Institut de Recherche
en Informatique de Toulouse
CNRS - INP - UT3 - UT1 - UT2J

Université
Fédérale
Toulouse
Midi-Pyrénées

Comparing energy-aware vs. cost-aware data replication strategy

Morgan Séguéla Riad Mokadem Jean-Marc Pierson

Journées Cloud 2019

Toulouse Institute of Computer Science Research, Toulouse University

Plan de la présentation

Context

Increase of data volume

ex: data traffic forecast for Cloud

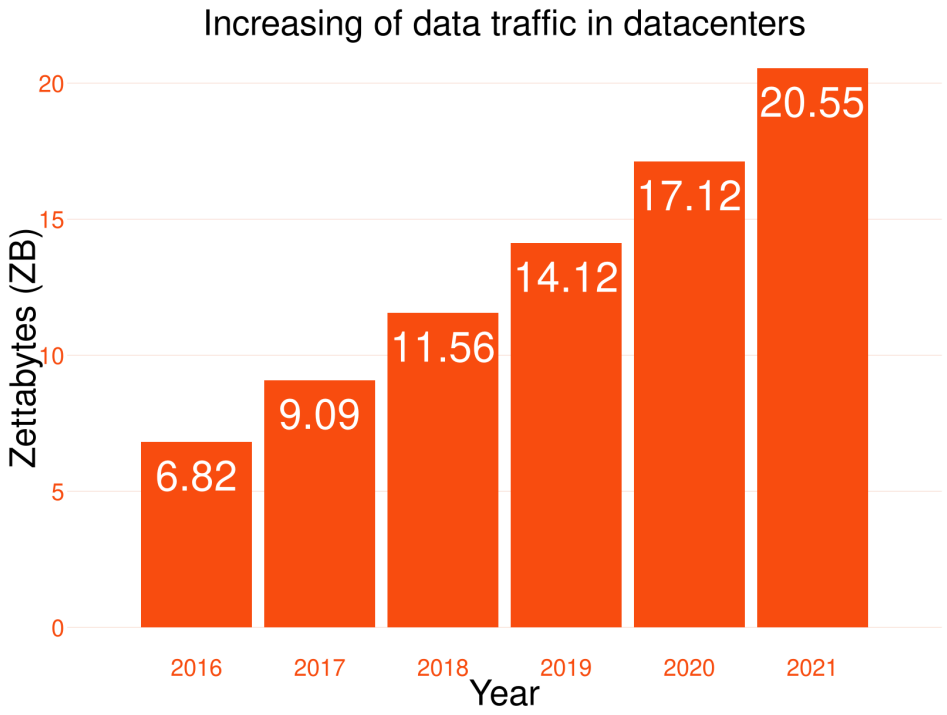

source: cisco, 2016

Context

Increase of data volume

ex: data traffic forecast for Cloud

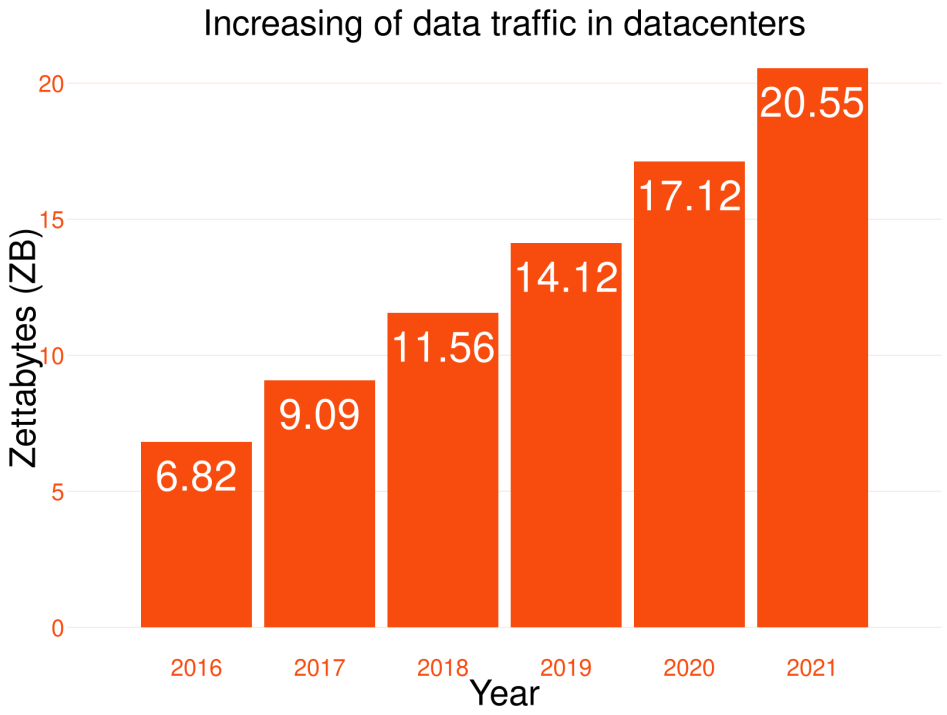

source: cisco, 2016

Needs

- Availability
- Performance
- Processing Capacity

Context

Increase of data volume

ex: data traffic forecast for Cloud

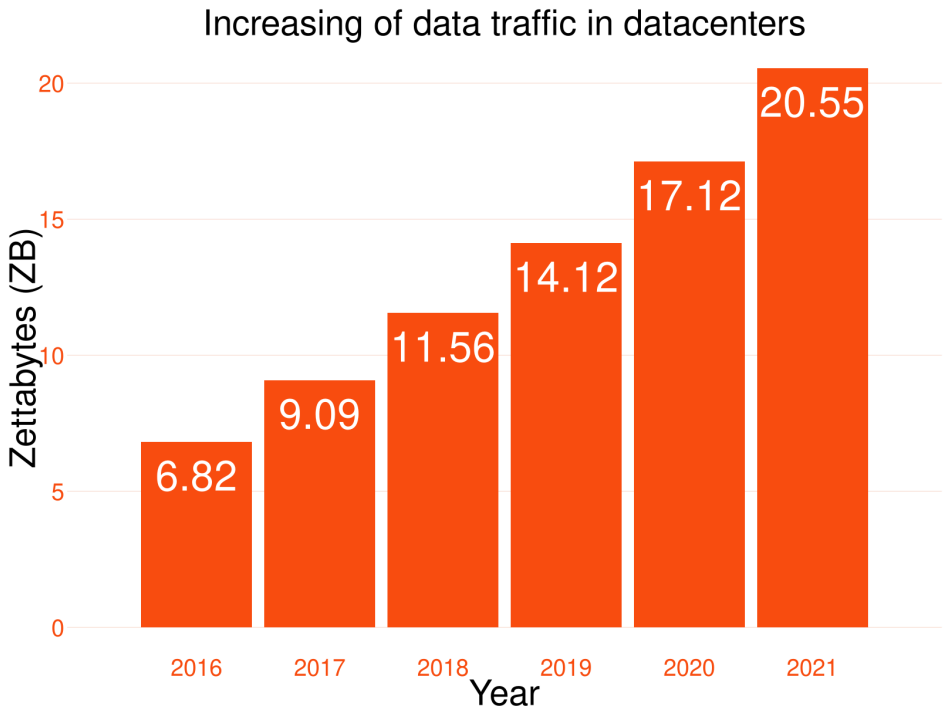

source: cisco, 2016

Needs

- Availability
- Performance

  → Data replication
- Processing Capacity

  → Cloud

Cloud Characteristics

Resources

Seemingly Unlimited

Cloud Characteristics

Resources

Seemingly Unlimited

Elasticity

Resources are automatically adapted

Available $\rightarrow$ Needed

Cloud Characteristics

Resources

Seemingly Unlimited

Elasticity

Resources are automatically adapted

Available $\rightarrow$ Needed

Pay-As-You-Go

Users pay what they consumed

Cloud Characteristics

Resources

Seemingly Unlimited

Elasticity

Resources are automatically adapted

Available $\rightarrow$ Needed

Pay-As-You-Go

Users pay what they consumed

Service Level Agreement

Contract

Provider and tenant

Cloud Characteristics

Resources

Seemingly Unlimited

Elasticity

Resources are automatically adapted

Available → Needed

Pay-As-You-Go

Users pay what they consumed

Service Level Agreement

Contract

Provider and tenant

Content

- Price
- Service Level Objective
- Penalties: SLOs not satisfied

Data Replication

Technique

Needs

- Availability
- Performance
- Fault tolerance

Data Replication

Technique

Needs

- Availability
- Performance
- Fault tolerance

Objectives

Beside answering to needs

- Profit
- Reduce energy consumption
- Reduce carbon footprint

State of the Art

Strategies	Availability	Perf.	Profit	Energy
[Long et al., 2014]	✓	✓	✗	✓
[Boru et al., 2015]	✗	✓	✗	✓
[Zhang et al., 2015]	✓	✓	✗	✓
[Xu et al., 2015]	✓	✗	✗	✓
[Li et al., 2011]	✓	✗	✓	✗
[Gill and Singh, 2016]	✓	✓	✓	✗
[Zeng et al., 2017]	✓	✓	✓	✗
[Tos et al., 2017]	✓	✓	✓	✗
