

Subject: Predict Performance Benefits of Packetized SCSI (Preliminary)
Date: Wed, 03 Dec 97 15:08:00 CST

Craig McCombs
PATeam
Symbios, Inc.
12-3-97

Non-disconnecting 4KB Sequential Reads:

NTIOGEN benchmarks shows 3 Seagate Barracuda disks can saturate a wide Ultra-I SCSI bus @ 6,200 IOPs or 24.8 MB/sec. This produces an Ultra-I Wide SCSI data transfer efficiency of 62%. It also show the 7200RPM Seagate Barracuda can sustain 8.27 MB/sec data rate from media.

Now let's convert this to Ultra-III (160MB/sec) data rate. This would be (1/6,200IOPs)-102.4us+25.4us=84.29us. This is 11,864 IOPs or 47.46 MB/sec with an efficiency of 29.7%. Now, let's assume 10K RPM disk drives that can sustain 12 MB/sec from media. This means an Ultra-III SCSI bus, without packetized SCSI, is saturated with 3.96 disks.

Now let's factor in the packetized SCSI savings. Western Digital's spreadsheet shows 34,955 ns/IO (or 28,608 IOPs) for non-packetized Ultra-III SCSI without a SCSI disconnect and it shows 33,080 ns/IO (or 30,230 IOPs) for packetized Ultra-III SCSI without a SCSI disconnect. This shows a savings of 1.875 us/IO. Now I subtract 1.875us from the 84.29 us/IO above to get 82.42 us/IO. This means we can do 12,133 IOPs and move 48.53 MB/sec with an efficiency of 30.3%. But notice the SCSI bus is still saturated at 4.04 disks.

Disconnecting 4KB Random Reads and Writes:

Western Digital's spreadsheet shows 41,955 ns/IO (or 23,835 IOPs) for non-packetized Ultra-III SCSI with SCSI disconnects and it shows 37,460 ns/IO (or 26,695 IOPs) for packetized Ultra-III SCSI with SCSI disconnects.

It is my understanding that current TPC-C benchmarks generate about 75 random IOs per disk with a 70% reads and 30% writes. The performance is limited by disk latencies and mixed read/write IO workload profiles further reduce disk performance below that achieved with 100% read or 100% write IO profiles.

Let's assume new high performance disk technologies that could more than double the above random IO performance to 150 IOs per disk. Let's further assume we could use SCSI extenders to connect more disk on each SCSI bus.

The non-packetized Ultra-III SCSI bus wouldn't become a bottle neck until you connected 159 disks to a single Ultra-III wide SCSI bus. The packetized SCSI bus would allow you to connect 178 disks to a single Ultra-III wide SCSI bus. That's an improvement of 19 disks per Ultra-III wide SCSI bus.

Anyone believe we can or would connect that many drives to a single Ultra-III wide SCSI bus?

Single-Threaded Client and Workstation Market Segments:

The IO workload profiles for customers in these market segments would gain no detectable performance advantage from packetized SCSI.