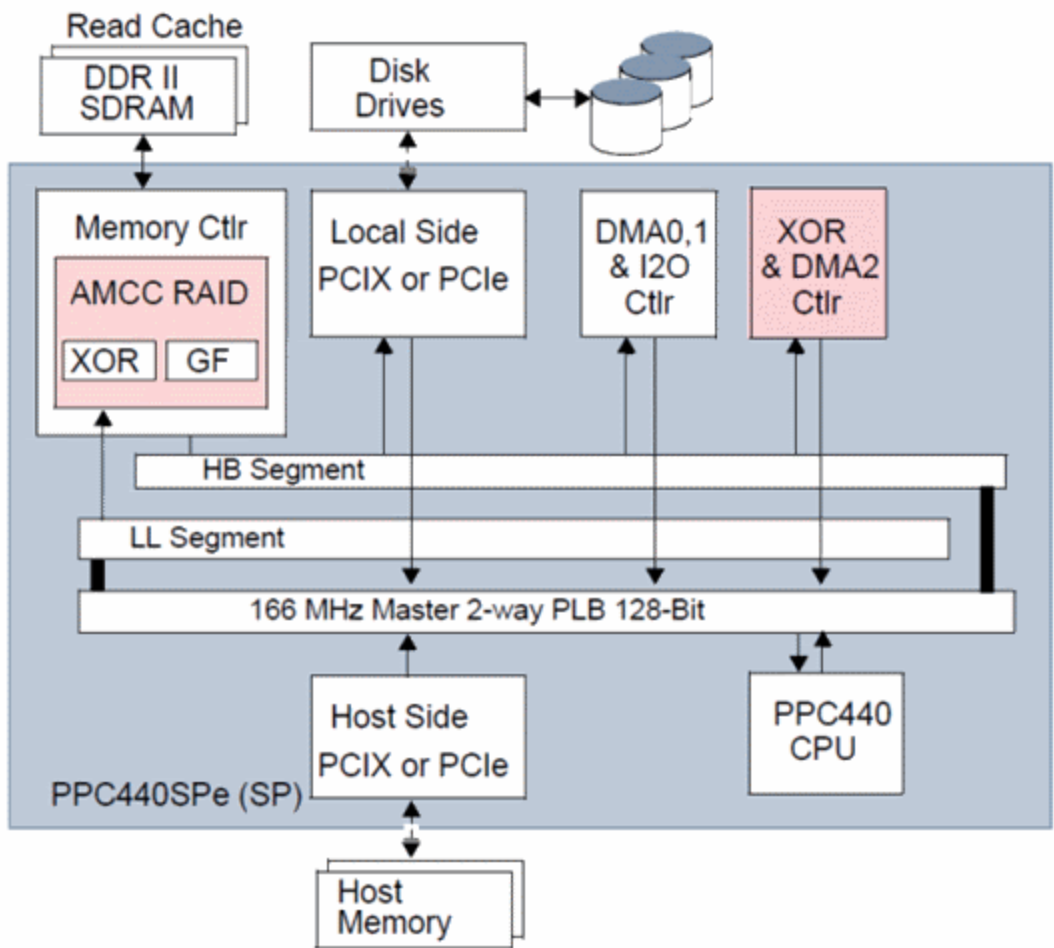





?:OSSLab Vxr??

??RAID?????????????Hardware RAID?Software RAID, ??????Hardware RAID???IOP(I/O Processor)???RAID ??, ??????XOR operation, ???RAID?????????. ??????parity value, ??????. ??AMCC?????-RAID 5/6 with PowerPC440SP/SPe, ???????:



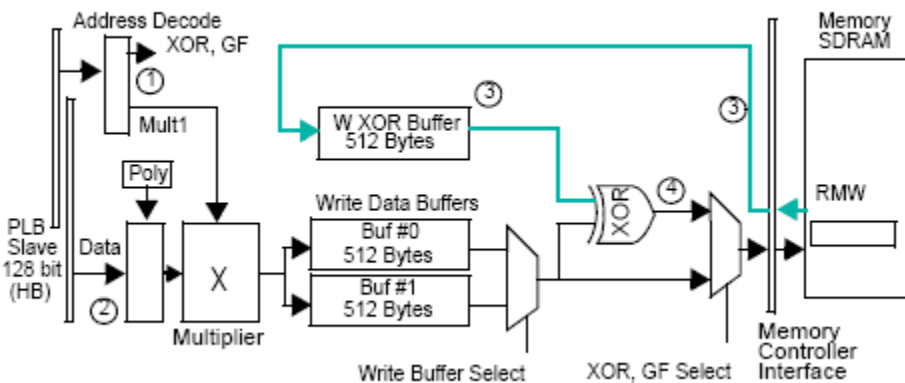
?????PPC440SPe?????IOP, ??????XOR operation, ???DMA Unit???XOR Assist, ?????? Memory Controller?XOR/GF function?????. ??????XOR operation?????, ?????? -Hardware RAID Assist, ??RAID????????????, ???RAID?. PPC440SPe?????????bus?, ??IBM? PLB(Processor Local Bus, CoreConnect?), Full-Duplex, crossbar?, 128bits, ?166MHz?, ??????10.4GB/s. ??????, ??????. Hardware RAID Assist????Intel?IOP????????, ???????:



????Intel????-Intel 81348 IO Processor Datasheet????, ??Three Application DMA Channels(ADMA;4KB???, Intel IOP333?AAU??1KB/512bytes????)????XOR operation??, ?????Memory Controller?, ??Cache Memory ??operand?result????, ?????bus(XSI): North?South, 128bits@400MHz, ???, ??????????. IOP348?RoC?, ?????????IO????????RAID Controller.

