Cold Storage Data Archives: More Than Just a Bunch of Tapes

Bunjamin Memishi, Raja Appuswamy, and Marcus Paradies

Introduction

Introduction

Data-intensive sciences
- high data volume generation

Genomics

Physics

Earth Observation

Radio Astronomy, etc.

Introduction

Data-intensive sciences
- high data volume generation

Genomics

Physics

Earth Observation

Radio Astronomy, etc.

Storage
- private infrastructure

Introduction

Data-intensive sciences
- high data volume generation

Genomics

Physics

Earth Observation

Radio Astronomy, etc.

Storage
- private vs. public infrastructure

Microsoft Azure

Google Cloud

amazon web services™

Introduction

Data-intensive sciences
- high data volume generation

Genomics

Physics

Earth Observation

Radio Astronomy, etc.

Storage
- private vs. public infrastructure

Microsoft Azure

Google Cloud

amazon web services

The problem is, what to store and where?

Data analysis

Data analysis

weather forecast

radio astronomy

earth observation

Data analysis

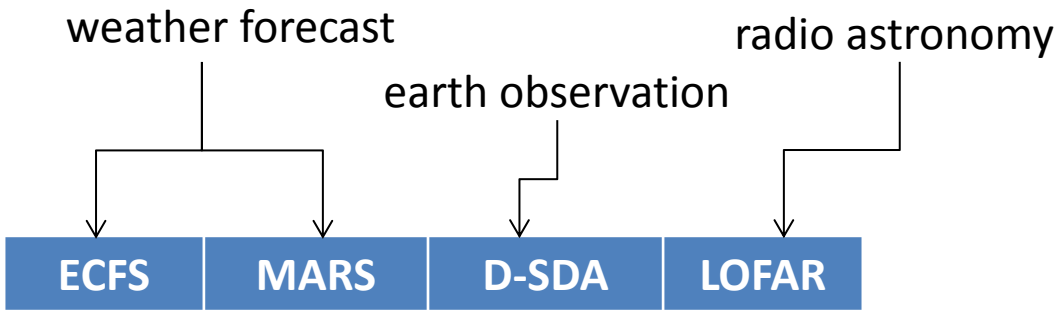

