Exertion-based billing for cloud storage access

Matthew Wachs*

Lianghong Xu*, Arkady Kanevsky†, Greg Ganger*

*PARALLEL DATA LABORATORY
Carnegie Mellon University

†VMWARE

Cloud accounting

- Infrastructure-as-a-Service (IaaS)
  - Provider makes resources available to clients
  - Clients pay for resources used

- Provider wants to recover costs
- Client wants to pay fairly for use

- This talk: Focus on storage

Provider: Recovering costs

- More bytes stored → more disks needed
  - Bill for capacity
  - ✓ Providers do this

- More time spent on requests → more disks needed
  - Bill for access
  - ✘ Not all providers do this

    vCloud Express
  - ✘ Those that do, use the wrong metric

Billing for access

- Providers currently bill for:

IOs	Amazon EBS, S3, Windows Azure
Bytes transferred	Amazon S3
Performance	Amazon EC2

✘ None matches time used × cost of resources

Unsustainable approach: Fixed cost per IO

Suppose provider charges $0.01 per IO
but it takes $0.01/ms to operate a disk

Request scenario	Disk time	Real cost	Bill
Cache miss, sequential	1 ms	$0.01	$0.01
Cache miss, long seek	20 ms	$0.20	$0.01
Cache hit	0 ms	$0.00	$0.01

Provider loses money

Client pays too much

Alternatives

- Charge per byte or performance
  - ✘ Both also vary as in previous example
- Charge for worst case ($0.20/IO instead of $0.01)
  - ✘ Clients with “easy” requests pay too much
- Charge for the average case
  - ✘ Some clients subsidize others
  - ✘ Will the average change over time?

Solution: Charge for disk time

Suppose it takes $0.01/ms to operate a disk and the provider charges $0.01/ms

Request scenario	Disk time	Real cost	Bill
Cache miss, sequential	1 ms	$0.01	$0.01
Cache miss, long seek	20 ms	$0.20	$0.20
Cache hit	0 ms	$0.00	$0.00

✓ Costs recovered fairly

But, a technical problem remains…

- Resources may be shared across tenants
- Workload A may be affected by Workload B
  - Workload A sequential, Workload B disrupts locality
  - Workload A cacheable, Workload B evicts A’s pages
  - ✘ Workload B drives up Workload A’s bill

Interference example

Workload	Standalone exertion	Exertion w/ interference
Sequential 1 MB/s	1.6%	23%
Random 1 MB/s	67%	74%

✘ Random workload increases Sequential workload’s exertion

How to solve?

- Bill Workload B for its impact on Workload A?
  - ✘ Not the “fault” of B that it coexists poorly with A

- Provider absorbs the cost?
  - ✘ Provider will either lose money or pass it along as a hidden cost to customers

✓ Avoid the interference in the first place

Performance insulation is needed

- *Performance insulation*: System property
  - Another workload is not able to affect your…
    - Efficiency
    - Exertion
    - Performance (in your share of time)
  - … beyond a small (e.g. 10%) factor

- Strictly limits transient influences

How to insulate storage

- Preserve locality
  - Disk-head timeslicing
    - Seek between workloads infrequently

- Provide predictable cache allocation
  - Cache partitioning
    - Partition sizes based on access patterns

- Argon storage system [FAST 2007]

Insulation example

Insulation limits impact of other workload to < 10%

Workload	Standalone exertion	Exertion w/ insulation
Sequential 1 MB/s	1.6%	≤ 1.8%
Random 1 MB/s	67%	≤ 75%

✓ Each workload's exertion is close to ideal because it receives dedicated disk time & cache space

Summary

- Clients should pay for resources used
- Storage: both capacity and access cost money
- Disk time is what costs, should be the metric
- Bills should be independent and predictable
- Performance insulation is needed to make it fair