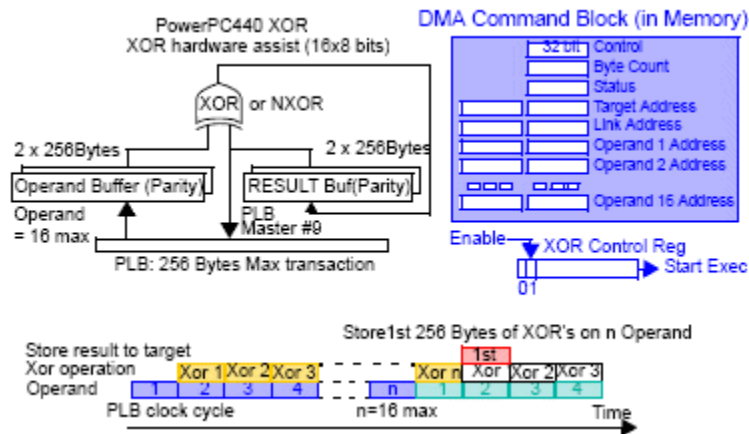
????????Hardware RAID???Hardware RAID Assist???, ?????I/O???, Hardware RAID Assist???????, ?????Cache?????operand?result????, ??????????????. ?????Hardware RAID Assist?????Memory Controller?? ?????DRAM buffer??Cache Memory. ???RAID?, ?????I/O???????, ??DRAM??Cache Memory????????? (WB&RA), ???Web Server???, DRAM buffer????????.

????AMCC????-AMCC RAID6 accelerator in a PowerPC IOP SOC, ?????DRAM buffer??Cache Memory??:



?????????, ??AMCC?RAID 6?XOR/GF?????, ??Memory Controller?????Hardware RAID Assist, ?????????? ??:XOR?GF function, RAID 6????????parity?????????, P?Q?parity, ??????????. RAID?????Stripping???

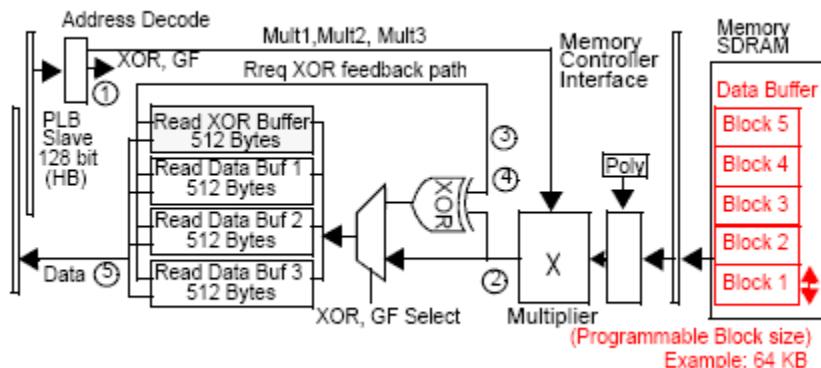
?????(Stripe)????, ?????64KB, ?????8??, ?????, 64/8=8KB, ??Stripe??8????8????(D1, D2, D3, ..., D8), ?????I/O?????. ?????RAID 0??. ??RAID 6????, ?????parity?, ??parity?????? ??, ?????D1, ..., D8, ???D1, ..., D6, P, Q. ?????, ?????N-2????, ?????, ?????. RAID 5????N-1????, ??parity????, ?????:



??DMA Unit??Hardware XOR Assist, ??DMA Unit?Cache Memory????XOR operation, ??result??Cache Memory. ???4??, ?????4-1????, ??Stripe Size??128KB, ?????, ?????32KB??, ?????32KB????parity? ?, ??D1, D2, D3, P???, P?value????D1 xor D2 xor D3 =P, ?????, AMCC????16????XOR operation??P value(????VD????), ??Cache Memory. DMA Unit????XOR operation? operand??.

??RAID 6???, ?????RAID 5??P value????, ??AMCC?RAID 6????Memory Controller?AMCC RAID?Hardware RAID Assist??GF??, GF function?Galois field????, ?????, ????? XOR????RAID 5??, ?????8??, ?RAID 6?, ??8-2????(????), P?Q????2/8, ????? ??Q????XOR????P value, ???D1 xor D2 xor D3 xor ... xor D7, ??Q parity??7/8??P????GF??, ? ??XOR operation????(????XOR buffer). ?????P?Q????, ?????, ??Cache Memory????? , ?????I/O????(????RAID 5/6?Cache Memory????).

????P, Q?parity????, ?????.




????Hardware RAID Assist????, ??RAID 6????, ?????, ?????. ?Hardware RAID????, ?????? ?????:??Hardware RAID Assist?Memory Controller?, ?????RAID????, ?????XOR/GF function????RAID?, Cache Memory????(WB/RA). ?????PPC440 CPU????built-in Processor, ???built-in Processor????, built-in Processor????I/O??RAID?, ??interrupt????Hardware Striping, ?????, ?????Cache Memory??

