

Solutions to the Fourth COSC 6360 Quiz for Fall 2012

Jehan-François Pâris
jfparis@uh.edu



RAID Part A

- Consider a RAID level 6 array with ten disks.
 - Which fraction of the array space can be occupied by data? (5 points)



Answer

- Consider a RAID level 6 array with ten disks.
 - Which fraction of the array space can be occupied by data? (5 points)
 - **Each stripe contains 2 parity blocks and $10 - 2 = 8$ data blocks**
 - **80 percent**



RAID Part B

- Consider a RAID level 6 array with ten disks.
 - How would you update a block d and its corresponding parity blocks p and q ?



Answer

- How would you update a block d and its corresponding parity blocks p and q ?
 - Read *old block d , old parity block p and old parity block q*
 - Compute
$$\textit{new } p = \textit{old } p \oplus \textit{old } d \oplus \textit{new } d$$
$$\textit{new } q = \textit{old } q \oplus \textit{old } d \oplus \textit{new } d$$
 - Write *new d , new p and new q*




LFS

- What is the cost of a write in a log-structured file system, when its segment cleaner has to clean four segments to produce two clean segments? (10 points)



Answer

- What is the cost of a write in a log-structured file system, when its segment cleaner has to clean four segments to produce two clean segments? (10 points)
 - **We use the formula $\text{Cost} = 1/(1 - u)$ and note that $u = 50$ percent**
 - **Cost is 4 disk accesses per new block being written**



Answer (long way)

- When we clean up 4 segments, we must
 - Read the contents of the 4 segments
 - Write somewhere else the 2 segments containing live data
- The process gives us two clean segments that we can use later for new writes
- The total cost of these new writes is
 - $4 + 2 + 2 = 8$
 - 4 disk accesses per block being written



Journaling file systems

- What are the main advantages and disadvantages of using the ***data mode*** in journaling file systems? (2×10 points)



Answer

- What are the main advantages and disadvantages of using the ***data mode*** in journaling file systems? (2×10 points)
 - ***Advantage: Safest solution***
 - ***Disadvantage: Slowest***



Soft updates

- Which ***dependency information*** do ***soft updates*** maintain? (10 points)



Answer

- Which ***dependency information*** do ***soft updates*** maintain? (10 points)
- They maintain ***dependency information*** about cached pieces of metadata **at the directory entry level**
 - ***This i-node must be updated before/after this directory entry***



NFS

- How can ***non-volatile RAM*** (NVRAM) improve the response time of an NFS server? (10 points)
- Is there a cheaper way to achieve the same result? (10 points) (Hint: just mention it)



Answer

- How can *non-volatile RAM* (NVRAM) improve the response time of an NFS server? (10 points)
 - **NVRAM allows servers to respond to write requests without waiting for the completion of their own write requests**
- Is there a cheaper way to achieve the same result? (10 points) (Hint: just mention it)



Answer

- How can *non-volatile RAM* (NVRAM) improve the response time of an NFS server? (10 points)
 - **NVRAM allows servers to respond to write requests without waiting for the completion of their own write requests**
- Is there a cheaper way to achieve the same result? (10 points) (Hint: just mention it)
 - ***Safe asynchronous writes***



NFS

- Why must NFS client requests be both ***self-contained*** and ***idempotent***? (2×10 points)



Answer

- Why must NFS client requests be both ***self-contained*** and ***idempotent***? (2×10 points)
 - They must be self-contained because NFS is *stateless* and keeps no record of previous requests
 - They must be idempotent to allow clients to *resend* any request for which they did not get a reply