

Storage Systems

NPTEL Course

Jan 2012

(Lecture 03)

K. Gopinath

Indian Institute of Science

Naming and Storing

- Two important functions:
 - Give a name to an object
 - Can involve some processing (eg. add name to an index)
 - Or, name itself may be computed based on contents
 - Store an object
 - May involve other reads or stores!
 - May involve significant computation
 - Compression, encryption, coding, or deduplication: remove redundancy
- May need to keep aux. information about object
 - Metadata vs data; recursion? (metadata about metadata...)
 - Metadata loss vs data loss
- Access (r/w): device specific aspects determine speed
 - Reads sequential, non-sequential or random
 - Writes in-place or out-of-place

Large persistent data structures

- For processing, parts need to be brought into mem.
 - Two copies: “in memory” and “on-disk”
 - Atomicity and Consistency issues
- Algorithmic aspects need to be carefully taken care of
 - Number of objects can be in billions
 - Size of object can be in gigabytes
 - As time progresses, newer algs needed as scale changes
 - Mail directories can have 1000's of msgs, each a file!
 - Creating a file requires locking directory
 - Concurrent creates to same directory may become lock-bound
 - Critical with web-level storage systems
 - Many newer models (eg. key-value stores) since c. 2000

Ext2 FS

(from Wikipedia)