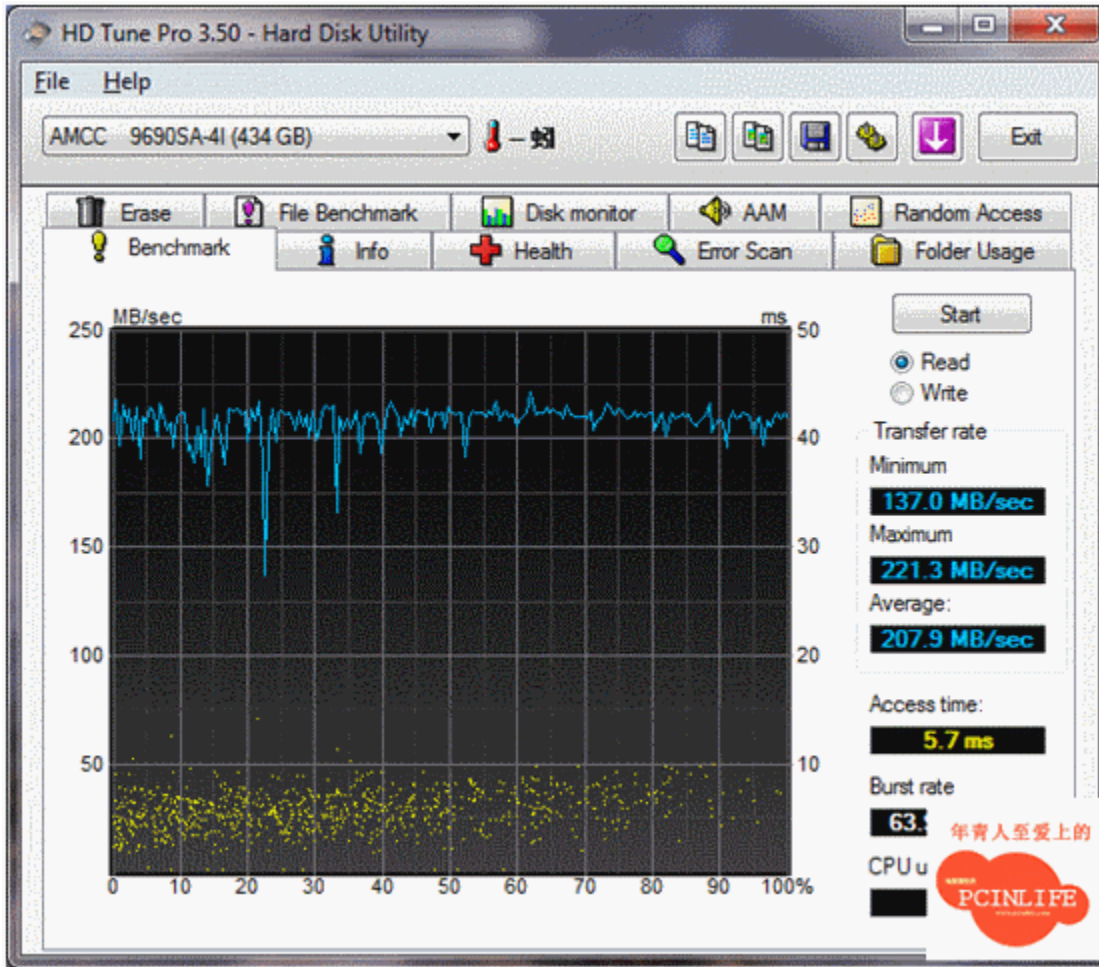
????, ??????????, Rebuild, OCE?ORLM?????. ??????????. ???built-in Processor?????RAID?????. ?????  
???RAID??-3ware StorSwitch, StorSwitch????????????????, ?????9650SE/9690SA, ???I/O Processor-  
PPC405CR(built-in processor), AMCC????DMA??(32?)XOR???-G133 RAID engine???I/O Controller????.  
9650SE???Marvell?IOC solution, 9690SA??Emulex?IOC500.



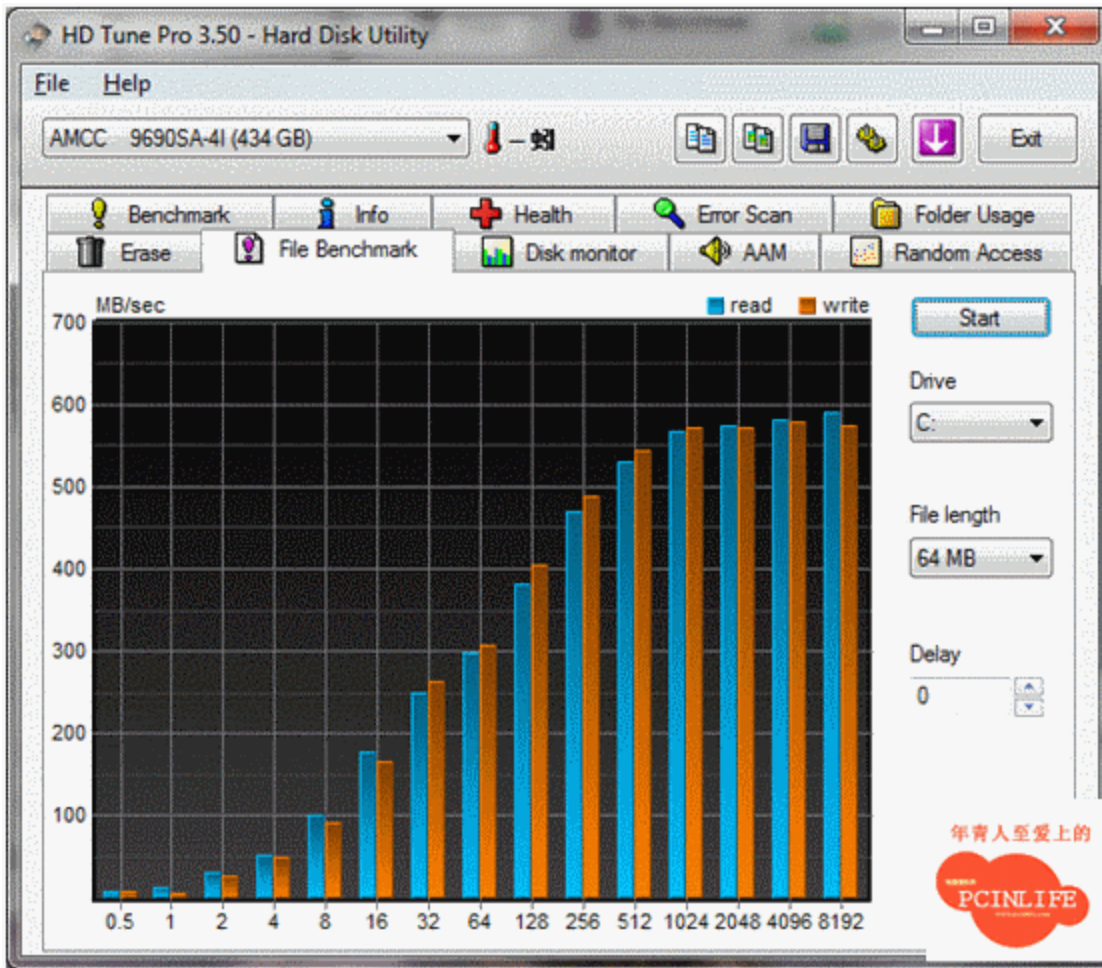
9690SA????????, ???StorSwitch????????????????LSI???, ???I/O????, ????. ?????9650SE RAID  
HBA?8?1TB FALS?HD TUNE????, ?????????????, ???LSI??????.



