

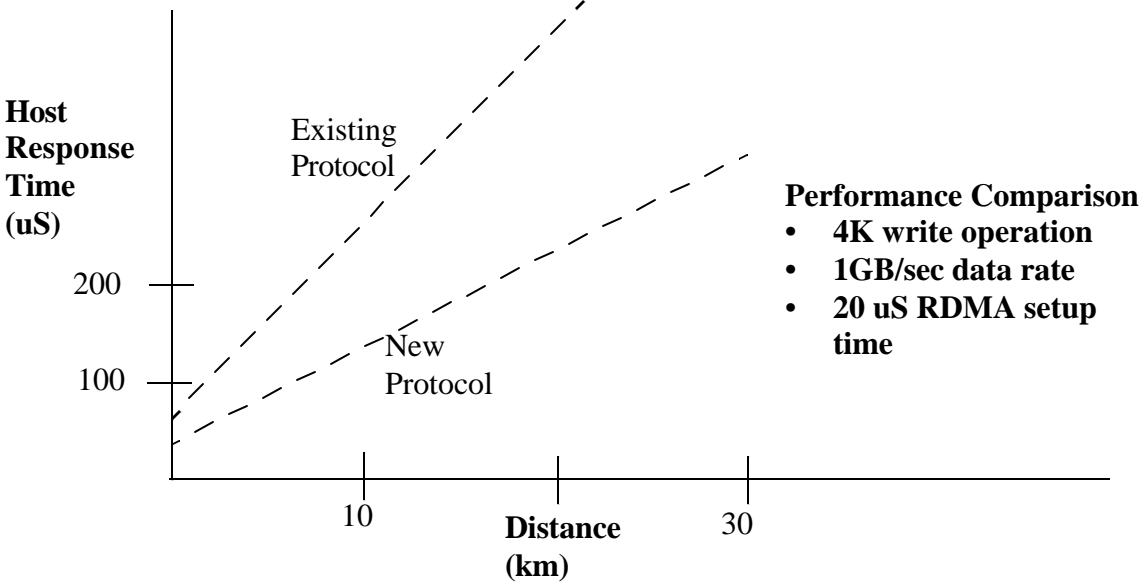
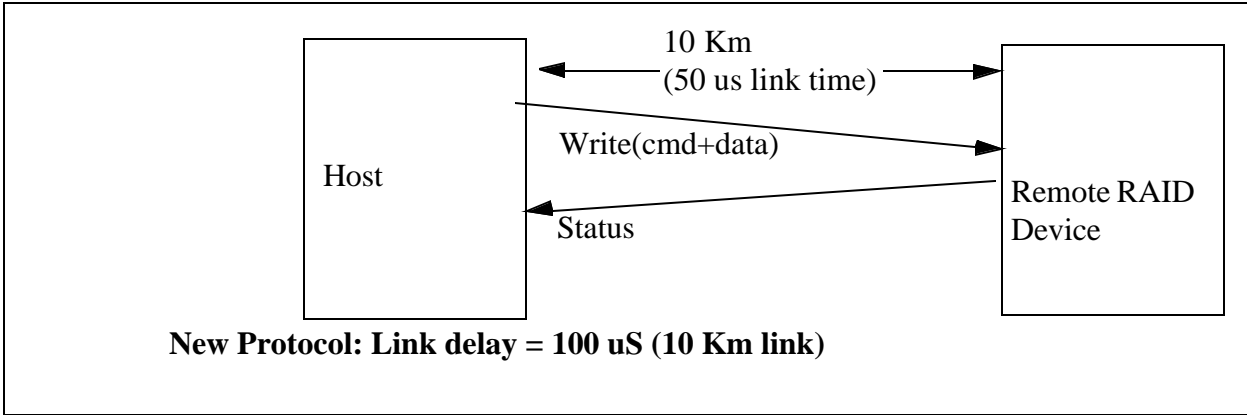
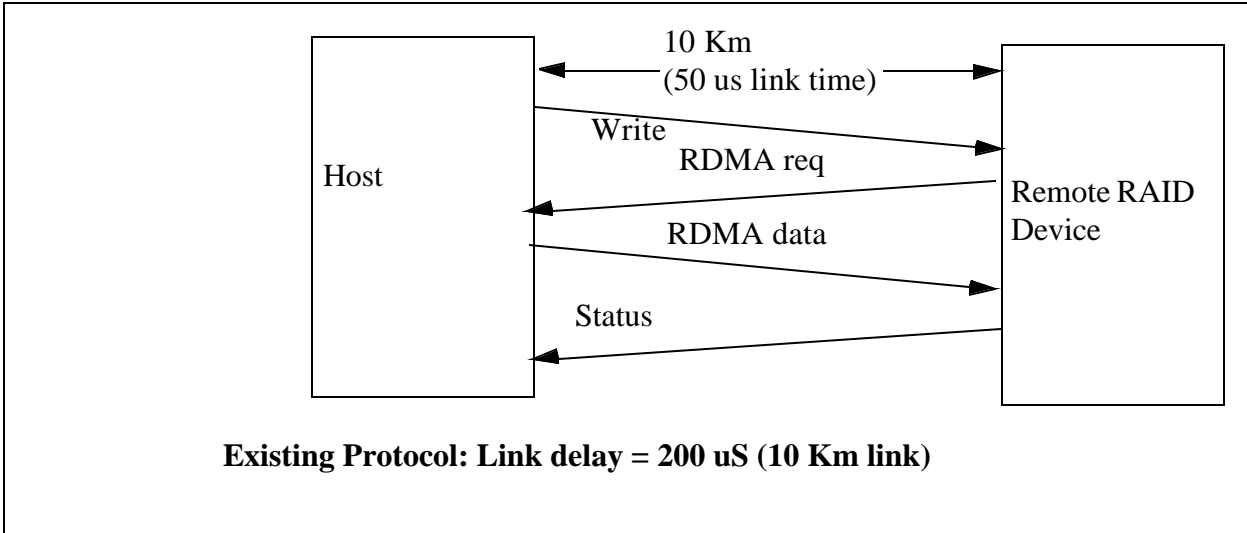
# **Sending Write Data with SRP Commands**

