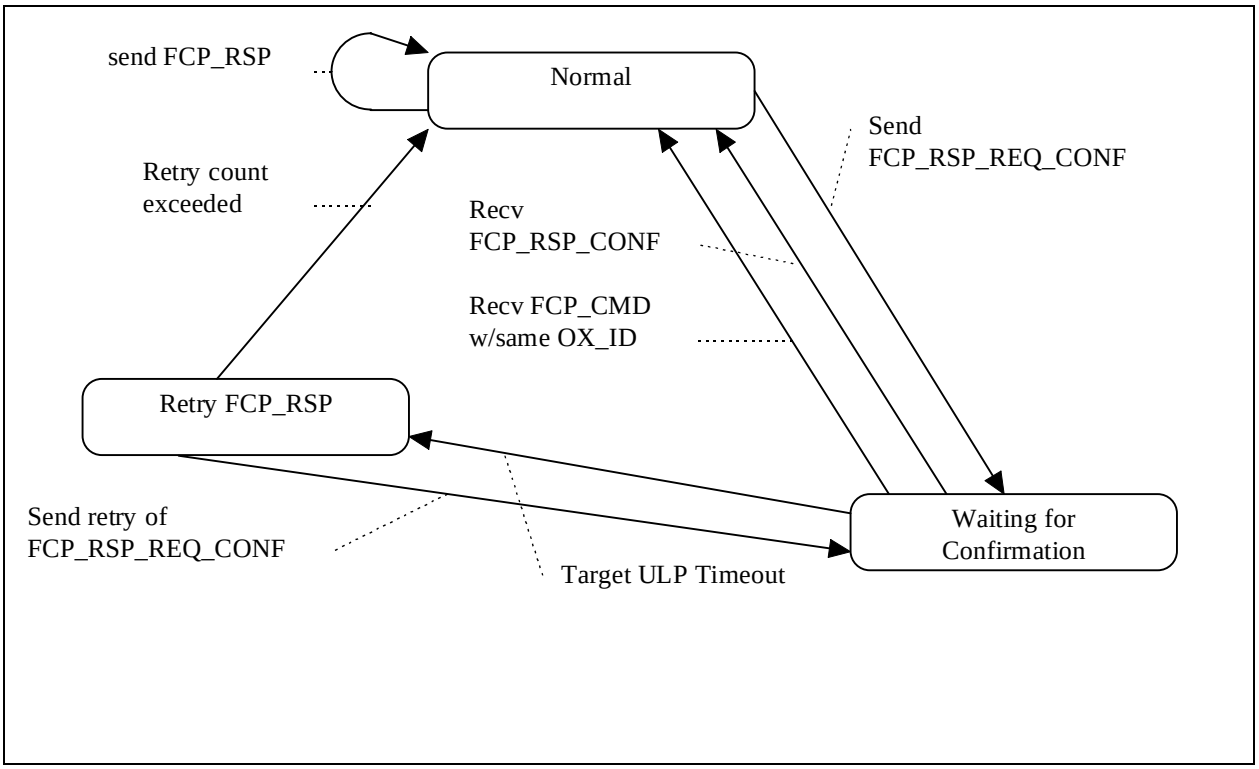

Figure 2: Target States

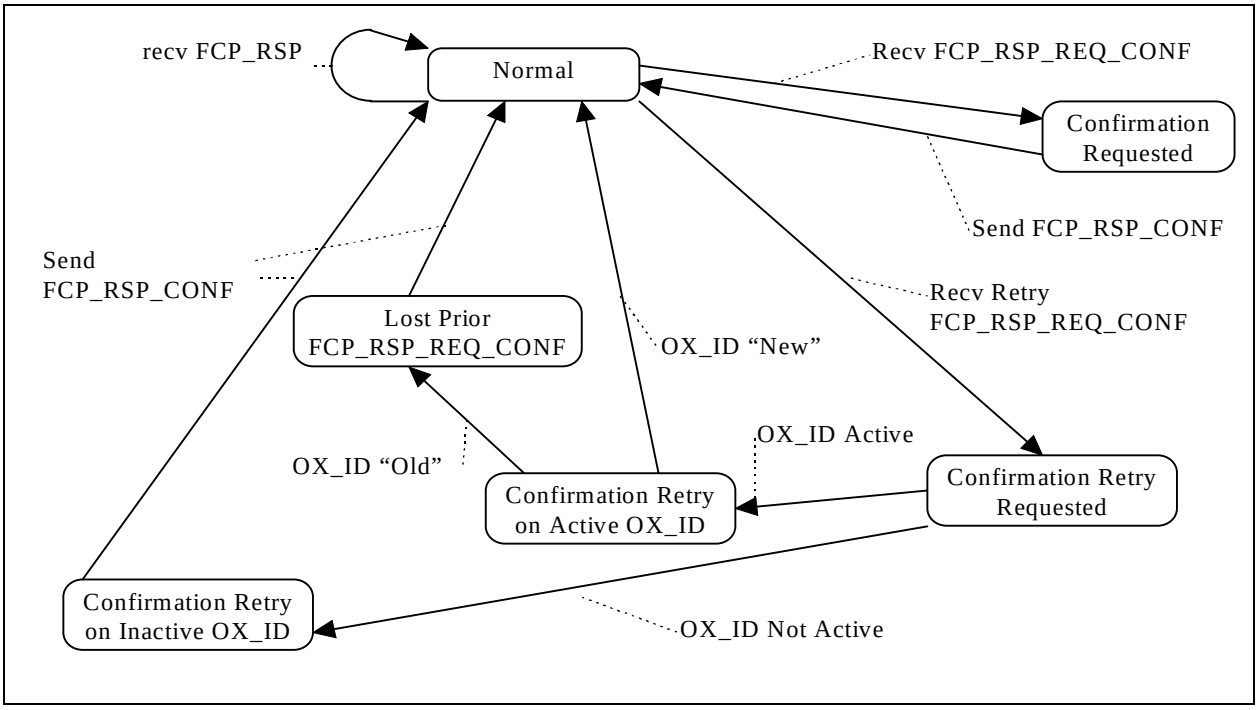

Figure 3: Initiator States