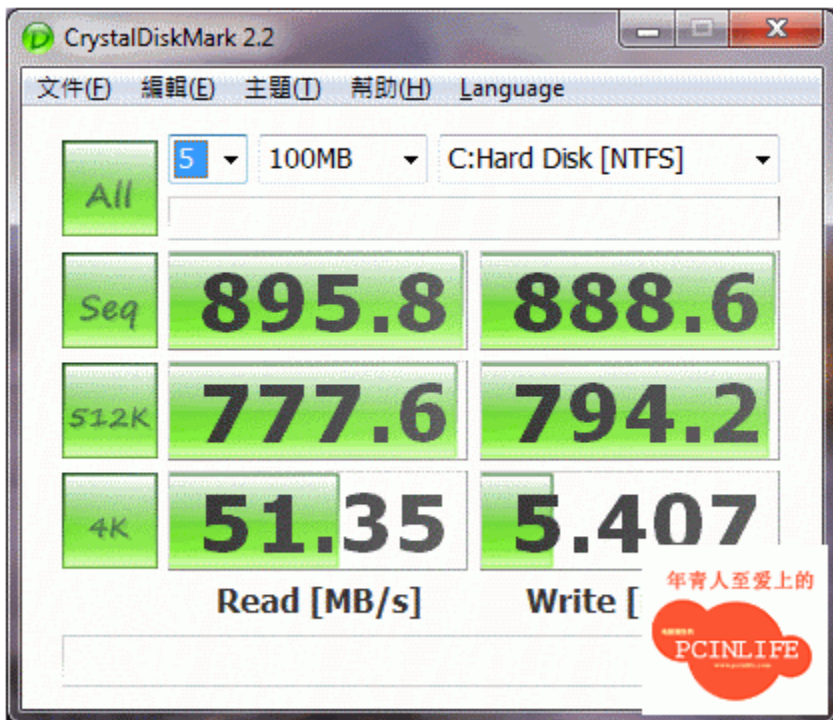
????9650SE??9690SA, ??????????LSISAS1078????, ??????????????. ???9690SA?4?SAS 15k. 6 146GB?RAID 5??  
???????




?????4?15k. 6 146GB?RAID 5???, StorSwitch?????????????, ??????????????. ??????????????




File Length?64MB, ?????????????, ??9690SA??????AMCC??XOR?????DDR2-533 512MB, ??64MB?????Cache Memory?ok?, ?????600MB?????????. ?????, ???Crystal Disk Mark???, ??????. ?100MB????????, ??? ????1GB???



???LSISAS1078?????, ??????. ???LSISAS1078???

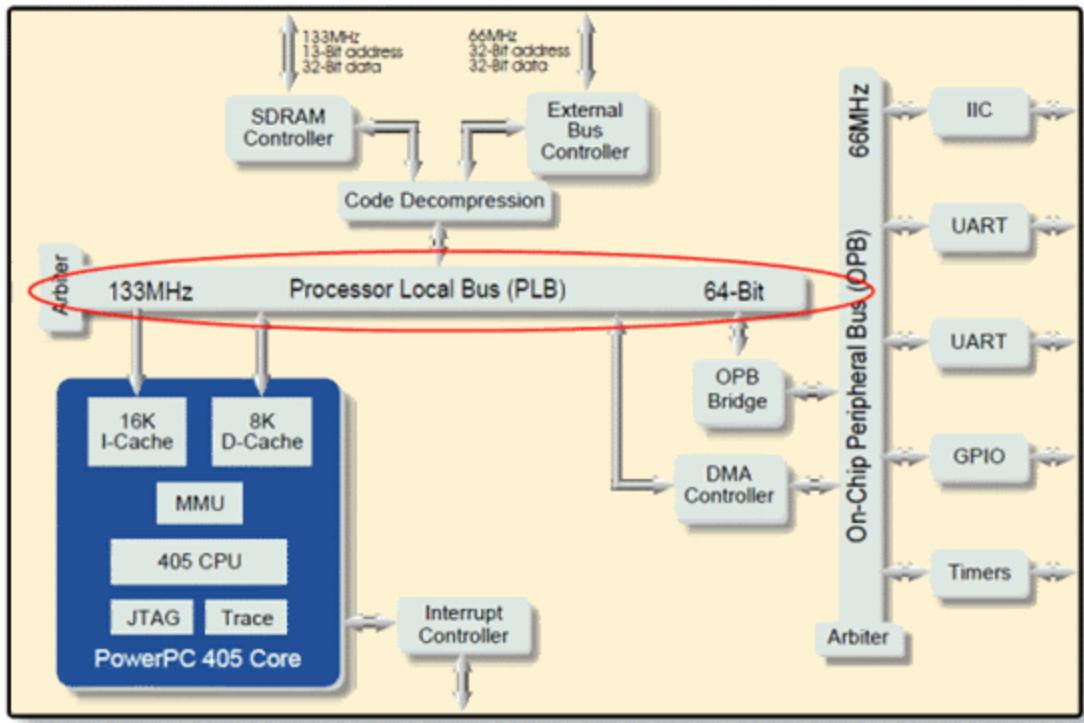


?????ATTO????????????????????, ????????:



?QD4????, ?????????????, ?????????700MB/s. ?????????, ??????, ?????????????, ?????????, ???IO????????, ??  
 IOC500??Intel?IOP34x?????????. ., ?????(?????????????. .), ?????????????IO??, ???AMCC DMA???XOR????????XOR???  
 ???2?DMA channels, ?????8????????????, ?????????????PM??, ??????????????. ??, ?????PPC405CR?I/O  
 Processor, ?????bottleneck?????I/O Processor?, ?????????I/O Processor?????, ?????I/O Processor?????:






PPC405CR?????I/O Processor, clock??266MHz, ??LSISAS1078?PPC440, ??????, ??AMCC?Product Selector Guide, ?????506~520DMIPS?, ??PPC440 Core????1000DMIPS?????, 2????????4????????(??), ??? MMU????????, ?????, ?????PPC Core?, ???PLB????, ???AMCC????????, ??PLB????????, 64bits@133MHz, ???SDRAM?clock?, ???crossbar?, ???, ?????1GB/s????, ?????????, ?????15k SAS?, ?????1GB/s?I/O???. 3ware??LSI??, ??RAID stack????LSI, StorSwitch????????????, ?????????.