[Alghamdi et al., 2017]	✓	✓	✓	✓

Table: Data replication strategies features for cloud

Issue

There are only few strategies that takes into account the provider's profit and the energy consumption simultaneously.

Provider's profit model

Proposed model in [Tos et al., 2017]:

Profit

$$Profit = Revenues - Expenditure$$

Provider's profit model

Proposed model in [Tos et al., 2017]:

Profit

$$Profit = Revenues - Expenditure$$

Revenues

$$Revenues = \sum_{P=0}^{N_P} Price_{Location}(Instance)$$

Provider's profit model

Proposed model in [Tos et al., 2017]:

Profit

$$Profit = Revenues - Expenditure$$

Revenues

$$Revenues = \sum_{P=0}^{N_P} Price_{Location}(Instance)$$

Expenditures

$$Expenditures = C_{Computation} + C_{Network} + C_{Storage} + C_{Penalties} + C_{Others}$$

Provider's profit model

Expenditure details

- $\mathbf{C_{Computation}}$: Execution cost of one task
- $\mathbf{C_{Network}}$: Data transfer cost
- $\mathbf{C_{Storage}}$: Data storage cost
- $\mathbf{C_{Penalties}}$: SLO violations penalty cost
- $\mathbf{C_{Others}}$: Any other costs (salary, energy, ...)

Energy consumption model

In order to compare different strategies $\rightarrow$ extend PEPR strategy to estimate the energy consumption

Energy consumption model

In order to compare different strategies → extend PEPR strategy to estimate the energy consumption

Energy consumption

$$E_{Tot} = E_{Idle} + E_{Dynamic}$$

Energy consumption model

In order to compare different strategies → extend PEPR strategy to estimate the energy consumption

Energy consumption

$$E_{Tot} = E_{Idle} + E_{Dynamic}$$

Idle

$$E_{Idle} = (P_{CPU_{idle}} + P_{RAM_{idle}} + P_{HDD_{idle}} + P_{NIC_{idle}}) * T_{\text{tot}}$$

Energy consumption model

Task execution

$$E_{Pro} = E_{CPU_{Act}} + E_{RAM_{read}}$$

Energy consumption model

Task execution

$$E_{Pro} = E_{CPU_{Act}} + E_{RAM_{read}}$$

File reading

$$E_{FT} = E^{src}_{HDD_{read}} + E^{tgt}_{RAM_{wr}} + \mathbb{1}_{src \neq tgt} * (E^{src}_{RAM_{wr}} + E^{src}_{RAM_{read}} + E^{src}_{NIC} + E_{Network} + E^{tgt}_{NIC})$$

Energy consumption model

Task execution

$$E_{Pro} = E_{CPU_{Act}} + E_{RAM_{read}}$$

File reading

$$E_{FT} = E^{src}_{HDD_{read}} + E^{tgt}_{RAM_{wr}} + \mathbb{1}_{src \neq tgt} * (E^{src}_{RAM_{wr}} + E^{src}_{RAM_{read}} + E^{src}_{NIC} + E_{Network} + E^{tgt}_{NIC})$$

