

07-416r1 T10 USB Queuing Transport (UQT) study group introduction

by Rob Elliott, HP
25 September 2007



USB and other external desktop interfaces

- USB dominates the market - 6.2 billion USB ports shipped as of 2007 (2.1 in 2006)
 - Pat Gelsinger, Intel, at Intel Developer forum 19 Sep 2007

External interface	Line rate	Throughput
USB Low Speed	1.5 MBit/sec	192 KB/sec
USB Full Speed	12 MBit/sec	1.5 MB/sec
USB 2.0 Hi-Speed	480 MBit/sec	60 MB/sec
1394	100 MBit/sec 200 MBit/sec 400 MBit/sec	12.5 MB/sec 25 MB/sec 50 MB/sec
1394b	800 MBit/sec	80 MB/sec
eSATA	1.5 Gbit/sec 3 Gbit/sec	150 MB/sec 300 MB/sec



USB storage performance

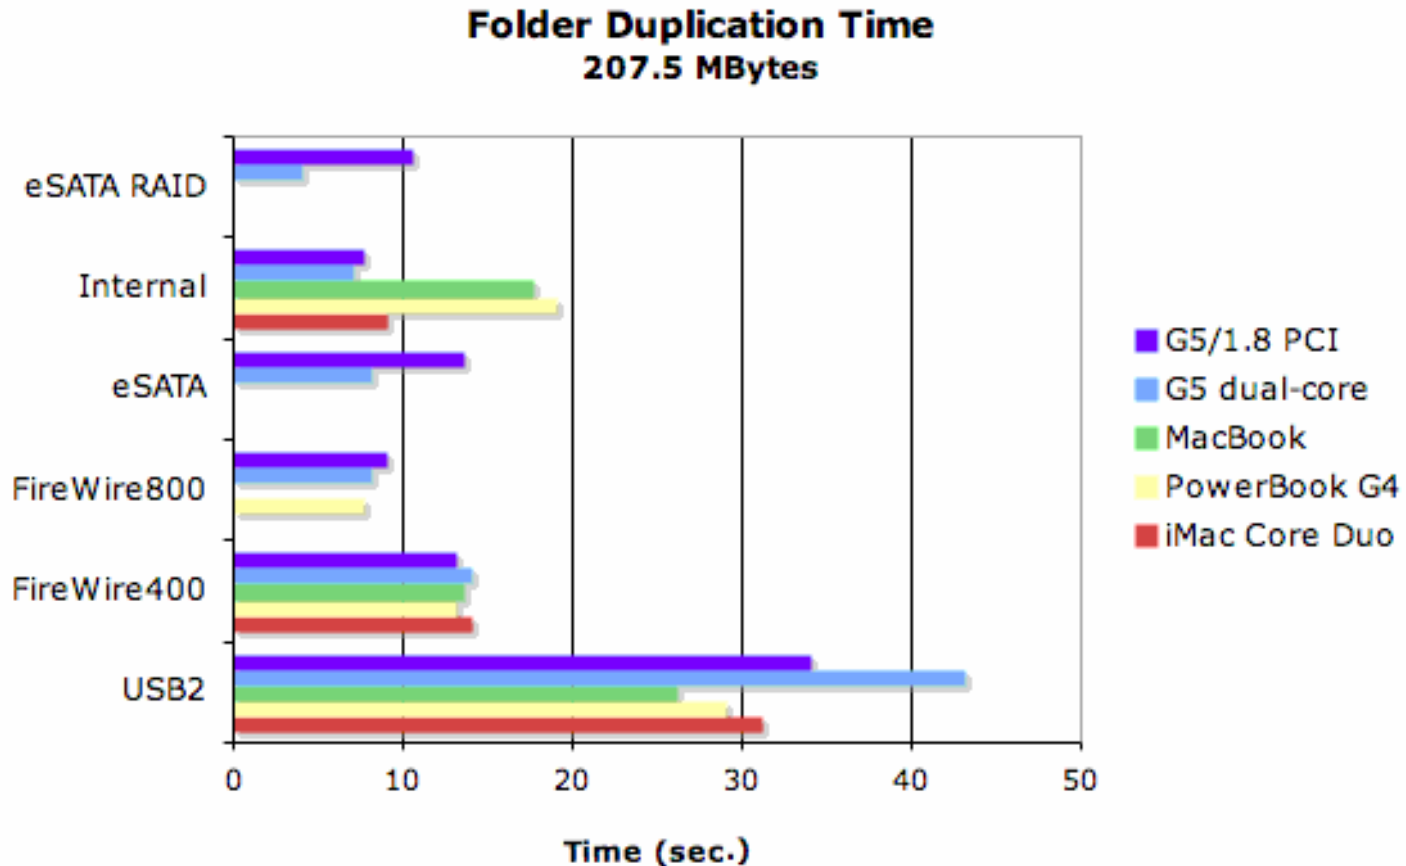
- USB Mass Storage Class Bulk-Only Transport is a SCSI transport protocol (not ATA)
- USB 2.0 mass storage performance is becoming an issue
 - Theoretical limit around 48 MBps (bus is half duplex and host-pollled)
 - Disk drives max out around 30-35 MBps with standard Microsoft Windows drivers
 - Protocol may not scale well to USB 3.0

