
 SafeCloud


## Using Several Clouds to Improve Storage Dependability: From DepSky to SCFS and Beyond


Miguel P. Correia

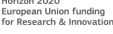
11th Cloud Control Workshop – June 2017


Joint work with [Alysson Bessani](#), B. Quaresma, F. André, P. Sousa, R. Mendes, T. Oliveira,  
N. Neves, M. Pasin, P. Verissimo, M. Pardal, E. Silva, F. Apolinário, D. Matos


 UNIVERSIDADE DE LISBOA

 TÉCNICO LISBOA

 European Commission


 Horizon 2020  
European Union Funding  
for Research & Innovation

 FCT Fundação para a Ciência e a Tecnologia  
MINISTÉRIO DA CIÊNCIA, TECNOLOGIA E ENSINO SUPERIOR

 inesc id lisboa

## Cloud-of-Clouds

- Consumer runs service on a **set of clouds** forming a **virtual cloud**, what we call a **cloud-of-clouds**
- Related to the notion of federation of clouds
  - **Federation of clouds** – a virtual cloud created by cloud providers; requires cooperation between providers
  - **Cloud-of-clouds** – an ad-hoc virtual cloud created by consumers; no cooperation between clouds needed



## Cloud-of-Clouds dependability+security

- There is **redundancy** and **diversity** between clouds
- so even if some clouds fail a **cloud-of-clouds** that implements **replication** can still guarantee:
  - **Availability** – if some stop, the others are still there
  - **Integrity** – if some corrupt data, data is still at the others
  - **Disaster-tolerance** – clouds can be geographically far
  - **No vendor lock-in** – several clouds anyway
- plus, although, not specific to cloud-of-clouds:
  - **Confidentiality** (from clouds) – encryption
  - **Confidentiality/integrity** (from users) – access control

4

## Outline

- **DepSky** – file storage in clouds-of-clouds
- **SCFS** – file system in clouds-of-clouds
- **SafeCloud-FS** – file system in clouds-of-clouds


5

## DEPSKY: FILE STORAGE IN CLOUDS-OF-CLOUDS

6

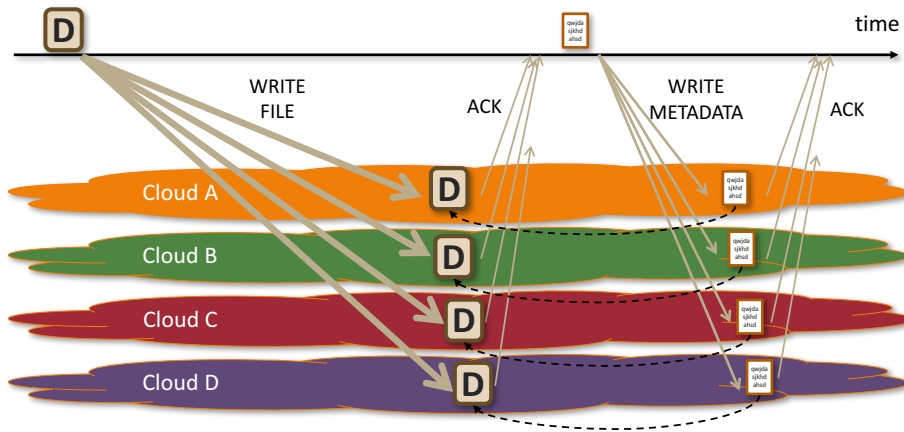
## DepSky

- Client-side library for cloud-of-clouds storage
  - File storage, similar to Amazon S3: read/write files, etc.
- Use storage cloud services (S3, etc.) as they are:
  - All code at the client
- Data is updatable
  - Byzantine quorum replication protocols for consistency