Data replication

$$E_{Repl} = E^{src}_{RAM_{read}} + E^{tgt}_{HDD_{write}} + \mathbb{1}_{src \neq tgt} * (E^{src}_{NIC} + E_{Network} + E^{tgt}_{NIC} + E^{tgt}_{RAM_{write}} + E^{tgt}_{RAM_{read}})$$

Setup

Simulator

Cloudsim → Extended : Data replication and profit [Thèse Tos]
→ Extended : Energy consumption

Setup

Simulator

Cloudsim → Extended : Data replication and profit [Thèse Tos]
→ Extended : Energy consumption

Architecture parameters

[Tos et al., 2017]

Power parameters

CloudMip platform from IRIT

Economic parameters

Google cloud pricing - 04/08/2019

Experiments available

Github : https://github.com/MorganSeguela/IGSC_2019_XPS
Docker : https://hub.docker.com/r/mseguela/igsc-2019-xps

PEPR - [Tos et al., 2017]

Objectives

1. Performance constraints
2. Taking profit into account

Architecture

Peer to Peer

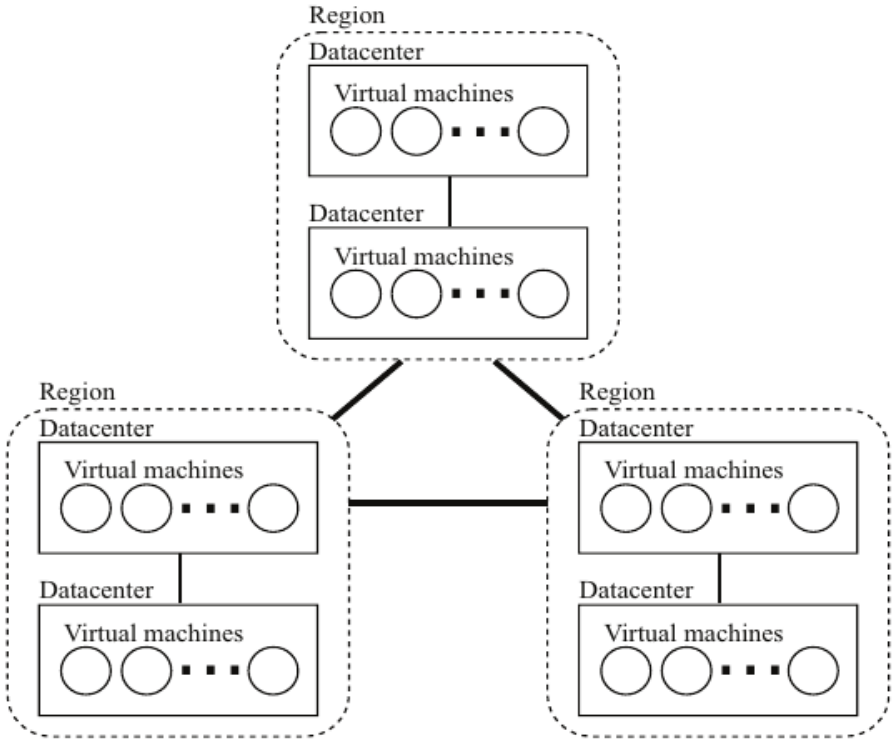

PEPR - [Tos et al., 2017]

Objectives

1. Performance constraints
2. Taking profit into account

Replication

Trigger:
Response time > Threshold

Decision:
Profit > 0

Architecture

Peer to Peer

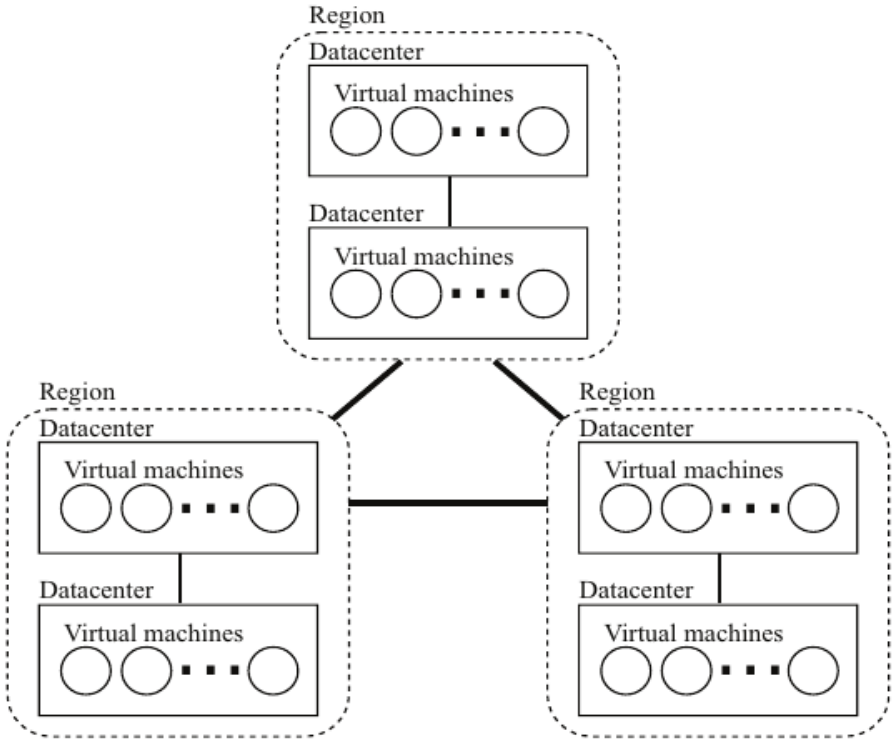

PEPR - [Tos et al., 2017]