Data analysis: **Volume**

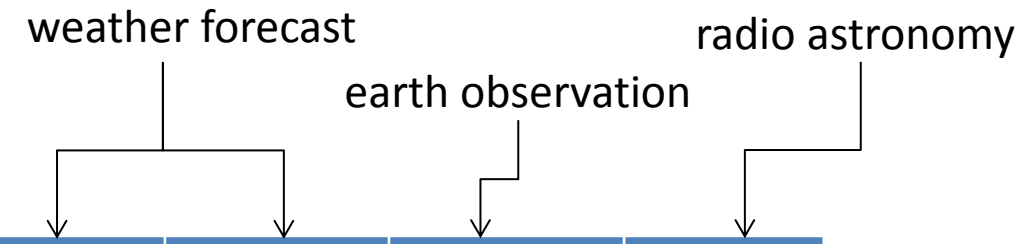

	ECFS	MARS	D-SDA	LOFAR
Total	14.8 PB	37.9 PB	15 PB	13 PB

Data analysis: Volume, **Variety**

weather forecast

radio astronomy

earth observation

	ECFS	MARS	D-SDA	LOFAR
Total	14.8 PB	37.9 PB	15 PB	13 PB

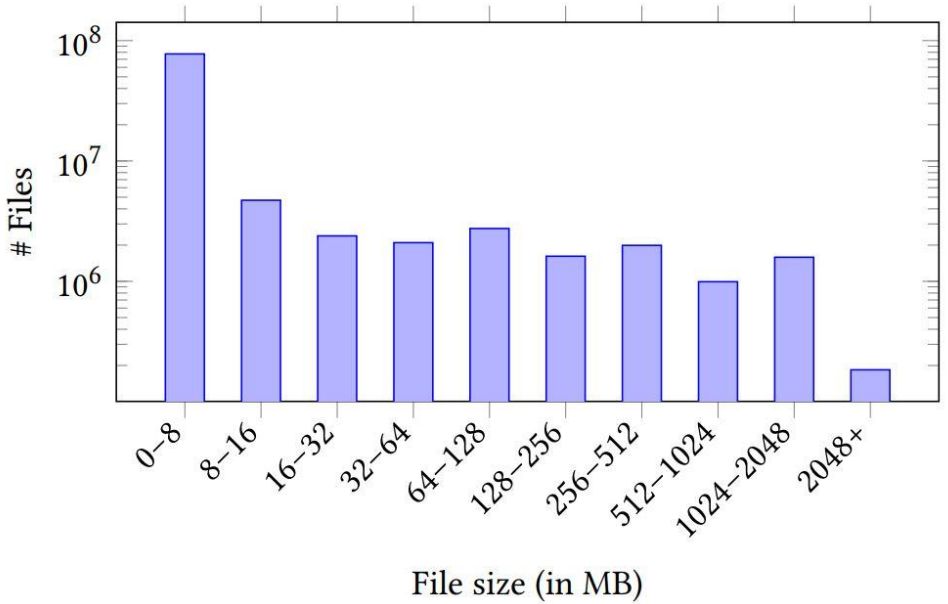

File size distribution: Main D-SDA product library

Data analysis: Volume, Variety, **Liveliness**

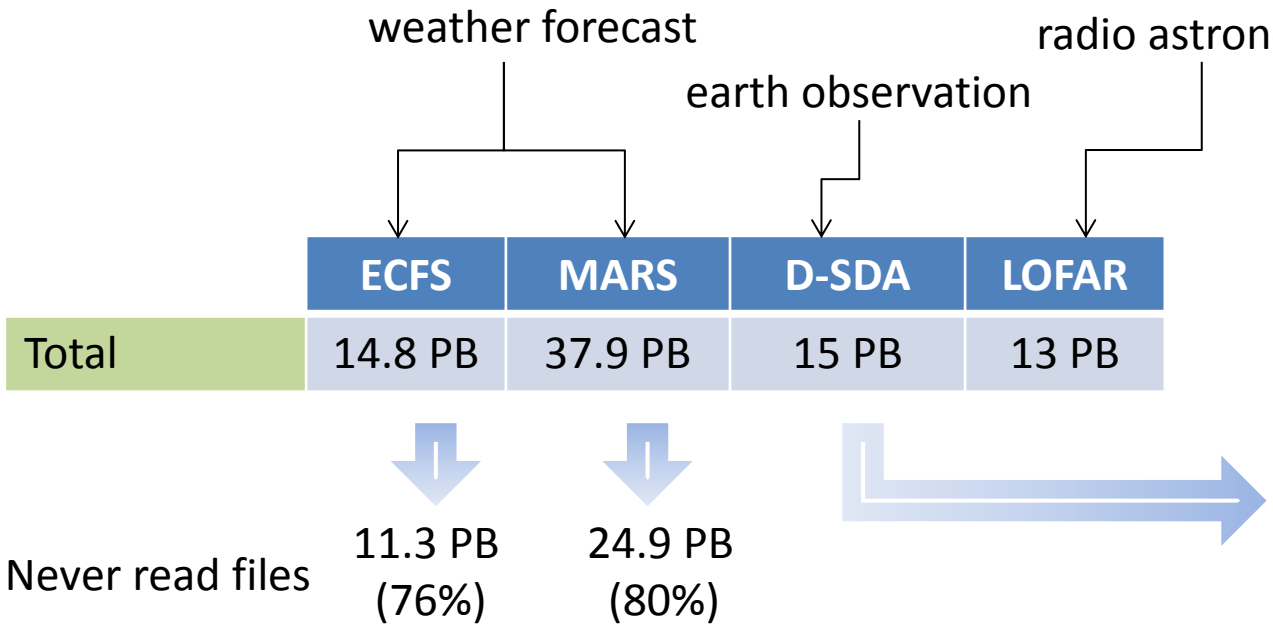

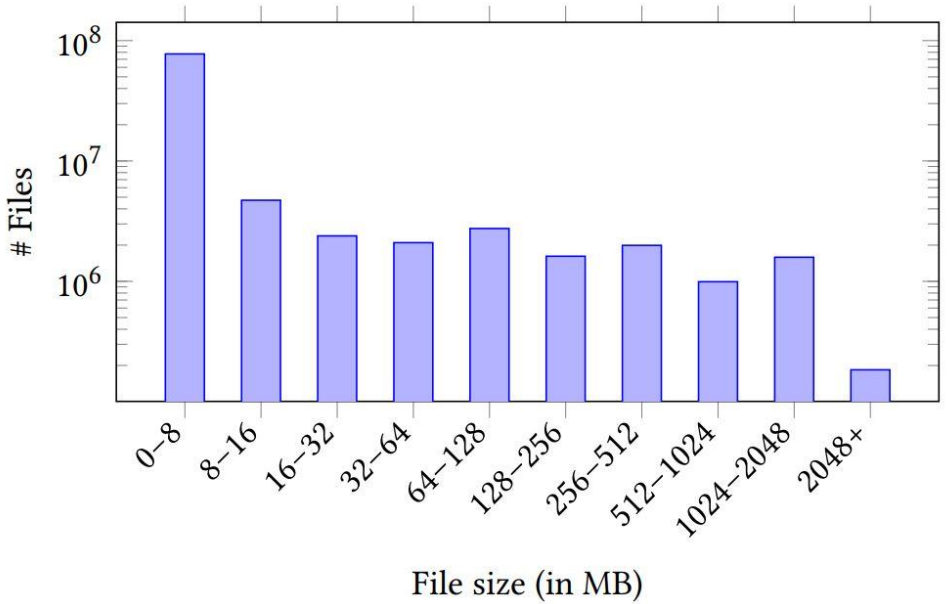

File size distribution: Main D-SDA product library

Data analysis: Volume, Variety, **Liveliness**

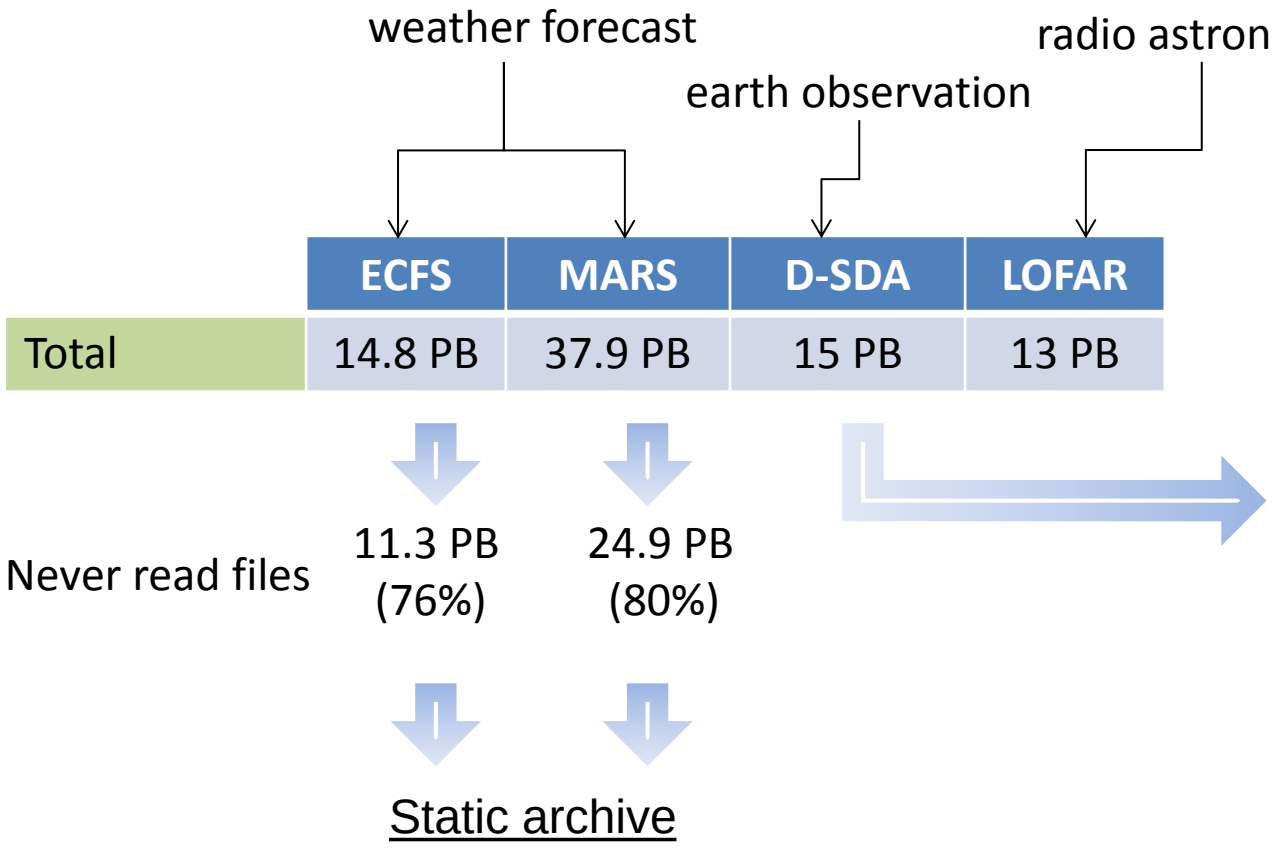

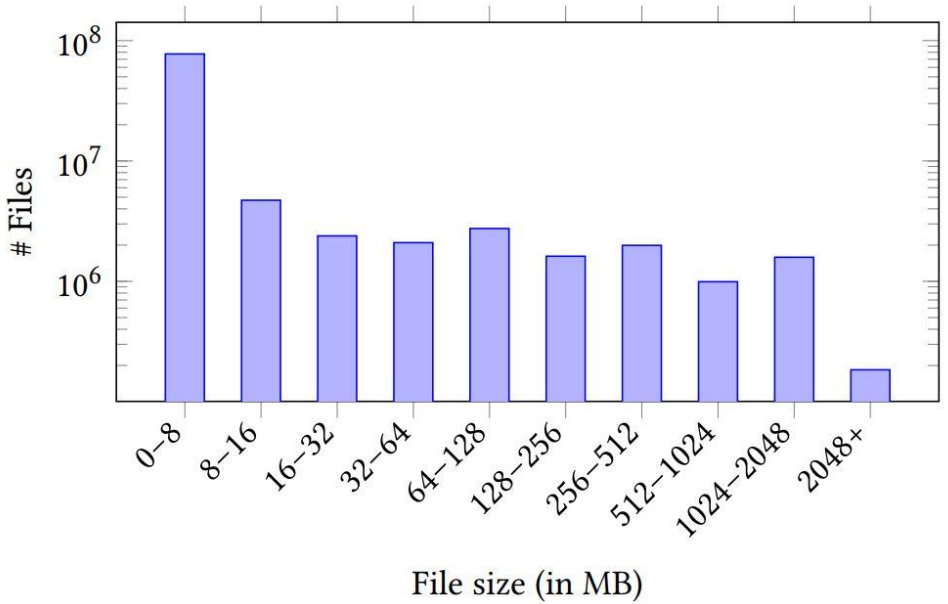

File size distribution: Main D-SDA product library

Data analysis: Volume, Variety, **Liveliness**

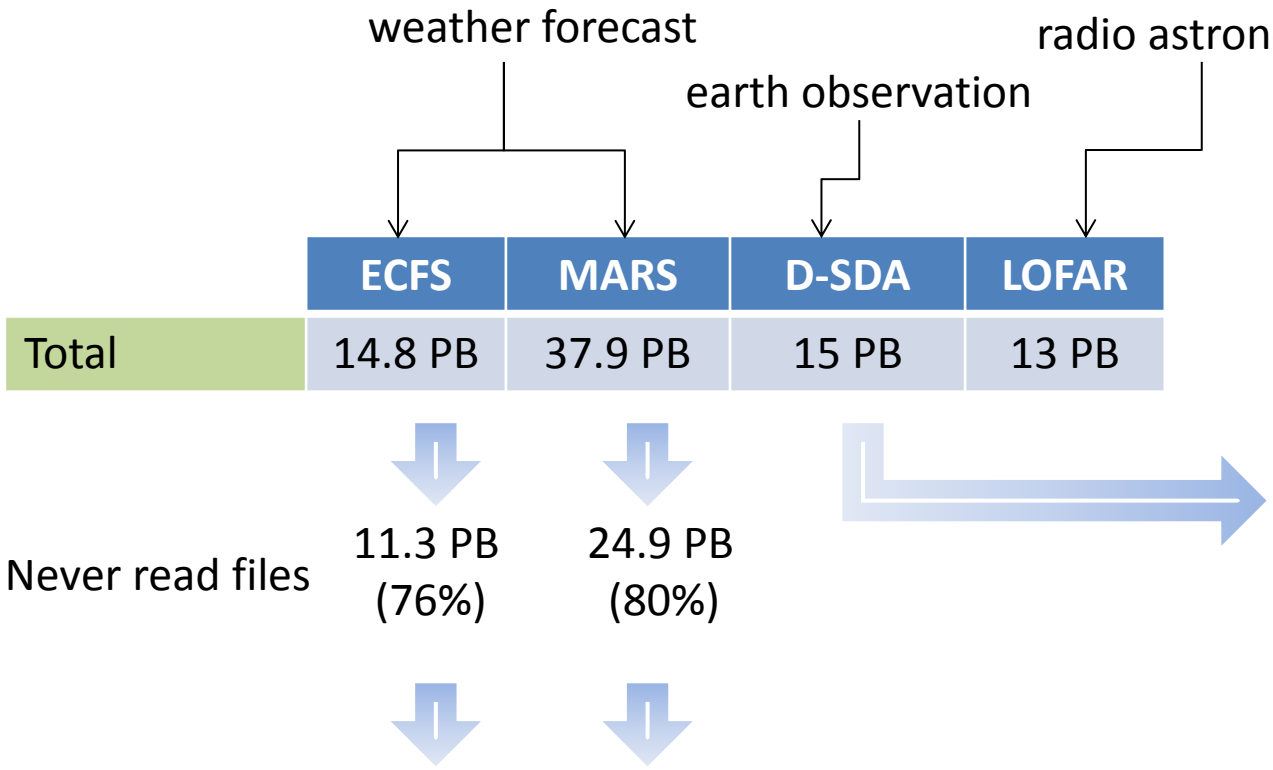

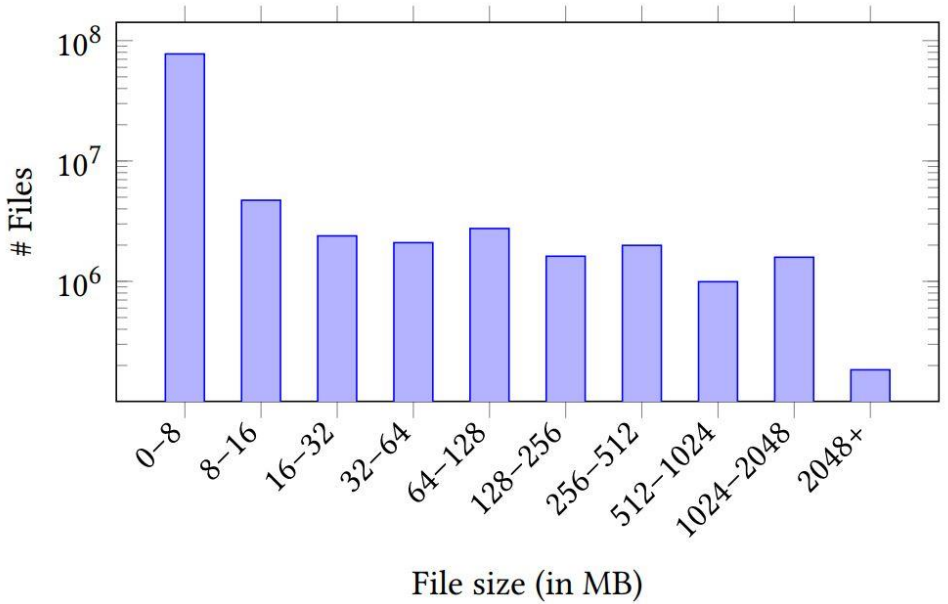

Active archive vs Static archive boundary: Important!

File size distribution: Main D-SDA product library

Deployment analysis

Deployment analysis: Private versus public infrastructures