?????????-built-in processor, ?????????????????????????????, ???RAID??, ??RoC?????, ??????:

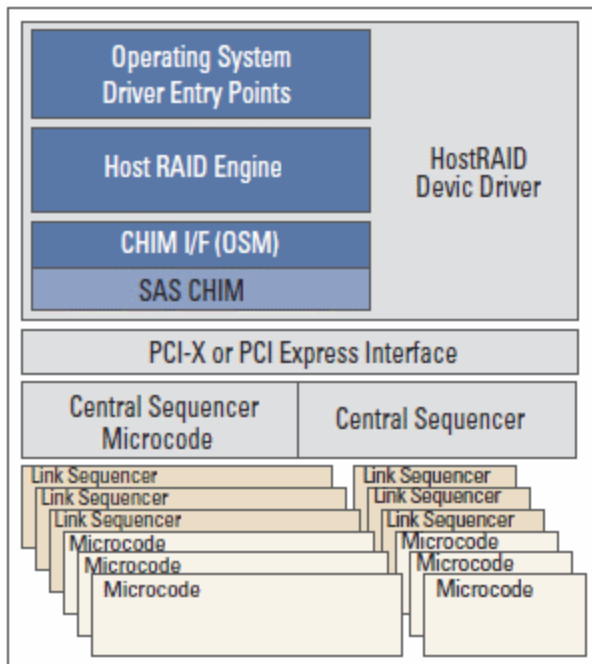
1. Layout????.
2. ??costdown.
3. ?????????Discrete RAID??.
4. ?????????????????.
5. ?????Hardware RAID Assist????????.

????????, ??I/O Processor, Hardware RAID Assist?IOC?RoC????????Discrete RAID?????. ?????LSISAS1078?? ??:

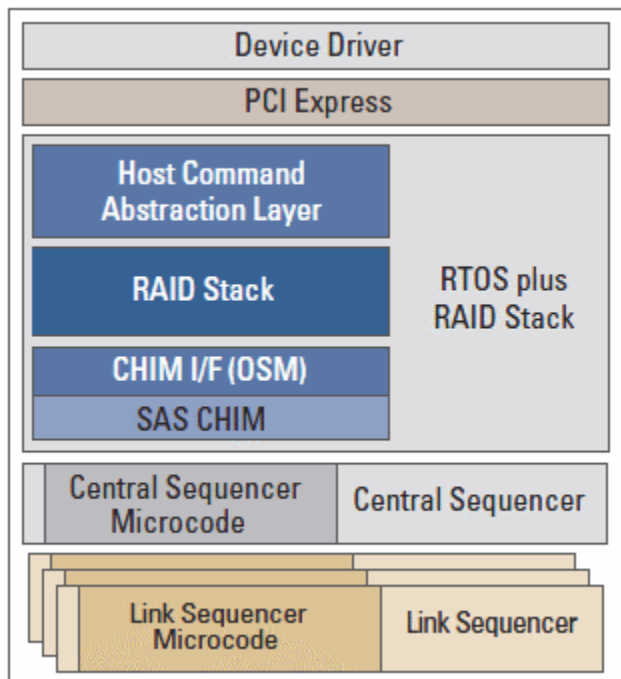
Figure 1. LSI SAS1078 Controller Block Diagram



?????, ???built-in Processor???, ??????????????????Hardware RAID Assist, ???RAID 0?????, ??????. ???????????  
 ?????:Driver-based RAID?IOP-based RAID. IOP-based RAID?????????Pure-Hardware RAID??, ??????????????????. ??  
 ??????????RAID stack;?????????????????Driver-based RAID??;????, ???Response Time;built-in Processor????, ???  
 ???Host-CPU?Driver-based RAID, ??????????. Driver-based RAID?Host-RAID Engine?????????????(driver)?, ??????  
 ??????????. ?????Driver-based RAID????, ??Adaptec by PMC?????:



HostRAID Device Driver???Host RAID Engine, ?????RAID????????????????, ?????????, ?????H/W Reset, ???????VD?. ??, ???Driver-based RAID??Hardware RAID, ?????????NVS RAM??, ???write journaling??, ?????log(????checkpoint)?NVS RAM, ?????????, ??????????????????. ??, ?????????Hardware RAID?stack??, ????:

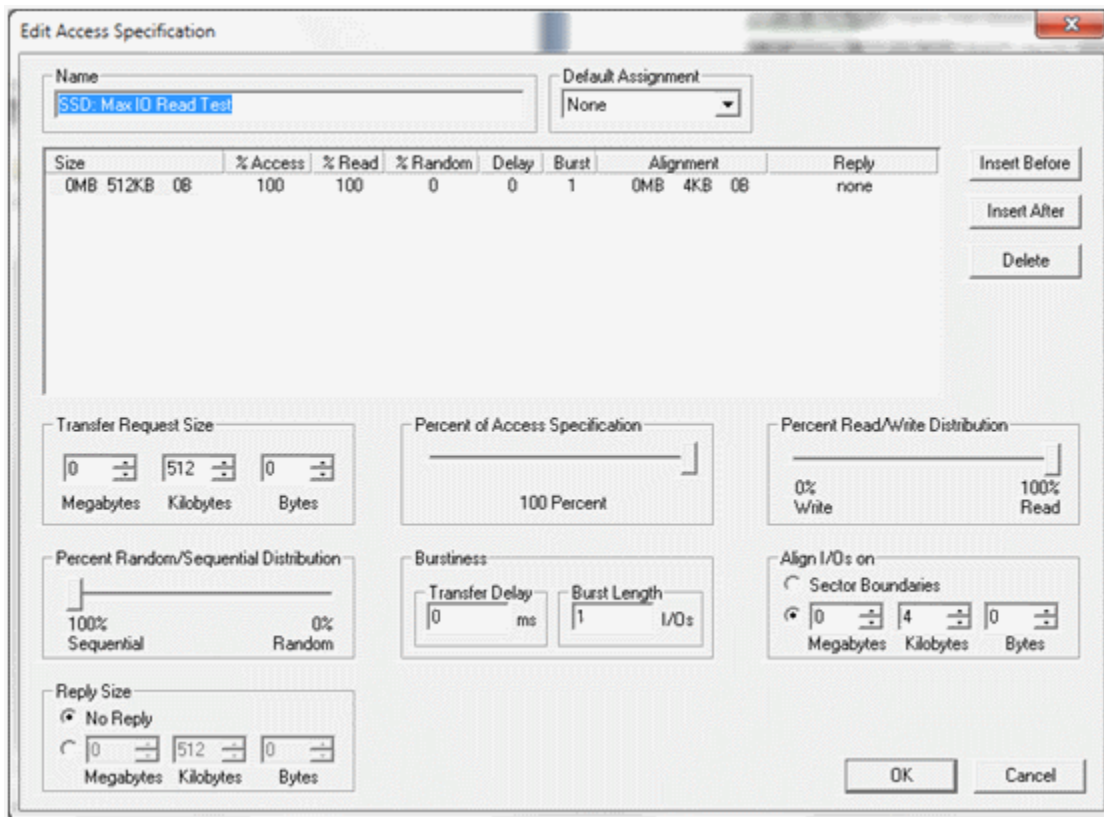


?????, Hardware RAID???, ??firmware????RAID stack????????????, ?????RAID????(Hardware RAID Assisr). ??Hardware RAID???, ??firmware????, ??LSI?LSISAS2008, ???RAIDKey??RAID5????, ?????? Hardware RAID Assist?. ???IMR???, ??Driver-based RAID??, ?????????, ???I/O????31?, ?LSISAS2108? 1008??I/O???, ??????. ???built-in Processor????, ?????????Host CPU. ?RAID 5???, ?????????????? ???. ???LSISAS2008????????HBA????????, ???IT???, ??built-in Processor????????.

????, ?????????, Driver-based RAID????Intel solution-ICH10R; IOP-based RAID????LSISAS2108?????. ???

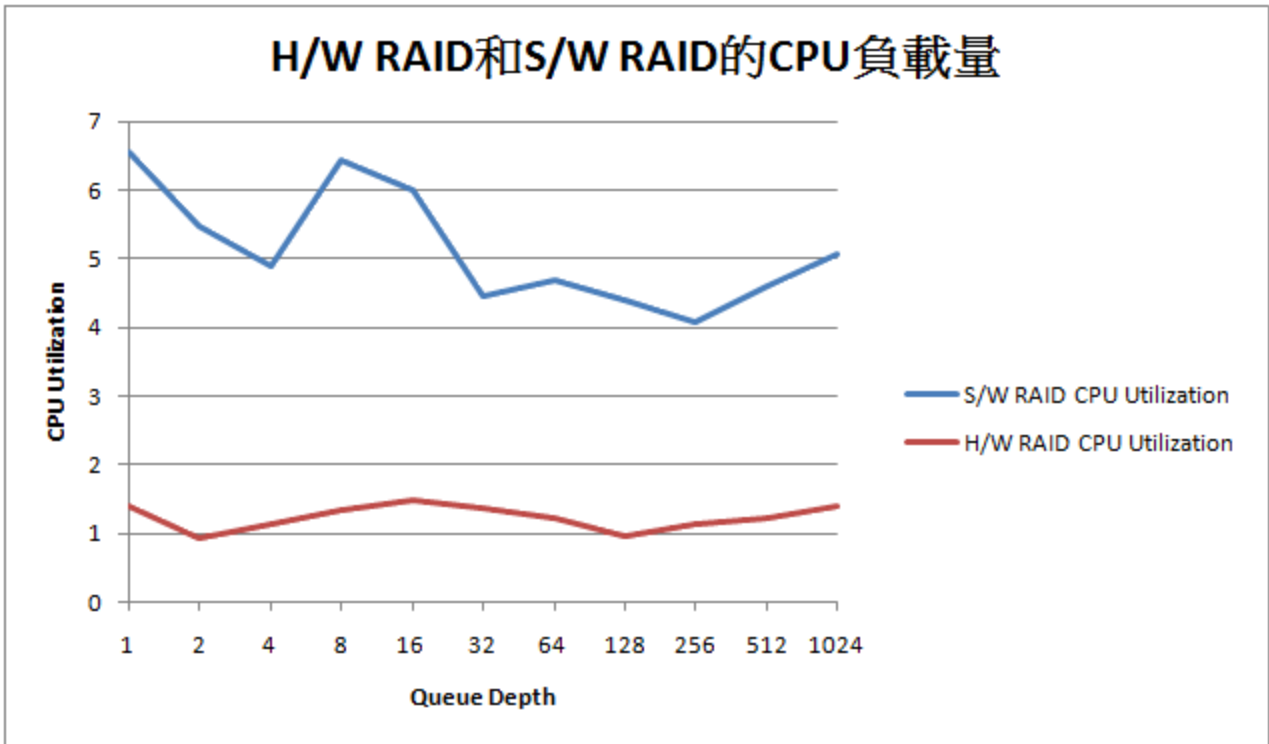
Crucial C300 128GB x2. ioMeter built-in Processor. Driver-based RAID. Host CPU, Intel Xeon W3570 4C/8T@3.20GHz CPU. Ok?

ioMeter I/O, :