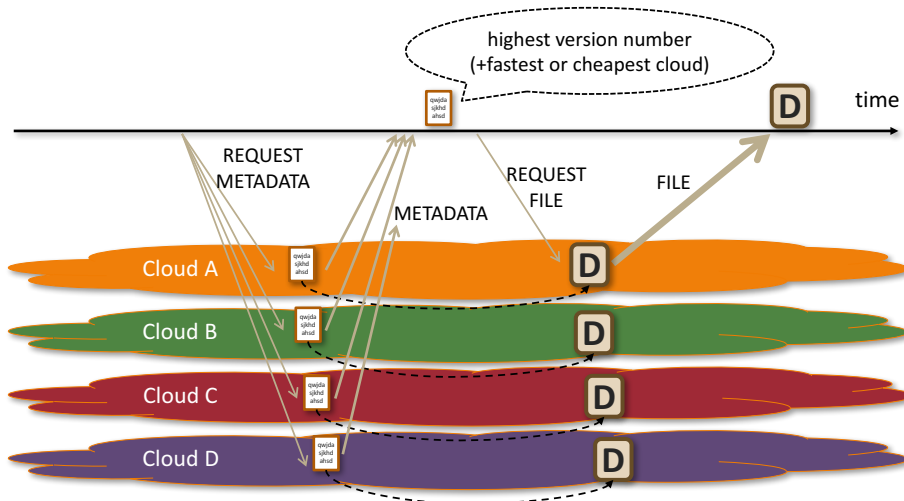
7

# Write protocol



8

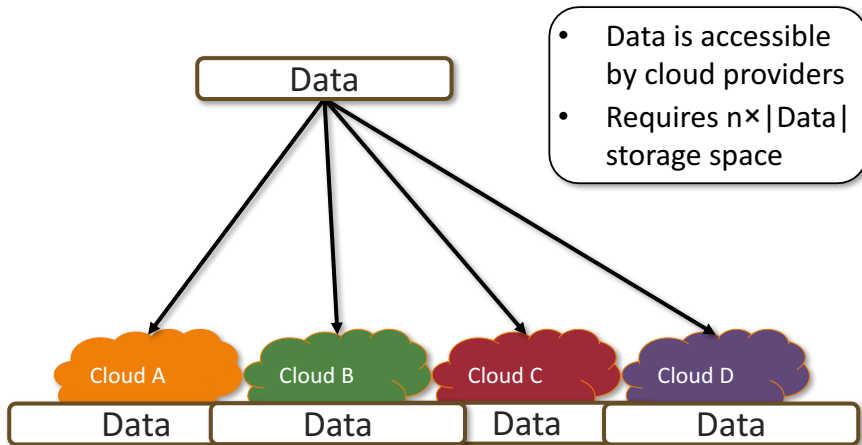
# Read protocol



File is fetched from other clouds if signature doesn't match the file

9

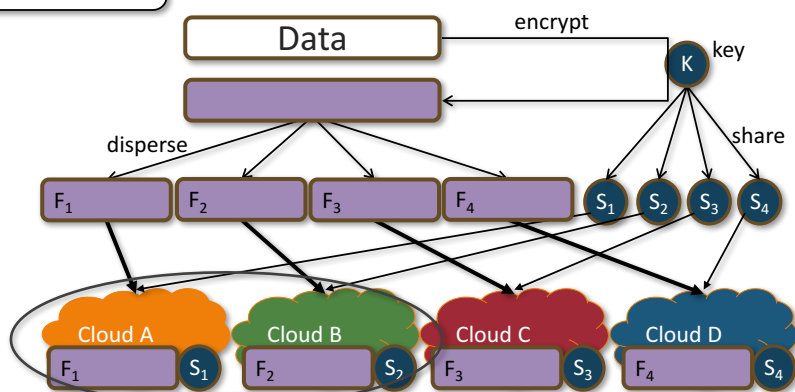
## DepSky-A: limitations




10

## DepSky-CA: combining erasure codes and secret sharing

Only for data, not metadata



Encrypted so data can't be read at a cloud!  
 Only 2x the size of storage, not 4x!