Deployment analysis: Private versus public infrastructures

At public cloud offerings, generally, 3 types of storage classes

- deep archival
- nearline archival
- online service

Deployment analysis: Private versus public infrastructures

At public cloud offerings, generally, 3 types of storage classes

- deep archival
- nearline archival
- online service

Azure use case

	Storage (GB/Month)	Retrieval (per GB)	Latency
Azure Archival Blob (deep)	$0.0045	$0.02	Several hours
Azure Cool Blob (nearline)	$0.0334	$0.01	61.4 ms
Azure Hot Blob (online)	$0.0422	$0	5.3 ms

Deployment analysis: Private versus public infrastructures

At public cloud offerings, generally, 3 types of storage classes

- deep archival
- nearline archival
- online service

+ a good fit for never read back data

Azure use case

	Storage (GB/Month)	Retrieval (per GB)	Latency
Azure Archival Blob (deep)	$0.0045	$0.02	Several hours
Azure Cool Blob (nearline)	$0.0334	$0.01	61.4 ms
Azure Hot Blob (online)	$0.0422	$0	5.3 ms

Deployment analysis: Private versus public infrastructures

At public cloud offerings, generally, 3 types of storage classes

- deep archival
- nearline archival
- online service

+ a good fit for never read back data
- expensive data retrieval

Azure use case

	Storage (GB/Month)	Retrieval (per GB)	Latency
Azure Archival Blob (deep)	$0.0045	$0.02	Several hours
Azure Cool Blob (nearline)	$0.0334	$0.01	61.4 ms
Azure Hot Blob (online)	$0.0422	$0	5.3 ms

Deployment analysis: Private versus public infrastructures

At public cloud offerings, generally, 3 types of storage classes

- deep archival
- nearline archival
- online service

+ a good fit for never read back data
- expensive data retrieval

Azure use case

	Storage (GB/Month)	Retrieval (per GB)	Latency
Azure Archival Blob (deep)	$0.0045	$0.02	Several hours
Azure Cool Blob (nearline)	$0.0334	$0.01	61.4 ms
Azure Hot Blob (online)	$0.0422	$0	5.3 ms

Storage vs. Retrieval cost

- Important implications for deploying scientific applications in a public cloud

Deployment analysis

Deployment analysis: Storage–access trade-off in a cloud

Deployment analysis: Storage–access trade-off in a cloud

Use case:

