



kDFS

Marko Obrovac

September 30, 2009

kDFS - Motivation

- Clusters (and their file systems) based on the historical model: storage nodes **vs** computational nodes




- kDFS** - make use of the whole storage capacity present in a cluster

kDFS - Idea

- ▶ kDFS - the symmetric file system
 - ▶ federate all hard drives present in the cluster
 - ▶ guarantee transparent and efficient disk usage: data striping, redundancy
- ▶ Improvement of overall cluster usage
 - ▶ integrate kDFS with other services: scheduler, hotplug, checkpointing
 - ▶ take advantage of services' complementarity

kDFS - Current status



- ▶ Federation of storage space achieved through KDDM sets
- ▶ No metadata server - the metadata is fully distributed amongst all participating nodes
- ▶ Nodes that do not provide storage can also access files (diskless nodes)
- ▶ Integrated local checkpointing/restart mechanisms

To do

- ▶ Striping and redundancy
 - ▶ the FS should provide an adaptive striping mechanism in order to boost cluster's I/O performance
 - ▶ data should be replicated as often as possible to avoid SPOF symptoms
- ▶ Integration with other cluster services
 - ▶ checkpointing - provide cluster-wide incremental snapshots of files
 - ▶ scheduling - cooperate with the SSI scheduler to improve data locality
 - ▶ hotplug - handle human-fashioned node addition/removal by transferring meta-data and content files