



Overview of Store-and-Forward Expander Device Operation

Bob Sheffield
Storage Group
11 September 2006

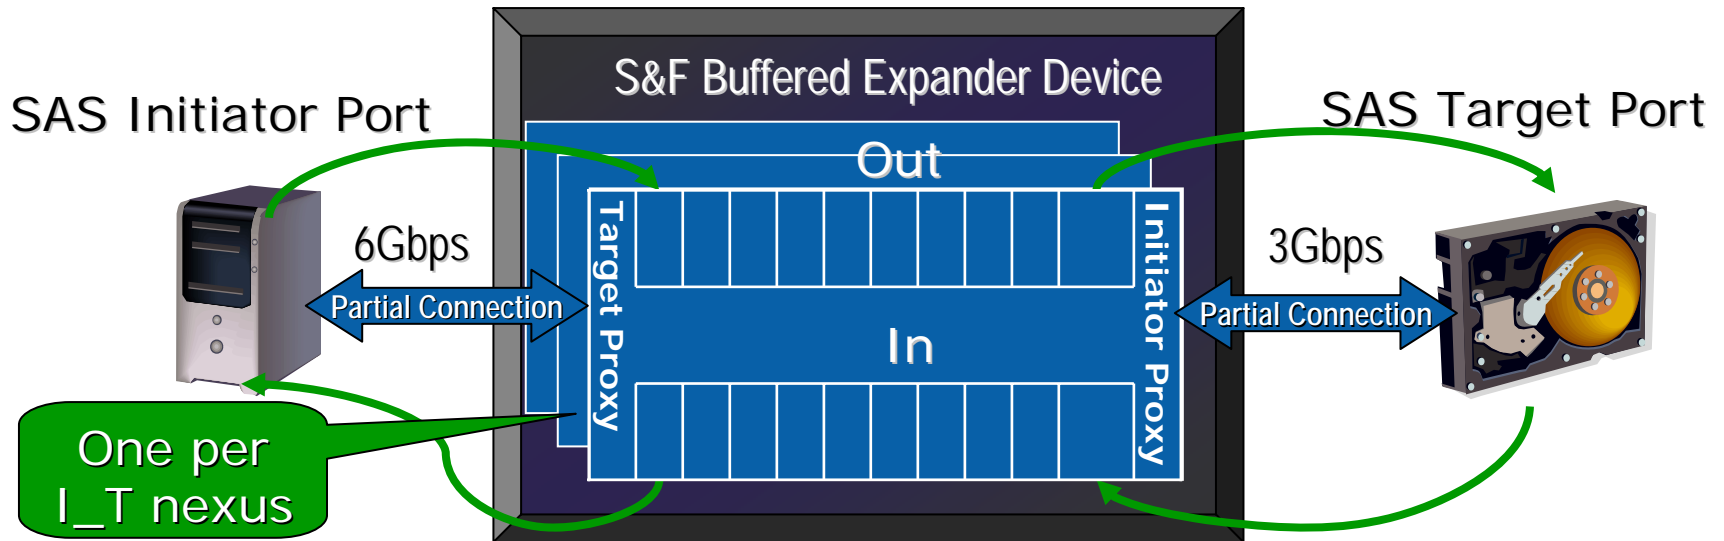
Overview

- Objectives
- General concept
- Basic elements
- Connections
- SSP flow model
- STP flow model
- SMP extensions
- SSP error handling
- STP error handling
- Summary