- 1 PB scientific archive
- Storage: 1 year
- Retrieval: all read once during that year

Deployment analysis: Storage–access trade-off in a cloud

Use case:

- 1 PB scientific archive
- Storage: 1 year
- Retrieval: all read once during that year

	Cost
Azure Archival Blob	$79K
Azure Cool Blob	$430K
Azure Hot Blob	$531K

Deployment analysis: Storage–access trade-off in a cloud

Use case:

- 1 PB scientific archive
- Storage: 1 year
- Retrieval: all read once during that year

	Cost
Azure Archival Blob	$79K
Azure Cool Blob	$430K
Azure Hot Blob	$531K

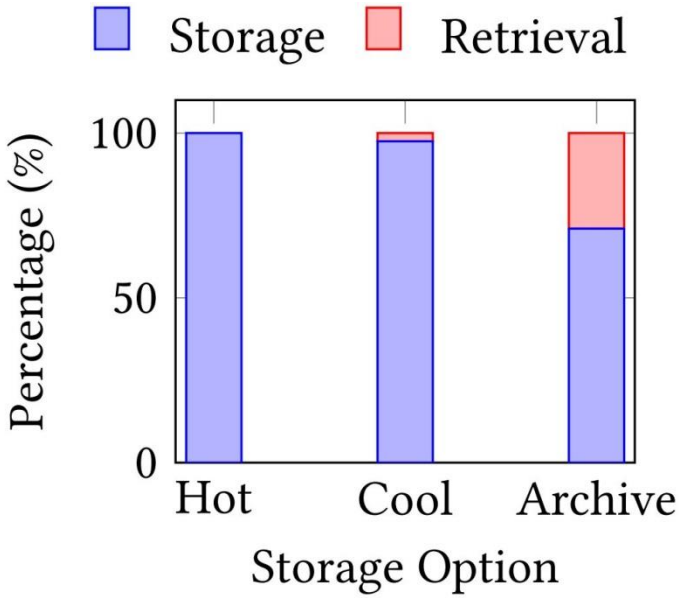

Deployment analysis: Storage–access trade-off in a cloud

Use case:

- 1 PB scientific archive
- Storage: 1 year
- Retrieval: all read once during that year

	Cost
Azure Archival Blob	$79K
Azure Cool Blob	$430K
Azure Hot Blob	$531K

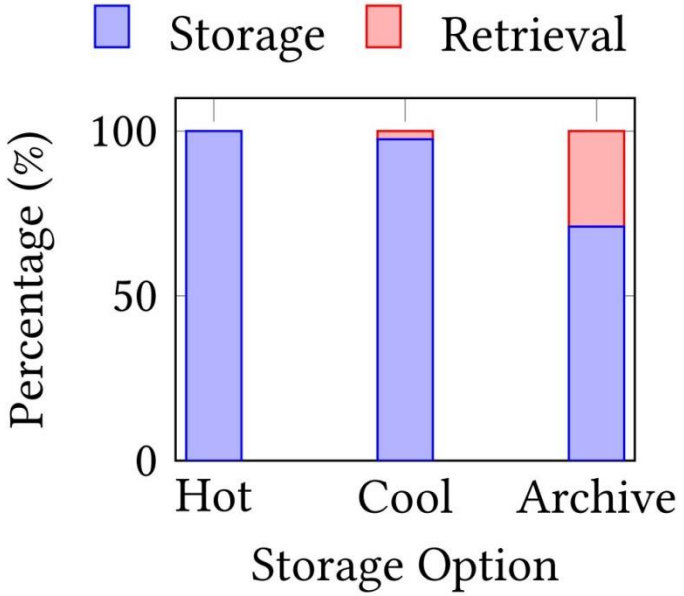

Archival blob store

- Data retrieval – 30% of the overall costs

Deployment analysis: Storage–access trade-off in a cloud

Use case:

- 1 PB scientific archive
- Storage: 1 year
- Retrieval: all read once during that year

	Cost
Azure Archival Blob	$79K
Azure Cool Blob	$430K
Azure Hot Blob	$531K

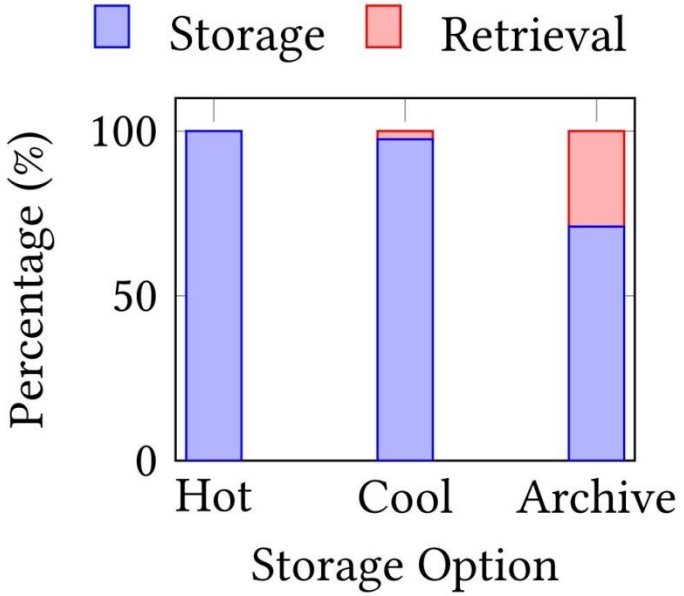

Archival blob store

- Data retrieval – 30% of the overall costs

Active archive migration

- storage is no longer the only concerning cost

Deployment analysis

Deployment analysis: Data scrubbing and vendor lock-in

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

- network utilization charges: storage vs compute nodes

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

- network utilization charges: storage vs compute nodes
- data scrubbing, if not built-in offer