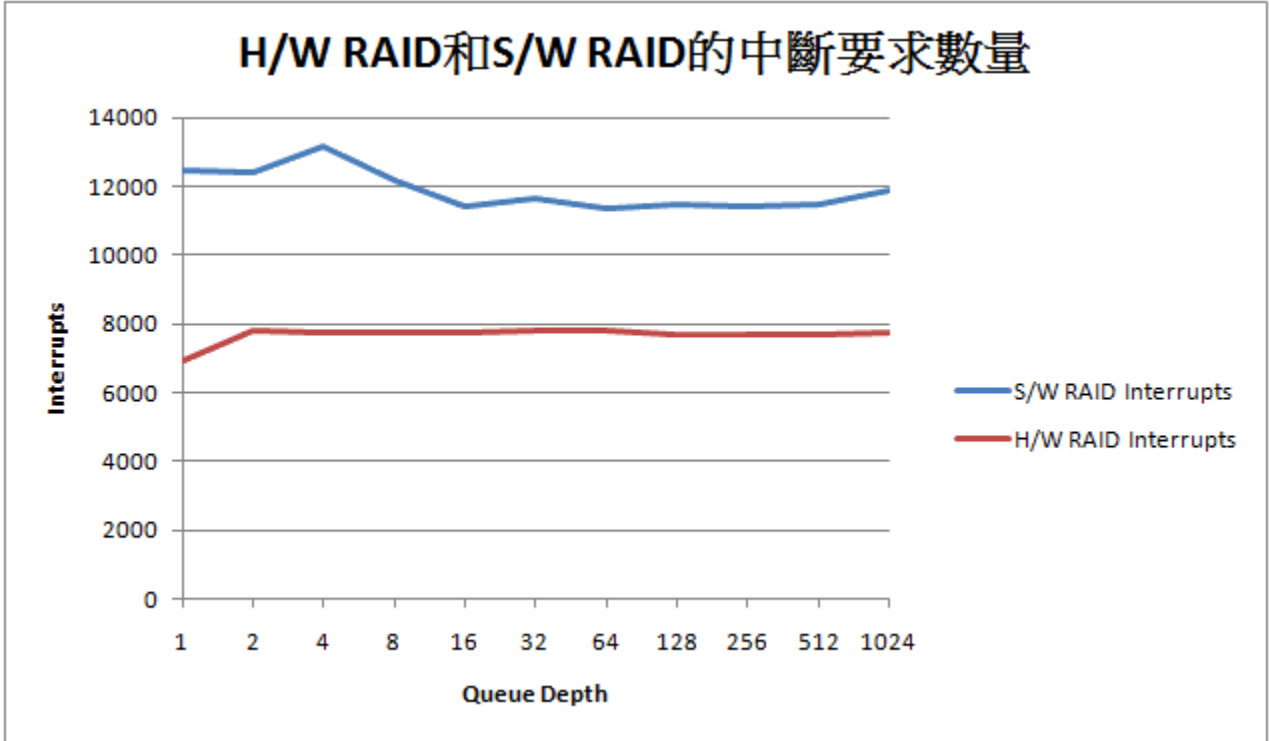
SSD, 4K, . QD. Stripe Size 64KB.

Driver-based RAID? IOP-based RAID? CPU:

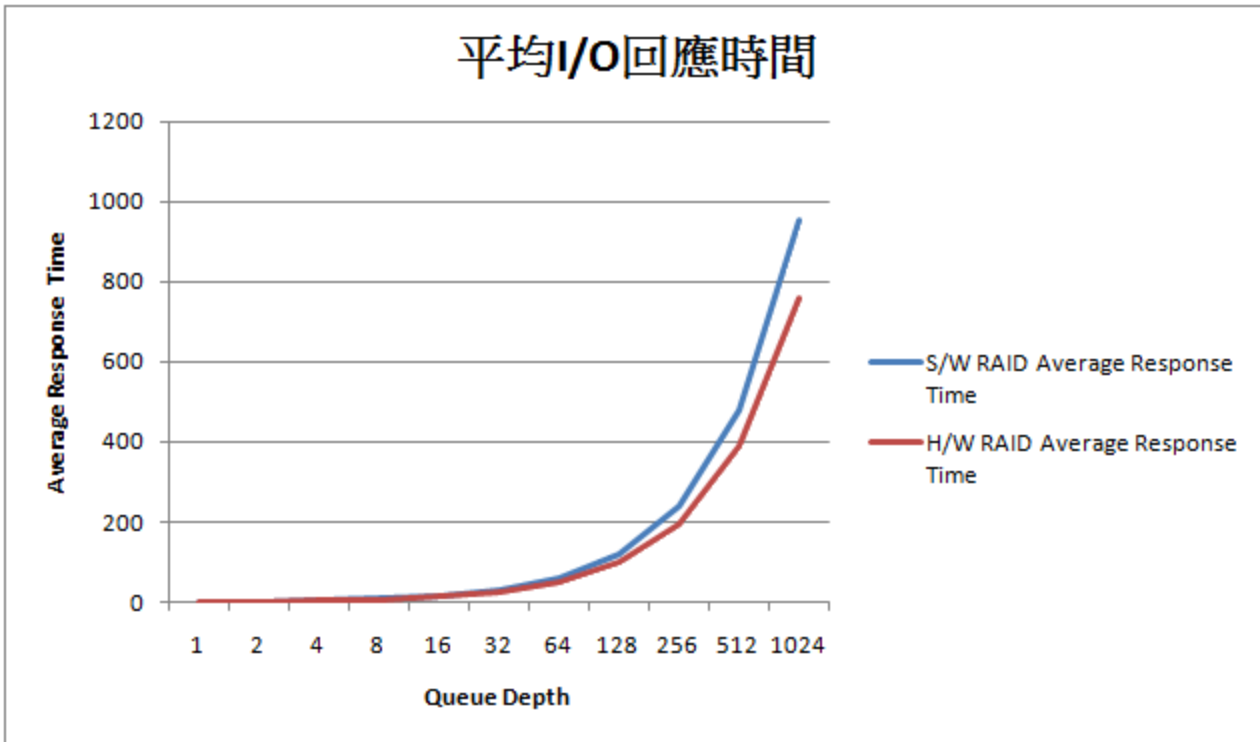


????????????????, ?????????, ?CPU?????????. ?????thread????????, ??????. ?????????, ???CPU Effectiveness??  
?????.