Objectives

1. Performance constraints
2. Taking profit into account

Replication

Trigger:
Response time > Threshold

Decision:
Profit > 0

Replica management
Remove replicas with the latest access

Architecture

Peer to Peer

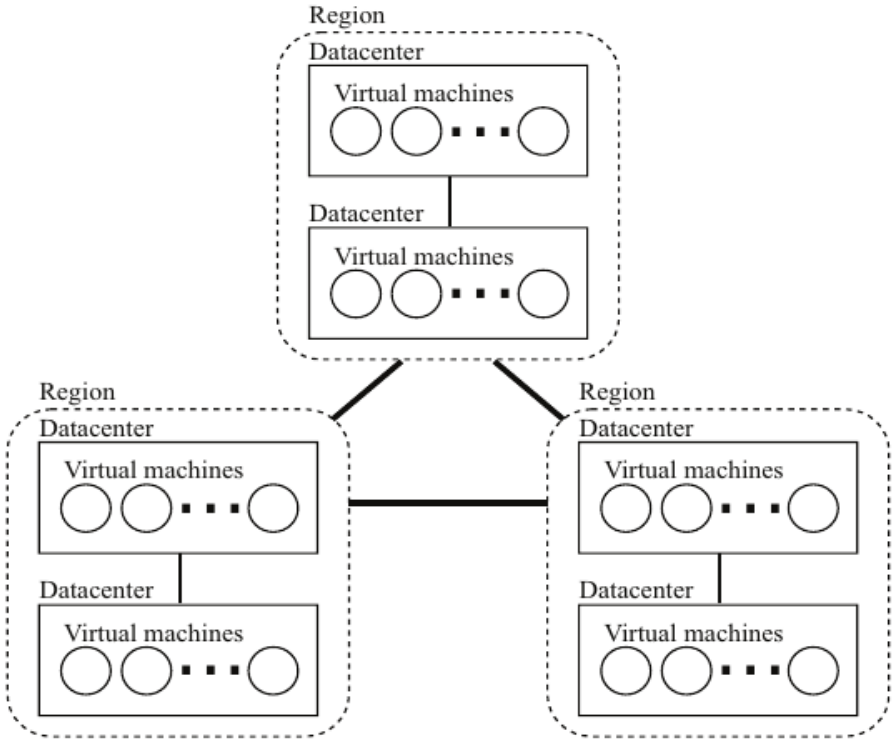

[Boru et al., 2015]

Objectives

- Minimizing energy consumption
- Minimizing bandwitdth usage

Architecture

Three Tier Fat Tree

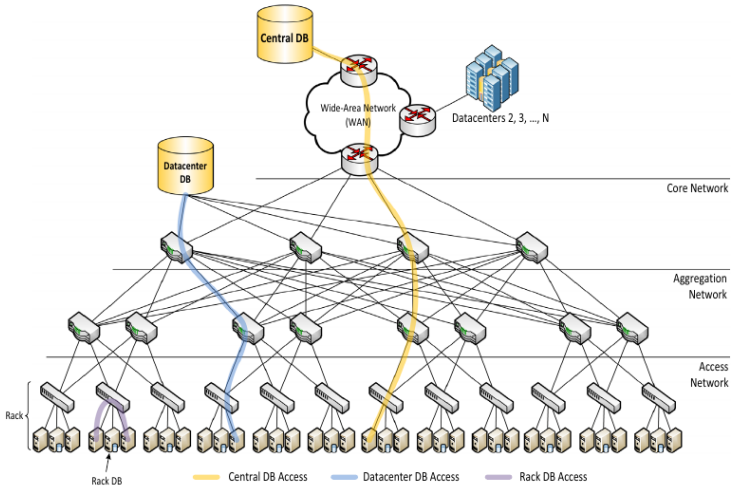

[Boru et al., 2015]

Objectives

- Minimizing energy consumption
- Minimizing bandwitdth usage

Replication

Trigger:
Requests number > Threshold

Decision:
Energy and Bandwidth consumption lower in the level below

Architecture

Three Tier Fat Tree

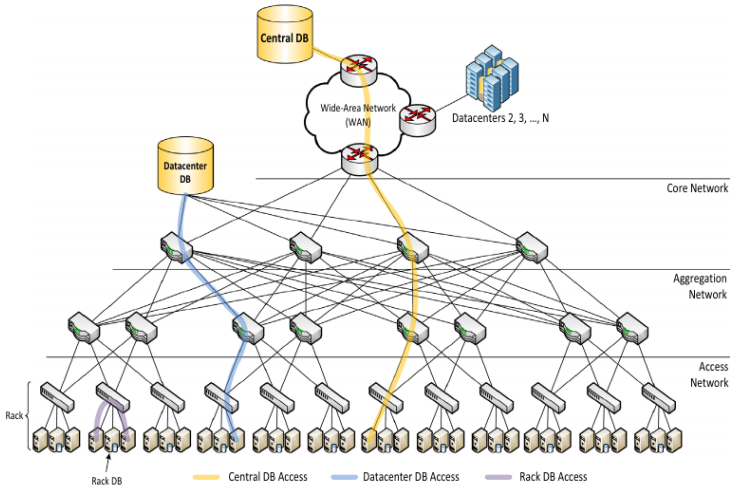

MORM - [Long et al., 2014]

Multi-objectives Optimization
Replication management

- Reduce file unavailability
- Reduce service time
- Reduce latency
- Reduce load variance
- Reduce energy consumption

Static strategy

- Replication done before the execution
- Do not adapt to the ongoing execution

Architecture

Peer to Peer

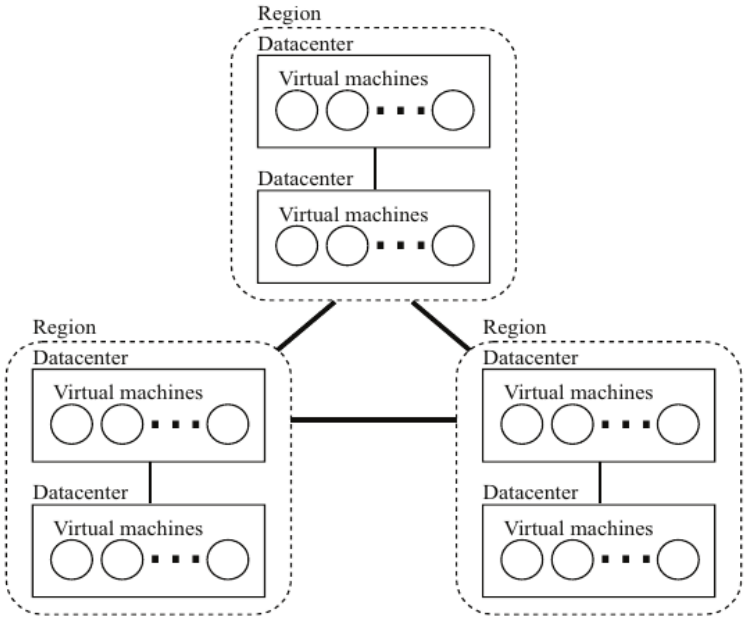

No Replication

Peer to Peer Architecture (P2PNorep)