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

- network utilization charges: storage vs compute nodes
- data scrubbing, if not built-in offer
- data migration, from one provider to another

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

- network utilization charges: storage vs compute nodes
- data scrubbing, if not built-in offer
- data migration, from one provider to another

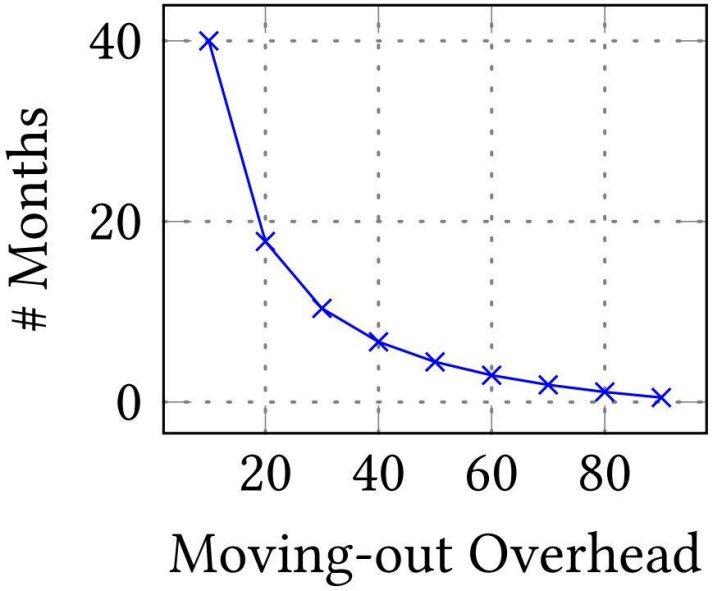

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

- network utilization charges: storage vs compute nodes
- data scrubbing, if not built-in offer
- data migration, from one provider to another

Migration costs (1PB)

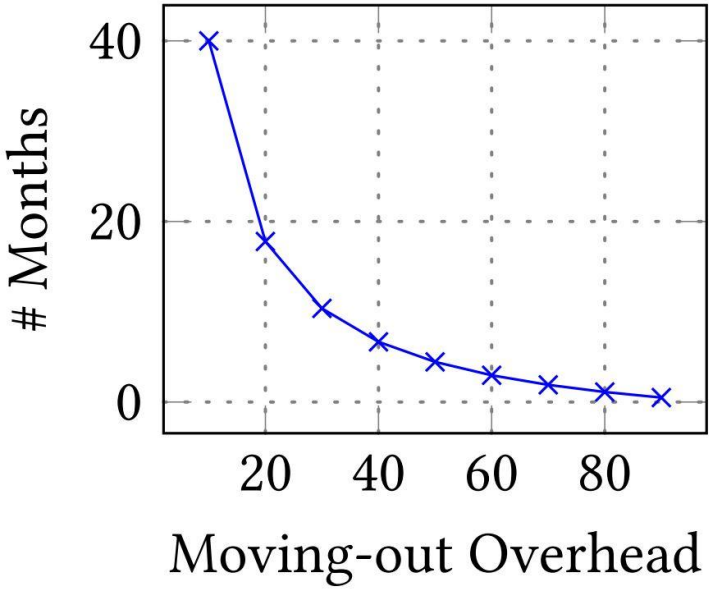

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

- network utilization charges: storage vs compute nodes
- data scrubbing, if not built-in offer
- data migration, from one provider to another

Migration costs (1PB)

- $23K (i.e. 5 months of additional storage)

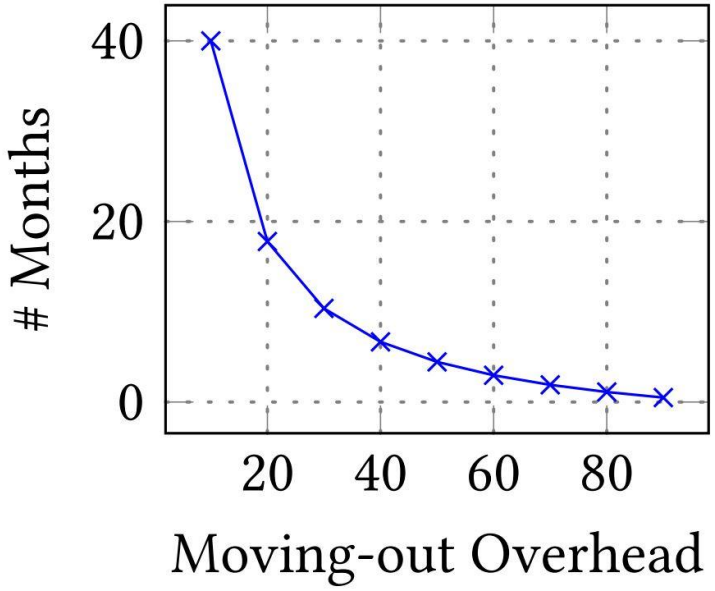

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

- network utilization charges: storage vs compute nodes
- data scrubbing, if not built-in offer
- data migration, from one provider to another

Migration costs (1PB)

- $23K (i.e. 5 months of additional storage)
- $75K (i.e. 16 months of additional storage), if migration charges included