## SCFS: FILE SYSTEM IN CLOUDS-OF-CLOUDS

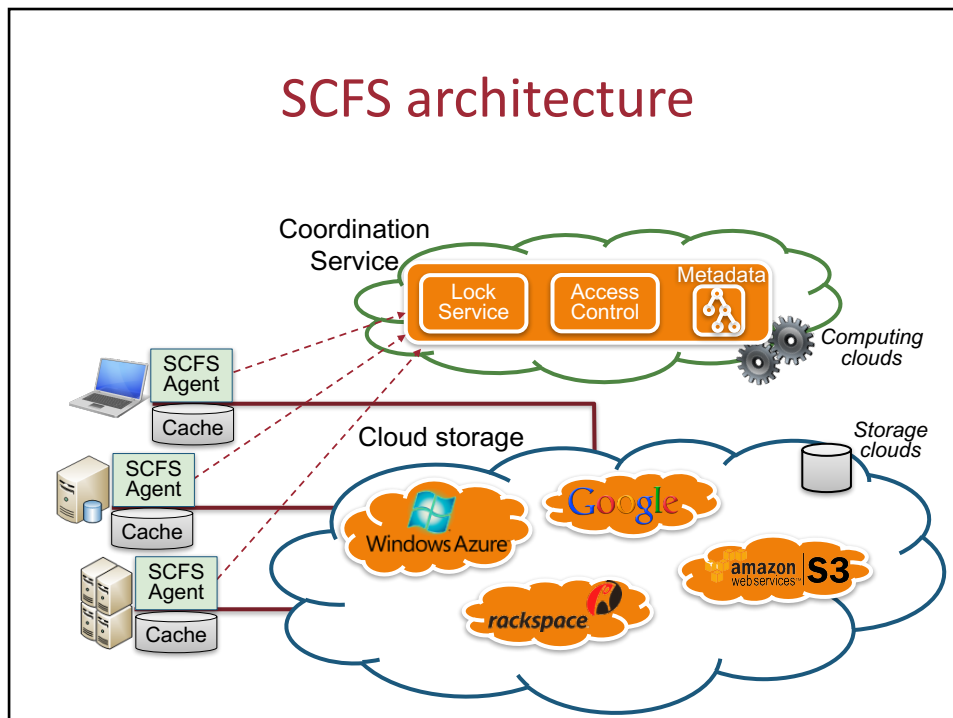
15

## Storage vs. File System (DepSky vs. SCFS)

- **Storage (DepSky)**
  - API: simple operations over data blocks
  - same consistency as clouds
  - **create**(id)
  - **read**(fd)
  - **write**(fd,block)
  - **delete**(fd)
  - **lock**(fd)
  - **unlock**(fd)
  - **setACL**(fd)
- **File system (SCFS)**
  - API: ~POSIX, so unmodified apps can use it (uses FUSE)
  - strong consistency
  - **open**(path,flags)
  - **read**(fd,buffer,length,offset)
  - **write**(fd,buffer,length,offset)
  - **chmod**(path,mode)
  - **mkdir**(path,mode)
  - **flush**, **fsync**, **link**, **rmdir**, **symlink**, **chown**,...