Uniformly randomized distribution of data between nodes

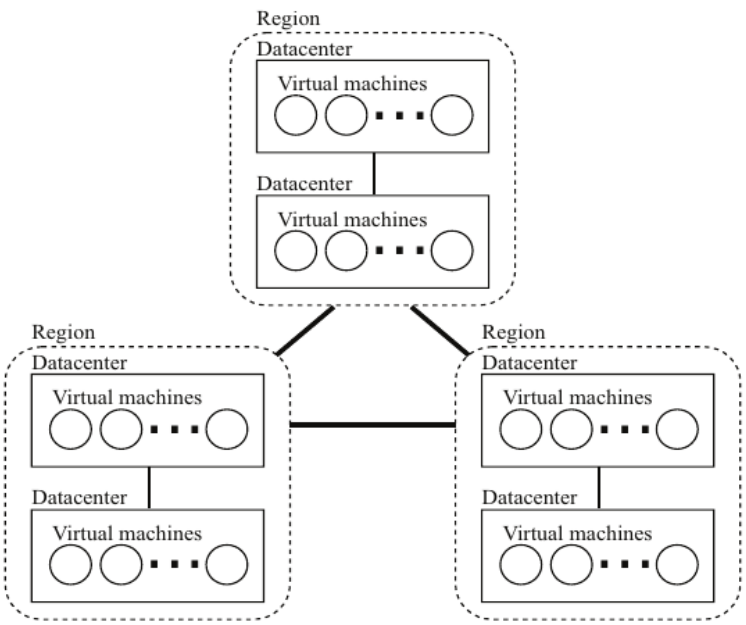

Three Tier Fat Tree Architecture (HNorep)