S&F Buffering Objectives

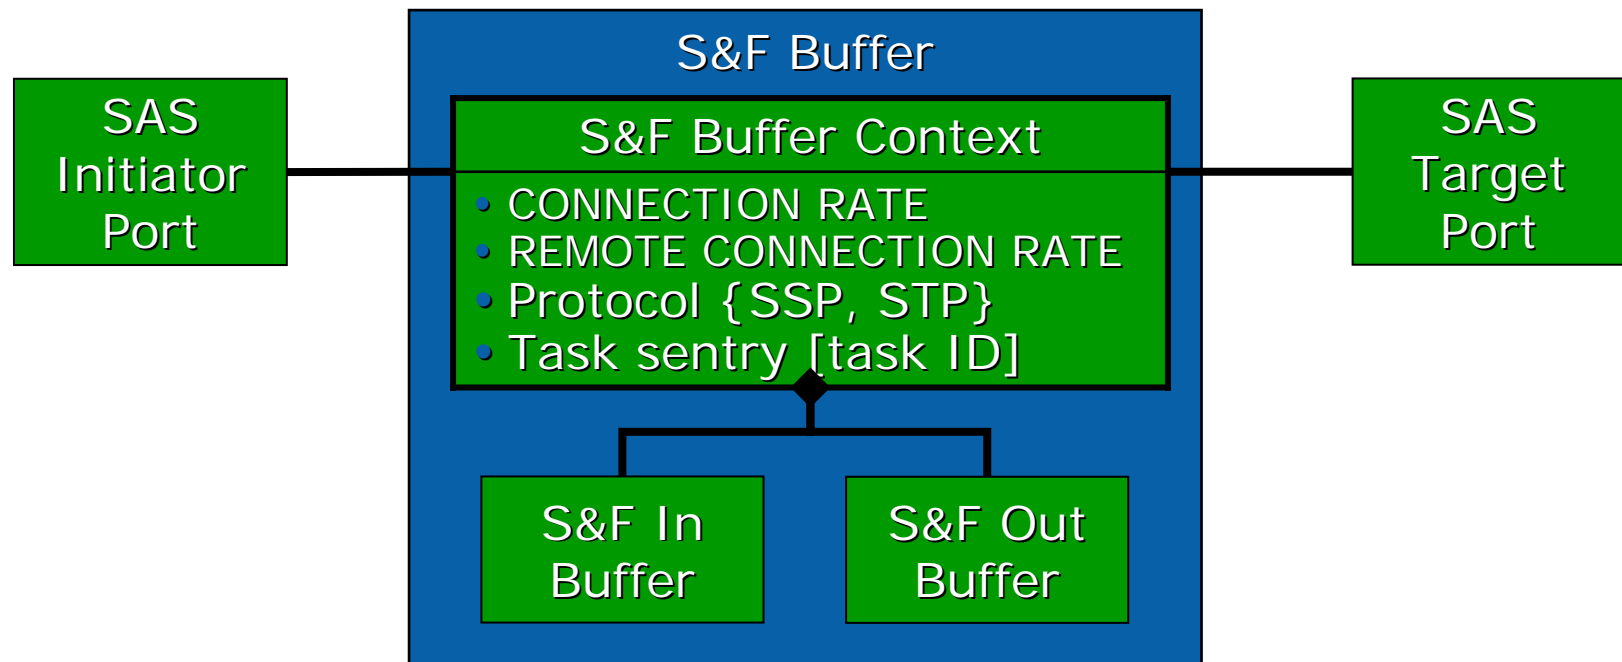
- Effective bandwidth aggregation
- Works for both SSP and STP
- No effect on SAS target or SATA device
- Minimum impact (SW/FW only) on initiator
- Add complexity only where it's needed
- Effective error detection and handling
- Backwards compatibility
- Accommodate higher speeds w/o change

S&F Basic Concept



- S&F buffer provides proxy initiator & target ports
- Separate buffers hold frames for transfers in and out
- Partial connections established to transfer to/from buffers
- Only one S&F buffer activated along I↔T pathway
- Out buffer closes connection to initiator port when full, accepts OPEN when empty
- In buffer closes connection to initiator port when empty,

Basic Elements



- OAF from initiator port w/ REMOTE CONNECTION RATE \neq 0
 - Establish S&F Buffer Context
 - Return AIP (SFT_TAKEN) to confirm S&F service
 - Forward OAF towards target port RCR \neq 0
 - If AIP (SFT_TAKEN) received, disable S&F
 - Attach S&F IN and Out buffers when connection is established
 - CLOSE on one end leaves a “partial connection” on the other end

