



Questions on Fifth Quiz

COSC 6360

Fall 2015



First question

- What is the main reason for partitioning an SMR disk into ***bands***?



Second question

- How is the persistent cache of a SMR disk drive using a set associative STL organized?



Third question

- What can we learn by writing blocks in the first 1GiB of a hard disk in random order and measuring the timings? (two-part answer)



Fourth question

- According to the authors of FAWN, where lies the sweet spot in CPU efficiency?



Fifth question

- Why does FAWN use a ***log-structured datastore***? (two-part answer)



Sixth question

- According to the authors of Xen, what is the main disadvantage of ***non-tagged TLBs***?
- How does Xen solve this problem?



Seventh question

- Consider a PIROGUE cluster with four servers. What would be the main advantage and disadvantages of replacing its lowest ranked server by a witness?
- ***Main advantage:***
 -
- ***Main disadvantage:***
 -