Data stored in the Central DB

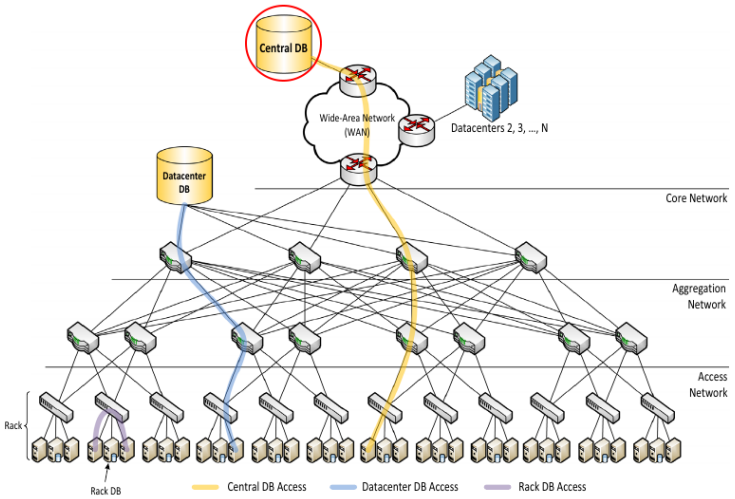

Results

Replications during the experiment

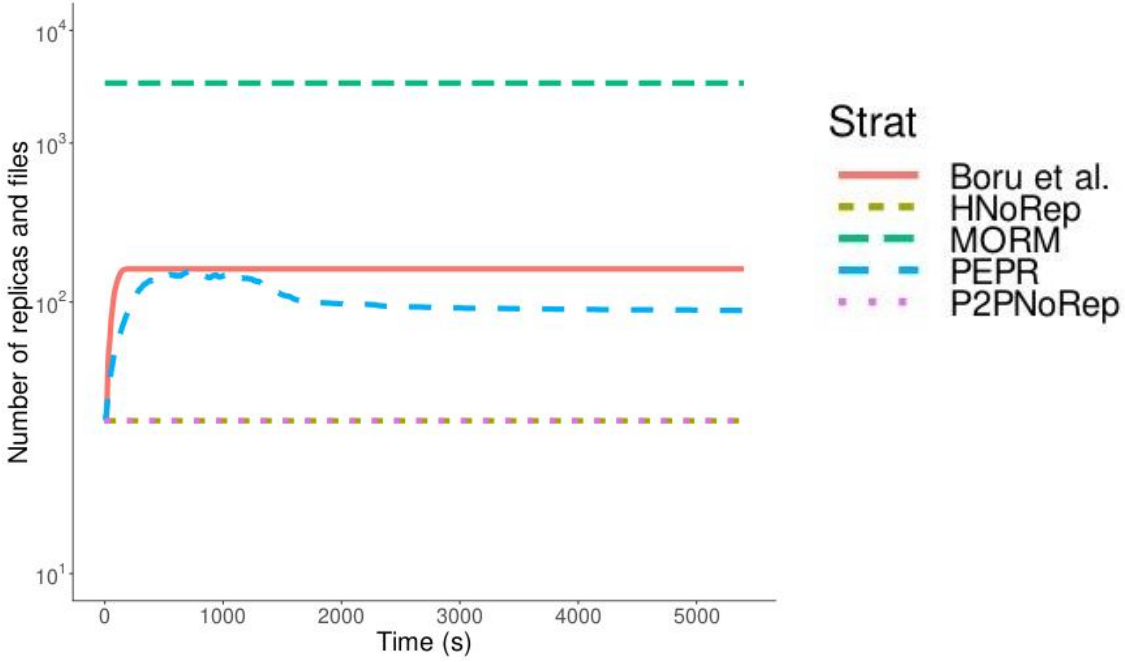

Violations during the experiment

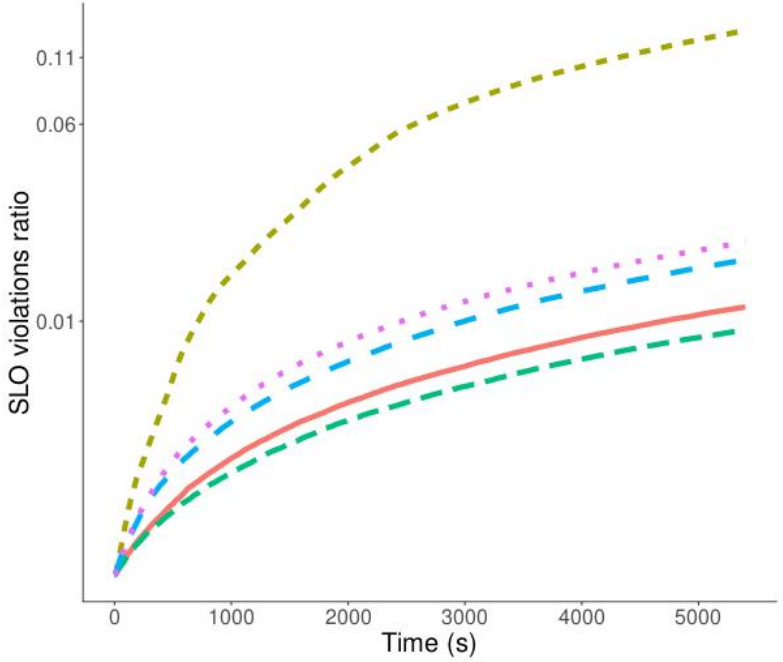

Results

Energy consumed

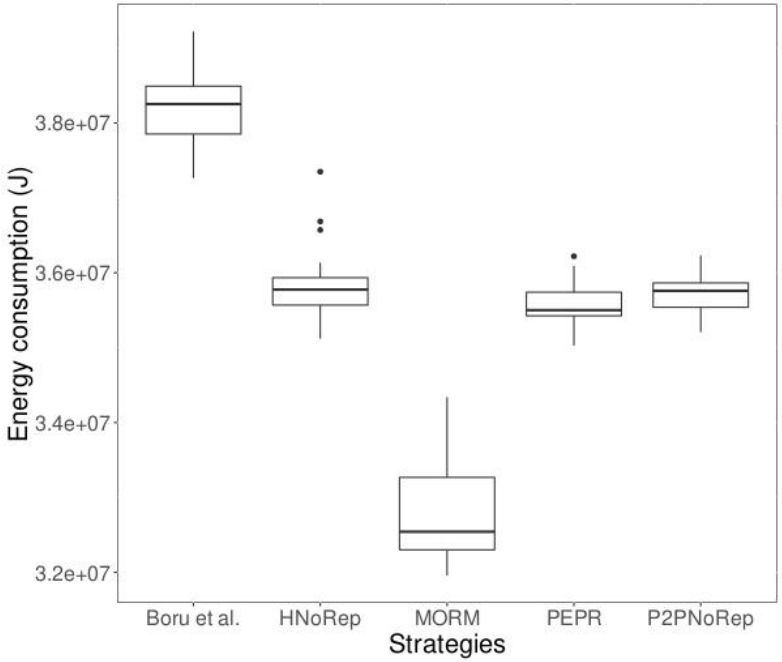

Results

Energy consumed

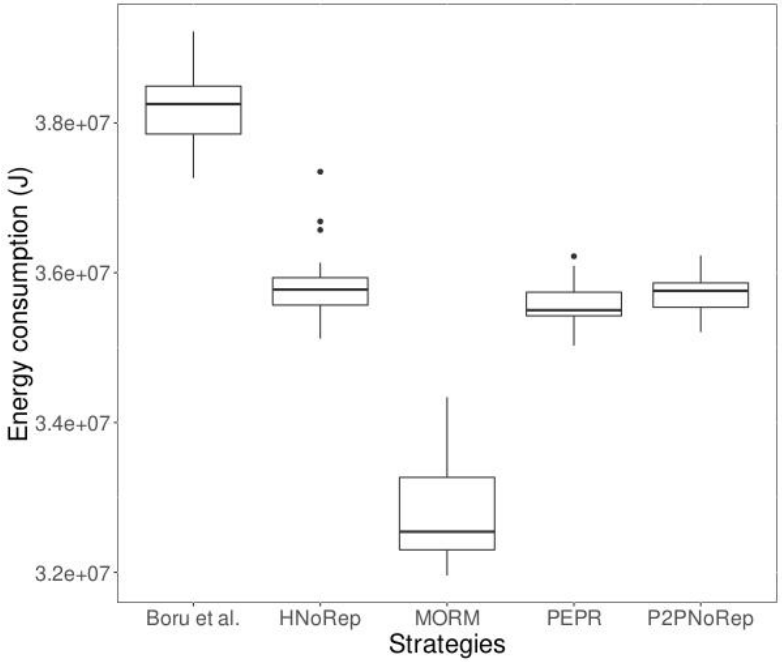

Analyses

- Boru et al. : Overhead made by replication

Results

Energy consumed

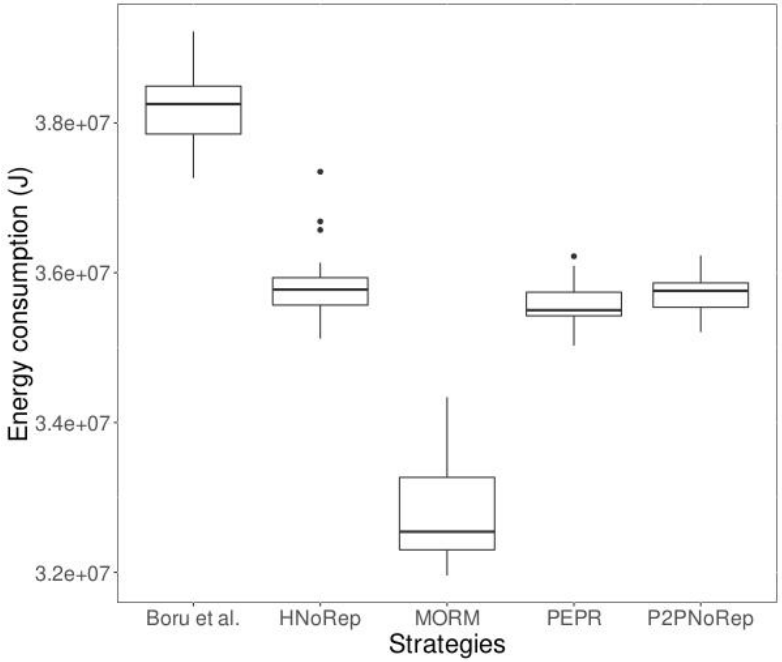

Analyses

- Boru et al. : Overhead made by replication
- PEPR : Significantly lower from the lack of replication