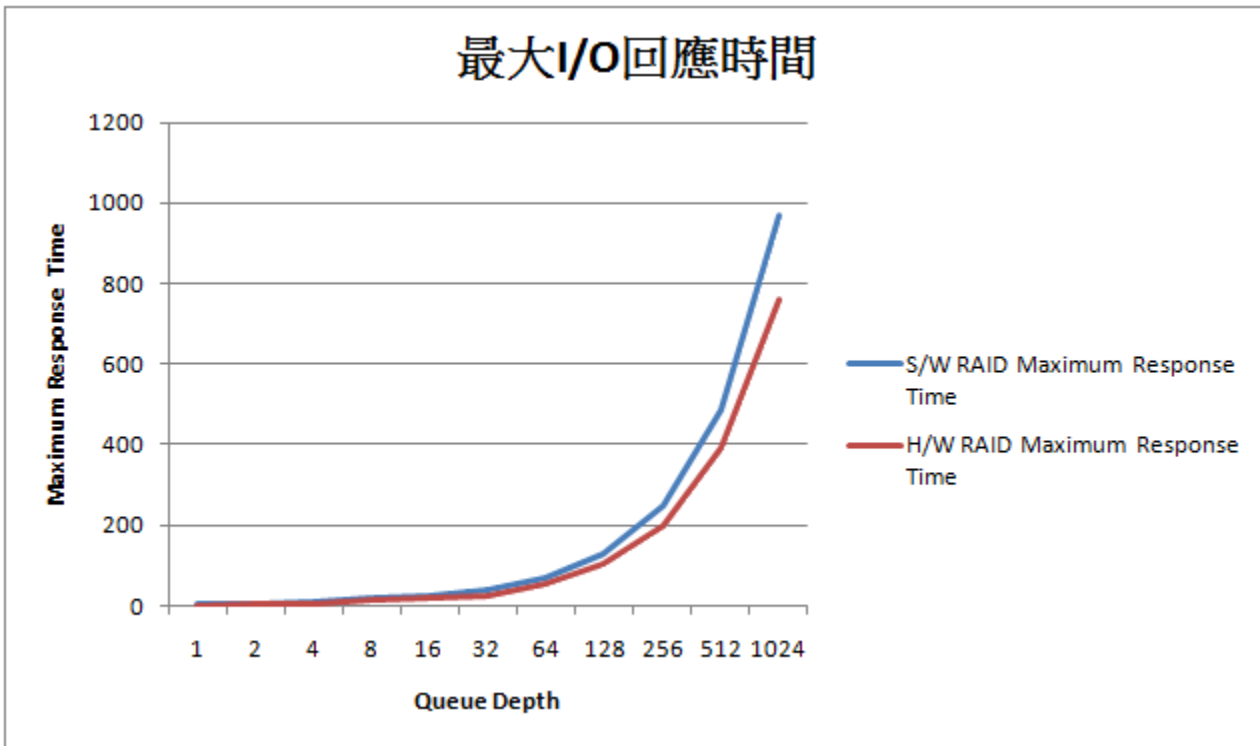
???????I/O?????Host CPU??????:



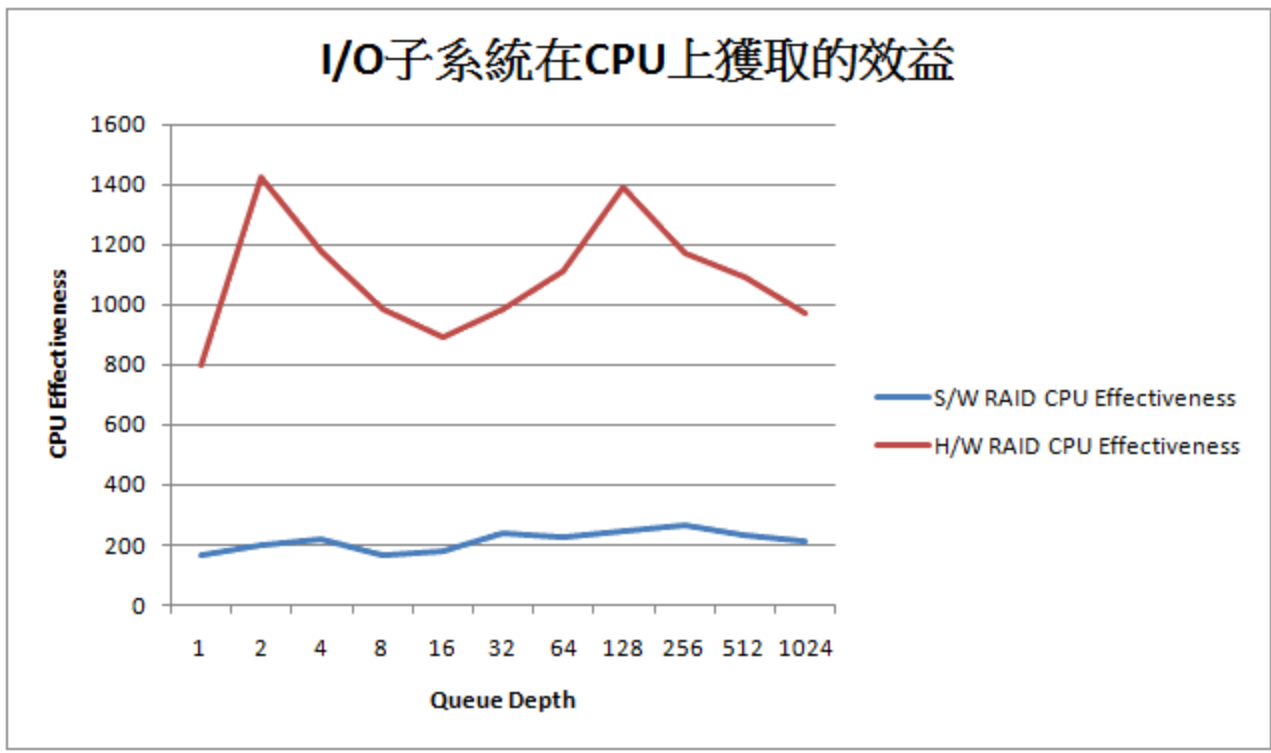
?????Hardware RAID????????????, ?Host CPU?????????????Driver-based RAID, built-in Processor?????????????. ?  
?????I/O?????????????:



?QD?????, Hardware RAID????, ?????I/O????, ???:



?????Hardware RAID????Driver-based RAID?????. ???????I/O?????CPU?????I/O??????:



OK, ...????????????, CPU Effectiveness????????, ?????????:

**CPU Effectiveness**= IOPS/CPU Utilization

????IOPS??, CPU????, ?????????, ?????????????I/O??, ??IOPS?????. ?????????IOP-based RAID??CPU????????

Driver-based RAID??, ?????????CPU Effectiveness.

??, ?????????Hardware RAID??Driver-based RAID?????, ?????????????, ?????built-in Processor?RAID??, ?????????.

H/W RAID?S/W RAID?CPU???(?????...):

