

PersiFS: A Continuously Versioned Network File System

Austin Clements, Dan Ports, Ben Schmeckpeper, Hector Yuen

Thursday, May 5, 2005

What is PersiFS?

- ▶ A network file system

What is PersiFS?

- ▶ A network file system
- ▶ A durable file system

What is PersiFS?

- ▶ A network file system
- ▶ A durable file system
- ▶ A continuously versioned file system

Continuous Versioning

- ▶ Access to complete file system at any point in the past
- ▶ Like AFS's OldFiles, but for every point in time
- ▶ Like CVS, but automatic and continuous

User interface

- ▶ Magic root directory
- ▶ `/persifs/current` – The latest, read/write version
- ▶ `/persifs/2004-05-25-12-00-00` – Read-only snapshot as of high noon, last towel day

FS Components

- ▶ Chunking
 - ▶ Divides data into content-sensitive chunks for efficient storage of modifications

FS Components

- ▶ Chunking
 - ▶ Divides data into content-sensitive chunks for efficient storage of modifications
- ▶ Superblob
 - ▶ Stores chunks, indexed by chunk hashes
 - ▶ Implements content-addressable storage
 - ▶ Chunk fusion!

FS Components

- ▶ **Chunking**
 - ▶ Divides data into content-sensitive chunks for efficient storage of modifications
- ▶ **Superblob**
 - ▶ Stores chunks, indexed by chunk hashes
 - ▶ Implements content-addressable storage
 - ▶ Chunk fusion!
- ▶ **Inode log**
 - ▶ Stores sequence of inode changes over time
 - ▶ Uses periodic inode map snapshots for efficient replay
 - ▶ Compact!