Results

Energy consumed

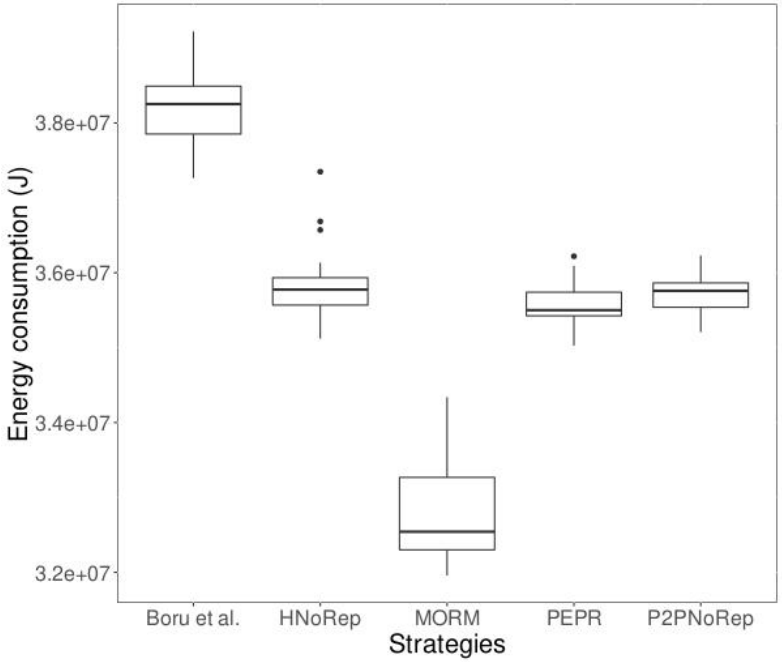

Analyses

- Boru et al. : Overhead made by replication
- PEPR : Significantly lower from the lack of replication
- MORM : Replication already done before the experience and made a lot of replicas that reduce the transfer energy consumption