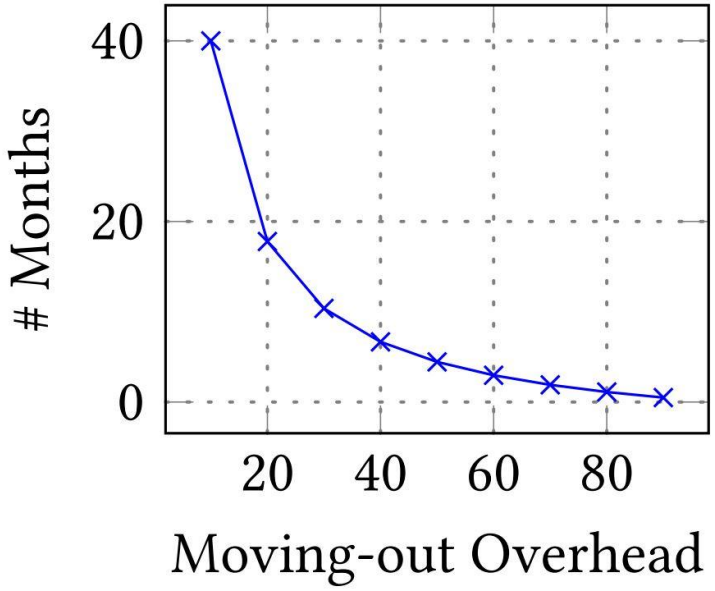

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

- network utilization charges: storage vs compute nodes
- data scrubbing, if not built-in offer
- data migration, from one provider to another

Migration costs (1PB)

- $23K (i.e. 5 months of additional storage)
- $75K (i.e. 16 months of additional storage), if migration charges included
- <10% moving-out overhead (>40 months storage)

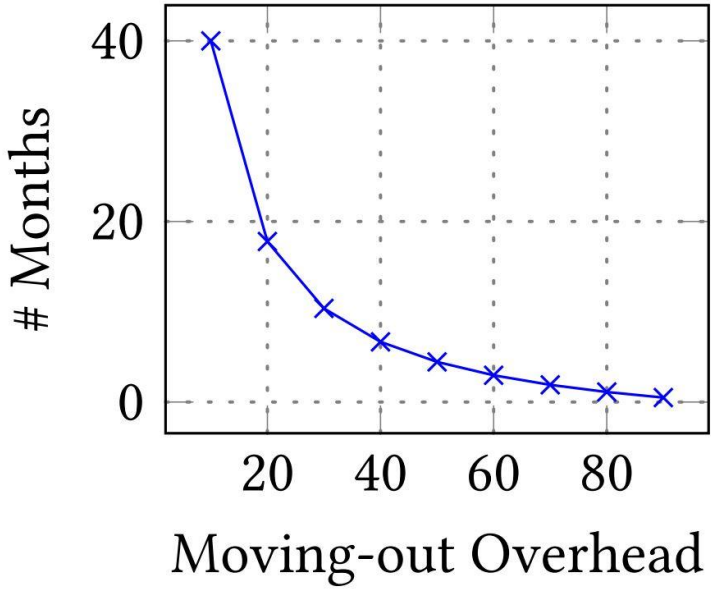

Deployment analysis: Data scrubbing and vendor lock-in

Cost added-values

- network utilization charges: storage vs compute nodes
- data scrubbing, if not built-in offer
- data migration, from one provider to another

Migration costs (1PB)

- $23K (i.e. 5 months of additional storage)
- $75K (i.e. 16 months of additional storage), if migration charges included
- <10% moving-out overhead (>40 months storage)

Once migrated, not easy to return *in-house*, or move to another provider

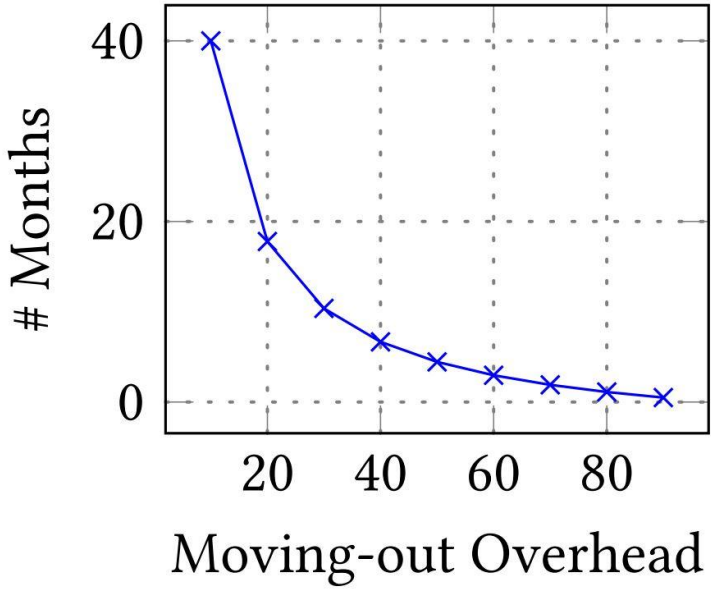

Deployment analysis

Deployment analysis: Tiered cold storage archive

Deployment analysis: Tiered cold storage archive

Private Infrastructure

Deployment analysis: Tiered cold storage archive

Private Infrastructure

Public Infrastructure

Deployment analysis: Tiered cold storage archive

Private Infrastructure

Google Cloud

Microsoft Azure

amazon web services

Public Infrastructure

Cache

+

Microsoft Azure

Google Cloud

amazon web services

Archive

Hybrid Infrastructure

Deployment analysis: Tiered cold storage archive

Private Infrastructure

Google Cloud

Microsoft Azure

amazon web services

Public Infrastructure

Cache

+

Microsoft Azure

Google Cloud

amazon web services

Archive

Hybrid Infrastructure

Two-tier, hybrid cloud infrastructure seems more appropriate

Deployment analysis: Tiered cold storage archive

Private Infrastructure

Google Cloud

Microsoft Azure

amazon web services

Public Infrastructure

Cache

+

Microsoft Azure

Google Cloud

amazon web services

Archive

Hybrid Infrastructure

Two-tier, hybrid cloud infrastructure seems more appropriate

- 1 copy locally, 1+ copy in cloud

Deployment analysis: Tiered cold storage archive

Private Infrastructure

Public Infrastructure

Cache

Archive

Hybrid Infrastructure

Two-tier, hybrid cloud infrastructure seems more appropriate

- 1 copy locally, 1+ copy in cloud
- eliminate data retrieval overheads

Deployment analysis: Tiered cold storage archive

Private Infrastructure

Public Infrastructure

Cache

Archive

Hybrid Infrastructure

Two-tier, hybrid cloud infrastructure seems more appropriate

- 1 copy locally, 1+ copy in cloud
- eliminate data retrieval overheads
- perform data scrubbing on local copy