Connections

- I→T OPEN: S&F service is requested w/ REMOTE CONNECTION RATE $\neq 0$
- T→I OPEN: 1st S&F expander enables S&F, then forwards OPEN w/ REMOTE CONNECTION RATE = 1
- If buffer context exists, S&F buffer accepts OPEN w/o forwarding
- Buffer confirms S&F service by AIP (S&F TAKEN)
- Buffer signals full condition to SSP Initiator port by sending DONE (CREDIT BLOCKED)
- Buffer uses CLOSE (NORMAL) to close STP Initiator connection when full
- S&F buffer closes connections as usual when no frames are available to transmit

SSP Flow Model

- S&F buffer respects interlocked rules for transmitted frames
 - i.e., requirement to track TAG field

For transfer out:

- Buffer accepts frames from initiator port
- DONE (CREDIT BLOCKED) returned to initiator port when buffer is full
- Initiator retries later (allowing time for buffer to drain)
- Buffer returns OPEN REJECT (RETRY) until buffer is (nearly) empty

For transfer in:

- Buffer disconnects [DONE (NORMAL)] from initiator port and receives frames from target port
- Reconnects to initiator port when buffer is full
- Disconnects [DONE (NORMAL)] from initiator port when buffer is empty

STP Flow Model

S&F buffer only buffers SATA frames w/ DATA FISes

- for non-DATA FIS, R_OK held off until FIS reaches destination

For transfer out:

- STP S&F buffer receives X_RDY and returns R_RDY
 - If a DATA FIS is received, buffer returns R_OK, SYNC,...
 - Continue until buffer fills or a non-DATA FIS is received
- CLOSE (NORMAL) returned to STP initiator port when buffer is full (i.e., cannot accept another full DATA FIS).
- STP initiator port retries (allowing time for buffer to drain)
- Buffer returns OPEN REJECT (RETRY) until buffer is (nearly) empty

For transfer in:

- Buffer disconnects [DONE (NORMAL)] from initiator port and accumulates DATA FISes from the STP target port
- Reconnects to STP initiator port when buffer is full or non DATA FIS is received and transmits FISes to STP initiator port
- Disconnects [DONE (NORMAL)] from initiator port when buffer is empty

SMP Extensions

- S&F CAPABLE bit added to REPORT GENERAL
- S&F ENABLED bit added to DISCOVER
- ENABLE/DISABLE field for CONFIGURE GENERAL
- ENABLE/DISABLE bit for PHY CONTROL
- New SMP S&F BUFFER CONTROL command
 - Query buffer (SSP or STP)
 - Buffers exist? Frames Present?
 - Frames present or sentry active for a specified tag?
 - Partial connection status (target and initiator)
 - Delete buffer (SSP or STP)
 - Abort Task (discards frames for a specified tag)
 - Set / Get parameter
 - STP Max Frame Size (default to 8K)

SSP Error Handling

- S&F Buffer transmits frame to destination port and receives NAK or ACK/NAK timeout:
 - S&F buffer already returned ACK to source port
 - Subsequent DATA frame gets offset error
 - SSP target will report the error
 - Worst case: ULP timeout
 - SSP initiator performs triage
 - Query Task to target device
 - SMP S&F BUFFER CONTROL (Query Buffer)
 - If context lost, send aborts to target and S&F Buffer to clear frames for the affected tag
 - Retry command at upper-layer

STP Error Handling

- For non-DATA FISes, handled normally
- R_ERR returned on transmitted DATA FIS:
 - S&F Buffer discards R_ERR and continues forwarding frames
 - Device detects and reports the error
 - STP initiator port retries
- Checking PIO SETUP parameters
 - S&F Buffer keeps parameters in buffer context
 - Returns an error FIS if next FIS doesn't match
- SYNC escape
 - When an S&F Buffer detects a SYNC, it will reply with SYNC and discard the affected frame.

Summary

- S&F Buffering applies complexity as needed
- Buffers (2/I_T) has proxy Initiator/Target ports
- Context and in/out buffers work with partial connections
- First connection goes end-to-end, then partials
- Initiator connects/disconnects to keep buffers busy
- STP only buffers DATA FISes (simplifies errors)
- SMP extensions provide control, visibility, recovery
- SSP error handling maps to ULP timeout when not detected by target device
- STP error handling usually detected by device. STP initiator recovers the usual way (RESET/RETRY).