Results

Total expenditure

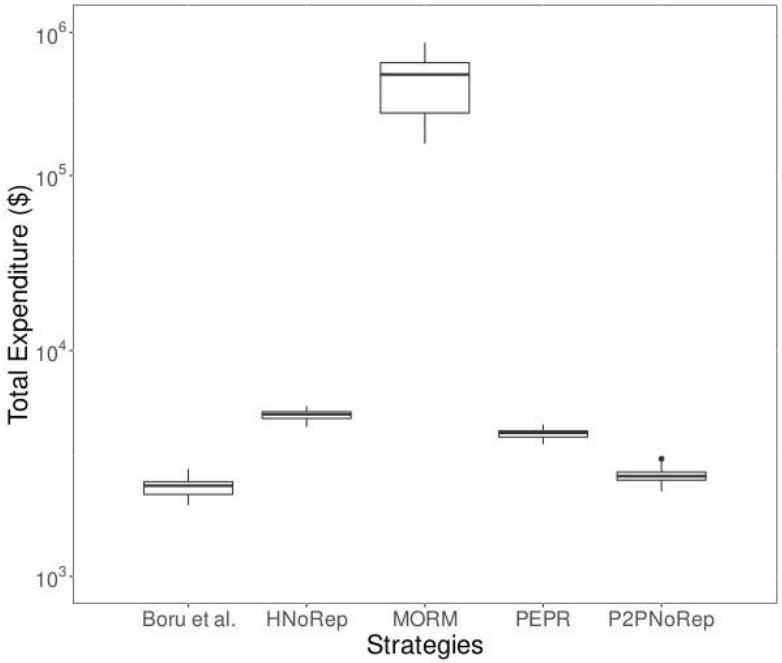

Results

Total expenditure

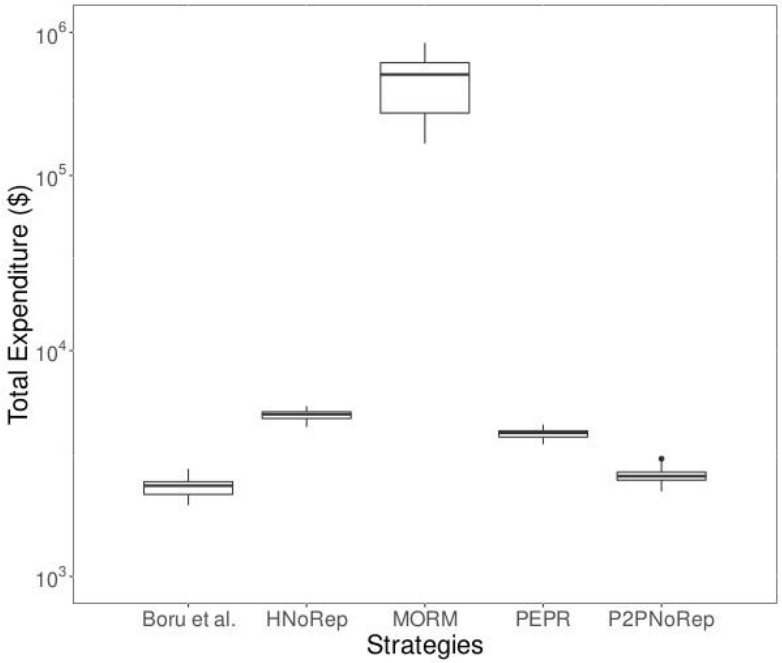

Explanation

- Boru et al. : Replicates in regions first

  $\rightarrow$ Reduction of transfer cost

Results

Total expenditure

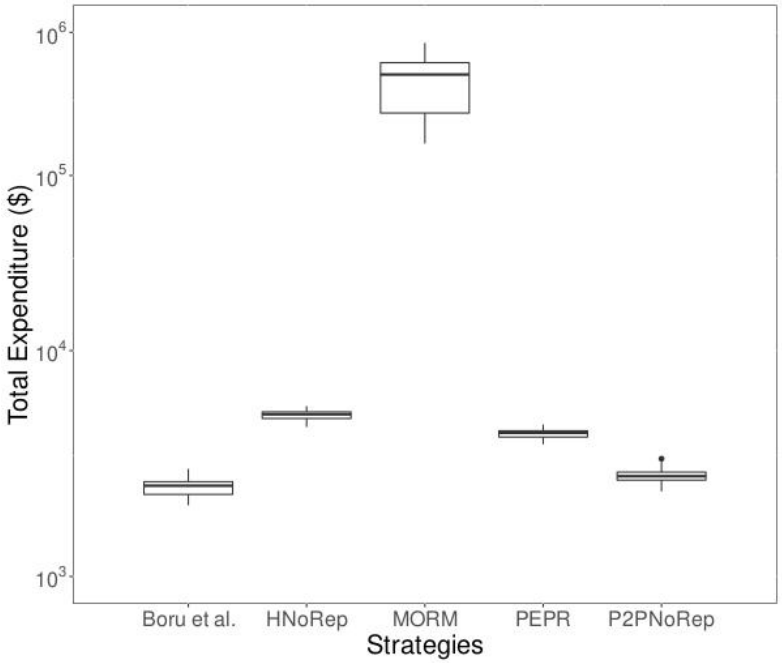

Explanation

- Boru et al. : Replicates in regions first

  $\rightarrow$ Reduction of transfer cost
- PEPR : Replicates closer to the processing node and removal mechanism

  $\rightarrow$ More transfer between regions

Results

Total expenditure

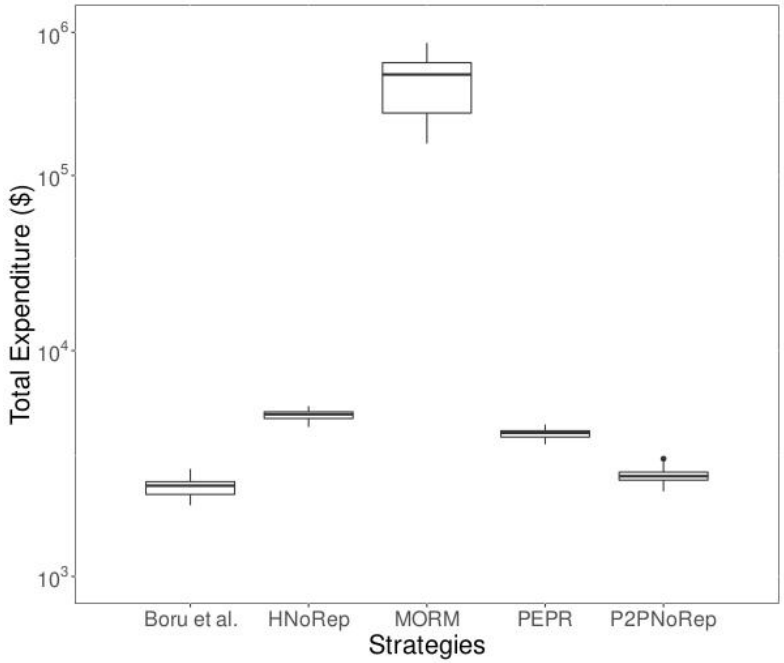

Explanation

- Boru et al. : Replicates in regions first

  → Reduction of transfer cost
- PEPR : Replicates closer to the processing node and removal mechanism

  → More transfer between regions
- MORM : Lot of replicas

  → Less transfers but high storage usage

Conclusion and future works

Conclusion

BORU : Low Expenditure but High Energy consumption
MORM : High Expenditure but Low Energy consumption
PEPR : Bit High Expenditure and bit Low Energy consumption

Guidelines :

- Replication during the execution
- Replica management mechanism

Conclusion and future works

Conclusion

BORU : Low Expenditure but High Energy consumption
MORM : High Expenditure but Low Energy consumption
PEPR : Bit High Expenditure and bit Low Energy consumption

Guidelines :

- Replication during the execution
- Replica management mechanism

Conclusion and future works

Conclusion

BORU : Low Expenditure but High Energy consumption
MORM : High Expenditure but Low Energy consumption
PEPR : Bit High Expenditure and bit Low Energy consumption

Guidelines :

- Replication during the execution
- Replica management mechanism

Future works

Proposing a data replication strategy that takes into account provider's profit and energy consumption

Thank you for your attention !

Any questions ?

Alghamdi, M., Tang, B., and Chen, Y. (2017).
Profit-based file replication in data intensive cloud data centers.
In *2017 IEEE International Conference on Communications (ICC)*, pages 1–7, Paris, France. IEEE.

Boru, D., Kliazovich, D., Granelli, F., Bouvry, P., and Zomaya, A. Y. (2015).
Energy-efficient data replication in cloud computing datacenters.
Cluster Computing, 18(1):385–402.

Gill, N. K. and Singh, S. (2016).
A dynamic, cost-aware, optimized data replication strategy for heterogeneous cloud data centers.
Future Generation Computer Systems, 65:10–32.

Li, W., Yang, Y., and Yuan, D. (2011).
A Novel Cost-Effective Dynamic Data Replication Strategy for Reliability in Cloud Data Centres.
In *2011 IEEE Ninth International Conference on Dependable, Autonomic and Secure Computing*, pages 496–502, Sydney, Australia. IEEE.

Long, S.-Q., Zhao, Y.-L., and Chen, W. (2014).
MORM: A Multi-objective Optimized Replication Management strategy for cloud storage cluster.
Journal of Systems Architecture, 60(2):234–244.

Tos, U., Mokadem, R., Hameurlain, A., Ayav, T., and Bora, S. (2017).
Ensuring performance and provider profit through data replication in cloud systems.
Cluster Computing.

Xu, Z., Deng, N., Stewart, C., and Wang, X. (2015).
CADRE: Carbon-Aware Data Replication for Geo-Diverse Services.
In *2015 IEEE International Conference on Autonomic Computing*, pages 177–186, Grenoble, France. IEEE.

Zeng, L., Xu, S., Wang, Y., Kent, K. B., Bremner, D., and Xu, C. (2017).
Toward cost-effective replica placements in cloud storage systems with QoS-awareness: Toward cost-effective replica placements in cloud storage systems with QoS-awareness.
Software: Practice and Experience, 47(6):813–829.

Zhang, L., Deng, Y., Zhu, W., Zhou, J., and Wang, F. (2015).
Skewly replicating hot data to construct a power-efficient storage cluster.
Journal of Network and Computer Applications, 50:168–179.