

Fabric Technology Can Enable and Ease Migration to 8Gig FC

Ali Ghiasi

June 7, 2004

FCIA Meeting

3151 Zanker Rd

San Jose, CA 95134

(408)922-7423

Connecting
everything™

aghiasi@broadcom.com



A. Ghiasi

BROADCOM CORPORATION

FC Meeting Chicago

Overview

- 8Gig MRD
- FC / Disk I/O
- Architectural implications Loop and Fabric
- Is an 8Gig homogeneous system path of least resistance
- Performance improvement with fabric

State of 8Gig MRD

□ IBM FCIA presentation Apr. 04

- ⇒ Components availability late 05 and product ship late 06.
- ⇒ Components required for IBM 8Gig deployment.
 - Backplane, optics, HBA, fabric, and copper cable.
- ⇒ There is no mention of the need for 8Gig disk drive in this time frame.
- ⇒ T11.2 initially set aggressive schedule based on IBM to develop 8Gig.

□ During recent FCIA call HGST said “8Gig drive needs 0.9 um CMOS to keep it cost and power competitive and the realistic schedule for it is 07/08“.

- ⇒ 4Gig drives are already late compare to other 4Gig FC components.
- ⇒ Based on HGST input T11.2 has pushed 8Gig development to 07.

□ The logical path forward is to decouple speed of disk drive from FC interface speed.


Connecting
everything™



I/O Throughput Improvement

□ To saturate an FC controller typically you need 10 disk drives.

⇒ Why burden the system operating disk I/O at 8Gig when real throughput is only 1/10.



Architectural Implications

❑ FC Arbitrated Loop enabled FC disk array (JBODs).

⇒ Arbitrated loop has significant performance weakness compare to fabric but allowed very simple hardware implementation (PBC).

❑ Due to lack of diagnostics and retiming most 4Gig disk array implementation are now based on intelligent devices (Loop Switch or intelligent PBC).

⇒ Silicon cost is the function of area and power dissipation.



- Power dissipation in a large port count device is driven primarily by the number of SerDes, speed, and memory.

⇒ At the higher speed the difference between a fabric switch and a loop switch diminishes.

⇒ Fabrics will deliver higher performance and lower cost if you operate the disk ports at lower rate, i.e. Controller operating at 8Gig and disk operating at 4Gig.

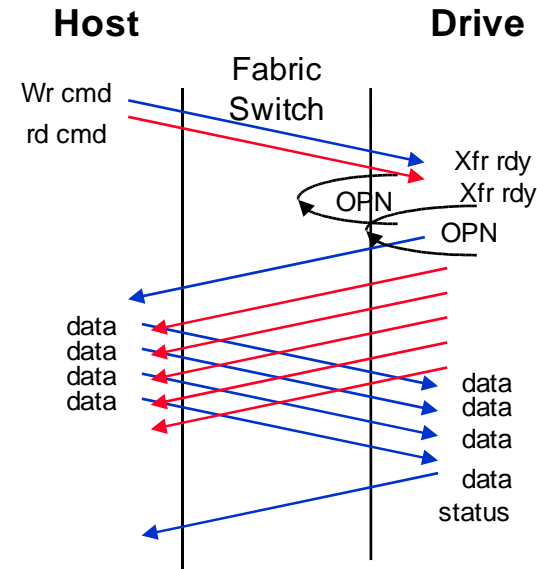

Loop Command Operation

- Advantage of Loop are simplicity of connecting device software.
- Disadvantage of Loop are simplex operation, single device can be opened by a given host, and all devices must operate at the same speed.




Fabric Command Operation

- Advantage of fabric are performance and scalability
- Disadvantage has been cost.



Improved Performance with Backend Fabric

MBps Comparison (2G HBA, 27 1G disk drives)



Connecting everything™



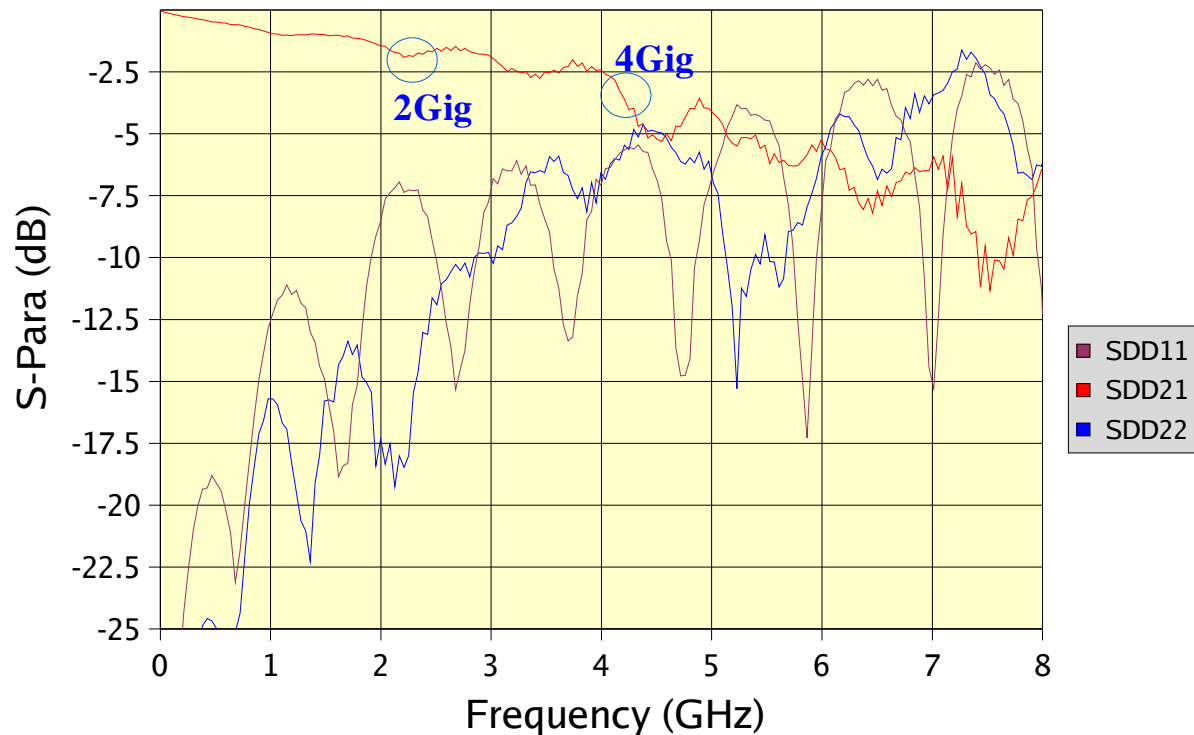
Should Disk Drive Lead the 8Gig

□ FC-PI2 had difficult time converging on Beta specification even 4Gig.

⇒ Operating disk interface at 8Gig is possible but more challenging.

⇒ We should not delay 8Gig FC due to lack of 8Gig drives.

SCA2 CH2



Summary

- ❑ **FCIA should not hold back roll-out of 8Gig FC for lack of 8Gig drive.**

- ⇒ Recommend FCIA to set the MRD date based on IBM request date of 05/06 if feasible.

- ⇒ Initial deployment of 8Gig can very well be based on 4Gig or 2Gig drives.

- ❑ **Fabric devices in backend decouple the controller speed from disk speed, while improving the performance, ease migration, and simplify the backplane.**

- ❑ **Fabric will be integral part of next generation arrays and we can't assume Fabric will reside just in the network anymore.**

- ❑ **Recommends FCIA to move forward with 8Gig Fabric vote.**

Connecting
everything™