Disk and flash drive bandwidth demands

- Desktop-class disk drive bandwidth faster than USB
 - WD Caviar: 54.4 to 97 MBps
 - WD Raptor: 55.4 to 88.3 MBps
 - Hitachi Deskstar: 46.4 to 86.9 MBps
 - Samsung SpinPoint: 41.5 to 83 MBps
 - Seagate Barracuda: 44.3 to 78.5 MBps
 - (Minimum to maximum read bandwidth per benchmarks on <http://www.storagereview.com>)
- Flash (solid state) drives still have headroom
 - Up to 10 MBps reads, 3 MBps writes in 2004 (<http://arstechnica.com/reviews/hardware/flash.ars/6>)
 - Up to 29.4 MBps reads, 19.23 MBps writes in 2006 (http://www.xbitlabs.com/articles/memory/display/10usbflash-roundup_9.html)

USB performance vs. other interfaces

- June 2007 comparison on <http://www.macintouch.com/specialreports/perfpack02>



Performance limitations of USB MSC Bulk-Only

- Some protocol limitations impair performance
 - Polled interface
 - No XFER_RDY like notification that drive is ready for write data
 - Possibility: advise host how long to wait before polling
 - cache lookup -> short wait
 - cache miss, seeking -> long wait
 - No queuing
 - eSATA now supports NCQ (native command queuing)
 - Data transferred in-order
 - When adding queuing, need to let commands finish out of order
 - Out-of-order data within a command not as important

Other limitations of USB MSC Bulk-Only

- Only 16 LUNs
 - SCSI architecture requires 16,384 LUNs (2-byte LUN field)
- No bidirectional commands
- 16 byte maximum CDB size
 - Object Storage Device (OSD) command set requires bidirectional long CDBs
- No SCSI status values
 - just Passed, Failed, and Phase Error
 - no BUSY, TASK SET FULL, etc. and no Retry Delay field
- No SCSI autosense data
 - REQUEST SENSE command required – CA concept obsolete after SAM-2
- No task attributes (SIMPLE, ORDERED, HEAD OF QUEUE)
- No task priority field
 - Serial ATA 2.6 includes a Priority feature
- No task management functions (ABORT TASK, LOGICAL UNIT RESET, etc.)
- Redundant transfer length field in command block wrapper
 - both good and bad

Proposal

- Start a new T10 project for a new USB Queued SCSI transport protocol
 - Support queuing and be SAM-4 compliant
 - Compatible with USB 2.0 and forward-looking
 - Devices and hosts must be able to implement both the current MSC and the new protocol for software compatibility
- Define the protocol in INCITS T10
 - T10 WGs open to all; T10 plenary votes for T10 members only
 - INCITS/ANSI patent policy only requires RAND (not RAND-Z)
 - USB Device WG members can participate
 - Avoid reusing existing USB MSC material to avoid need for a copyright release
 - T10 can reference USB specifications, just not quote extensively from them

References

- **USB Mass Storage Class Specification Overview**, Revision 1.2, 23 June 2003
- **USB Mass Storage Class Control/Bulk/Interrupt Transport**, Revision 1.1, 23 June 2003 (used by floppy drives only)
- **USB Mass Storage Class Bulk-Only Transport**, Revision 1.0, 31 September 1999
 - USB mass storage specifications are available at the USB Implementers Forum at <http://www.usb.org>
- **SCSI Architecture Model – 4 (SAM-4)**, latest working draft on <http://www.t10.org/drafts.htm>
- T10/07-263 proposes a SAM annex listing the choices about optional features that each SCSI transport protocol has made (see <http://www.t10.org/doc07.htm>)
 - highlights some of the areas that this protocol will have to cover
- USB 3.0 – announced at Intel Developer Forum 18 September 2007



i n v e n t