Giles Frazier

## **Problem**

- Short write data operations are costly in SRP
  - 3 round-trips on link before first write data arrives at target
  - RDMA setup required for relatively small data transfers
  - Significant performance degradation for many applications
    - \* Real-time remote mirroring
    - \* Backups requiring multiple updates of small data segments
    - \* Remote transaction processing
    - \* Batch data backup operations
  
- Other protocols have eliminated these delays
  - FCP
  - FC-SB-2
  - SDP
  - iSCSI
  
- SRP needs this capability to be competitive

### Performance Comparisons:



## **Solution**

- Define a new Command-data IU (or modify existing Command IU):
  - Payload includes a command + immediate data
  - Amount of data limited by maximum initiator to target IU length
  
- Discovery of support:
  - New Login bit: “immediate data requested/supported”
  - or: Non-zero “first burst” parameter in Disconnect/Reconnect mode page

## Alternative I: Modify Existing Command IU

Byte	Bit	7	6	5	4	3	2	1	0
0		TYPE(02)							
1-4		Reserved							
5		DATA-OUT BUFFER DESCRIPTOR FORMAT				DATA-IN BUFFER DESCRIPTOR FORMAT			
6		DATA-OUT BUFFER DESCRIPTOR COUNT							
7		DATA-IN BUFFER DESCRIPTOR COUNT							
8-15		TAG							
16-19		Reserved							
20-27		LOGICAL UNIT NUMBER							
28		Reserved							
29		Reserved				TASK ATTRIBUTE			
30		Reserved							
31		ADDITIONAL CDB LENGTH = n						Reserved	
32 - 47		CDB							
48 - 47+4*n		ADDITIONAL CDB							
48+4*n 47+4*n+do		DATA-OUT BUFFER DESCRIPTOR							
48+4*n+do 47+4*n+do+di		DATA-IN BUFFER DESCRIPTOR							
48+4*n+do+di ... (<Max IU Size)		IMMEDIATE DATA OUT							

## Alternative II: New Command-Data IU

Byte	Bit	7	6	5	4	3	2	1	0
0		TYPE(0x)							
1-7		Reserved							
8-15		TAG							
16-19		Reserved							
20-27		LOGICAL UNIT NUMBER							
28-31		Reserved							
32 - 47		CDB (max 16 bytes)							
48 - n		Command data (limited by max IU size)							

**Potential advantage: Start of data is always at a fixed location in IU.**

**- May enable the target to more easily route the command data to appropriate location**

# Large Write Operations

- Primary usage of this new capability is for short writes
  - ..but could be used for large writes also

## Three alternatives for long writes:

- Do not send any immediate data, use RDMA exclusively
- Send small amount of immediate data, use RDMA subsequently
  - Probably not desirable--small performance improvement
- Adjust login parameters to allow data to be sent with command:
  - Large MAXIMUM INITIATOR TO TARGET IU SIZE
  - Small REQUEST LIMIT

**The third option eliminates start-up delays for write operations without wasting excessive buffer space at the target**