16

## SCFS architecture



## Features

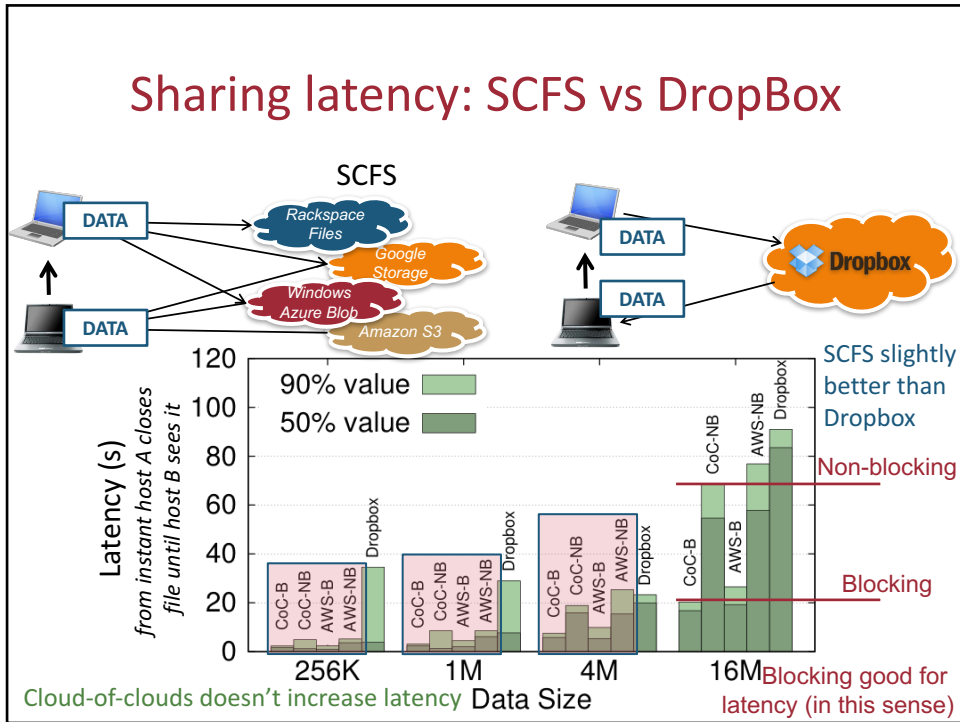
- Data layout/access pattern
  - Each file is an object (single-block file)
  - Multiple versions of the files are maintained
  - Always write, avoid reading (exploiting free writes)
- Caching
  - File cache: persistent (to avoid reading)
    - Local storage is used to hold copies of all/most client files
    - Opened files are also maintained in main-memory
  - Metadata cache: short-lived, main-memory
    - To deal with bursts of *metadata* requests

## Features

- Consistency
  - Consistency-on-close semantics
    - when user closes a file, all updates he did become observable by the rest of the users
  - Locks to avoid write-write conflicts
- Modular coordination
  - Metadata is stored in a coordination service
    - e.g., Apache Zookeeper (crash fault-tolerant), our own DepSpace (Byzantine/intrusion-tolerant)
  - Also used for managing file locks
  - Separate data from metadata



## Sharing latency: SCFS vs DropBox



## SAFECLOUD-FS – AN ENHANCED CLOUD-OF-CLOUDS FILE SYSTEM

## SCFS: opportunities

- **Coordination service** stores metadata (e.g., filenames, directories) in clear
- **Integrity verification** of data stored in a cloud requires first downloading the data
- **Intrusion recovery** – when a user account is compromised and data corrupted, recovery has to be done manually

26

## SafeCloud-FS

- Based on SCFS, with the features just explained, but:
- **Coordination service** HomomorphicSpace
  - Based on DepSpace but supports homomorphic operations
  - Based on the MorphicLib library (Java)
    - Operations: searchable, order preserving, summable, multipliable
  - Stores file metadata encrypted
- **Integrity verification:** SafeAudit
  - integrity verification of stored data without downloading it, using homomorphic signatures
- **Intrusion recovery** automatically with SafeRCloud

27

## WRAP-UP

28

## Conclusions

- Masking faults / intrusions using clouds-of-clouds
- **DepSky**: storage clouds-of-clouds
  - Availability, integrity, disaster-tolerance, no vendor lock-in, confidentiality
- **SCFS**: a cloud-backed file system
  - Based on DepSky and providing similar guarantees but near-POSIX API
- **SafeCloud-FS**: an enhanced cloud-backed file system
  - Ongoing work

29



## Thank you

- Papers:
  - DepSky: Dependable and Secure Storage in a Cloud-of-Clouds. ACM Transactions on Storage, 2013 (also EuroSys 2010)
  - SCFS: a Shared Cloud-backed File System. Usenix Annual Technical Conference (ATC), 2014
- Code:
  - DepSky: <http://cloud-of-clouds.github.io/depsky/>
  - SCFS: <http://cloud-of-clouds.github.io/SCFS/>
- My web: <http://www.gsd.inesc-id.pt/~mpc/>

