

FCP-2 Problem Recovering the Proper Command

Dave Baldwin
Emulex Corporation
June 30, 2000

I propose the following solution to this problem:

For this proposal, I am defining "FCP-2 target" as a FC SCSI/FCP target device in a system where both the FCP-2 target and the Initiator have set the Retry bit in their PRLI/ACC (Process Login) frames.

This proposal attempts to minimize the severity of the changes, yet keep the performance at an acceptable level. This proposal uses the Parameter field to define an FCP2_Handle that is used to uniquely identify the proper command. The initiator assigns a unique FCP2_Handle to every new command, and transmits this handle in the Parameter field of the FC header for the command. The F_CTL Relative Offset Present bit shall not be set, as this may confuse some hardware. The FCP-2 target shall save the FCP2_Handle for use in error recovery.

The FCP2_Handle is used in the REC and SRR commands in order to recover the proper command. The Initiator and the FCP-2 target shall transmit the FCP2_Handle in each REC and SRR command frame. The FCP2_Handle is used to uniquely identify the command being recovered. The S_ID, OX_ID, and RX_ID fields shall still be verified before returning the ACC, LS_RJT, FCP_ACC, or FCP_RJT frame.

In order for the Initiator to properly control the FCP2_Handle, the FCP-2 target is required to age completed commands, and to release all resources associated with completed commands within a dependable timeout period. The timeout should be RR_TOV, but a maximum time needs to be specified. For this proposal, I have picked 30 seconds for the maximum RR_TOV time period. Even though the probability is very low, if old information is left around forever, then you could come to a point where you reuse the FCP2_Handle and have a command identification problem.

The Initiator can control the FCP2_Handle in many ways. An appropriate timestamp could be used. In today's devices you may find it sufficient to just have an incrementing counter put in to each new exchange. As we move to faster speeds, like 10Gb, then the Initiator may want to have an incrementing counter for every FCP-2 target it is communicating with. The designer would ensure FCP2_Handle uniqueness based on the above timeout of completed commands.

There are recovery problems created if both the Initiator and the FCP-2 target don't implement this new behavior at the same time. The receiving device wouldn't be able to identify the command when "garbage" is sent in the Parameter field. We could choose to say always update all devices at the same time and leave the coordination to system administrators. Or, we could choose to make it a negotiated behavior to ensure error recovery doesn't degrade. The negotiation should be done with a new "Identify" bit in the PRLI command and in the ACC. The Identify bit should go in Word 3, bit 9, to put it right next to the Retry bit. The FCP-2 device must not recover based on the FCP2_Handle unless both FCP-2 devices have set the Identify bit. In fact, some implementations may choose to not talk to an FCP-2 device that doesn't set the Identify bit, or send up red flares to let someone know there is a problem.

ABTS Command Identification

We need to guard against having several outstanding exchanges with the same OX_ID from the target's point of view (20 tape drives, individual LUNs within one target, whose last exchange executed just happen to have the same OX_ID within the timeout period). Otherwise, we have command identification issues with AB TS because it is not LUN specific. This will become a problem when Class2 behavior is implemented.

I think a good solution is for the FCP-2 target to release all resources associated with the old command with OX_ID = n (which the target believes has been completed), when it gets a new OX_ID = n frame in with a new command (R_CTL = 6). The reuse of the OX_ID by the initiator is a confirmation that the old command has been completed. Since the target and initiator both think the old exchange is complete, this should be sufficient confirmation to get rid of the old information in the target.

These changes are a simple yet very effective way to ensure proper command identification in FCP-2 recovery.