Deployment analysis: Tiered cold storage archive

Private Infrastructure

Google Cloud

Microsoft Azure

amazon web services

Public Infrastructure

Cache

Archive

Hybrid Infrastructure

Two-tier, hybrid cloud infrastructure seems more appropriate

- 1 copy locally, 1+ copy in cloud
- eliminate data retrieval overheads
- perform data scrubbing on local copy
- solve vendor lock-in problem

Discussion

Discussion

Active Archive—Polystore

- active archives (and not just performance-sensitive analytics), an interesting candidate for adopting polystore architecture

Discussion

Active Archive—Polystore

- active archives (and not just performance-sensitive analytics), an interesting candidate for adopting polystore architecture

Provisioning & Configuring Data Archives

- for private storage infrastructures, important to having simulation, configuration tools, benchmarking, full-system monitoring tools

Discussion

Active Archive—Polystore

- active archives (and not just performance-sensitive analytics), an interesting candidate for adopting polystore architecture

A Cold Storage Benchmark

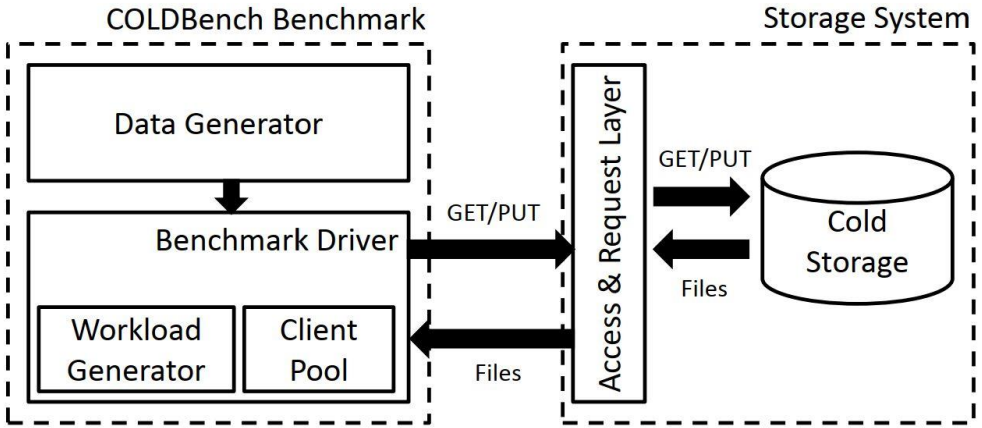

Provisioning & Configuring Data Archives

- for private storage infrastructures, important to having simulation, configuration tools, benchmarking, full-system monitoring tools

Discussion

Active Archive—Polystore

- active archives (and not just performance-sensitive analytics), an interesting candidate for adopting polystore architecture

A Cold Storage Benchmark

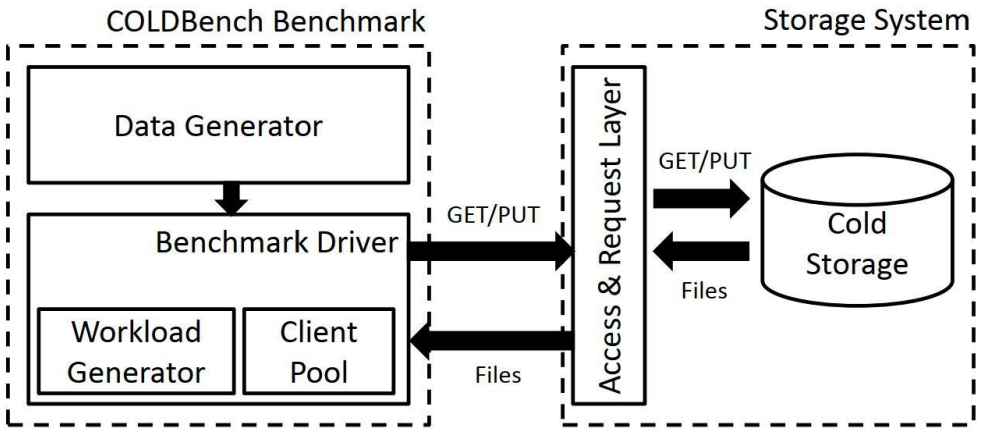

Provisioning & Configuring Data Archives

- for private storage infrastructures, important to having simulation, configuration tools, benchmarking, full-system monitoring tools

Archive Profiling & Monitoring

- current methodologies, inappropriate and time-consuming, due to complexity and privacy

To summarize

Cold storage data archives, a fundamental building block for scientific application domains

Data archive analysis of three application domains

- weather forecast (ECFS, MARS), earth observation (D-SDA), radio astronomy (LOFAR)

A hybrid two-tier approach of a private/public cold storage infrastructure, most promising for a reasonable cost/performance trade-off

Several areas need further exploration: scientific data archival – not just a storage problem

Contact

Bunjamin Memishi
German Aerospace Center
Institute of Data Science, Jena
Data Management Technologies Group
bunjamin.memishi@dlr.de

And (or)

- Raja Appuswamy (raja.appuswamy@eurecom.fr)
- Marcus Paradies (marcus.paradies@dlr.de)