

End-to-end Performance Isolation through Virtual Datacenters

Mateusz Piotrowski

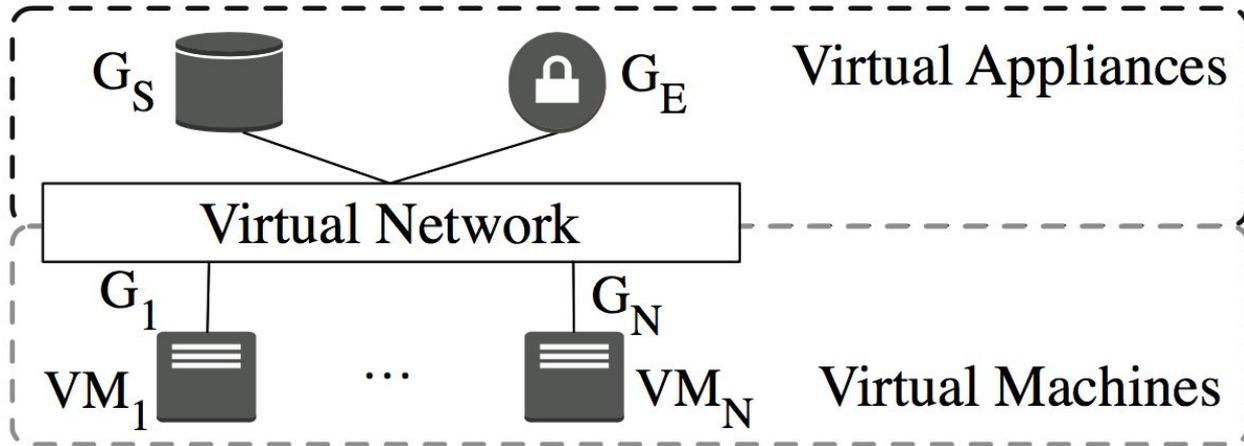
Problems

- Tenants move from enterprise datacenters to **cloud datacenters**.
- Cloud datacenters provide **services** like storage and load balancers with **appliances**.
- Tenants' **applications are not isolated** from each other so they **impact each other's performance**.
- The standard ways to measure performance like **requests/second** or **bits/second** are **inadequate**.

Virtual datacenters

Virtual datacenters (VDC)

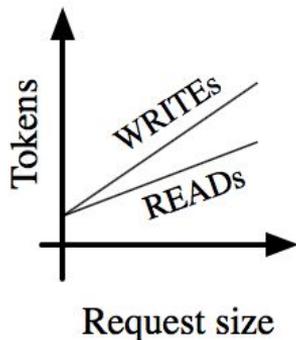
- Guarantees can be fixed or minimal.
 - **Fixed** guarantee gives the tenant predictable performance. It **does not change** over time.
 - **Minimal** guarantee allows tenants to elastically obtain **additional resources when needed**.



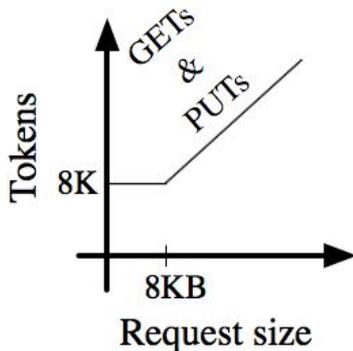
Virtual cost function

Use tokens/sec instead of requests/sec

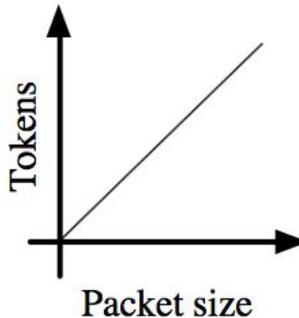
- Simple way to introduce **allocation policies**.
- Unified throughput metric based on virtual request cost functions.
- Bottlenecks change over time and space.



(a) Filestore



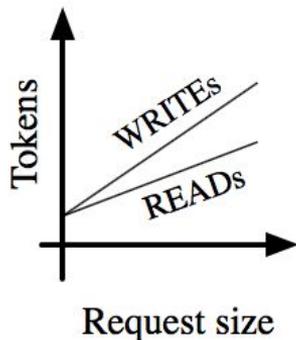
(b) Key-value store



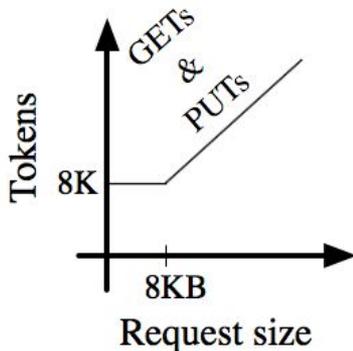
(c) Network

Virtual cost function

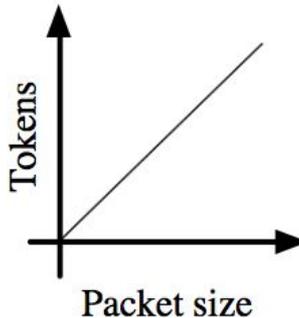
- Provider must determine a function for each appliance. It is typically **based on benchmarks and historical data**.
- **Doesn't have to be perfect!** Sometimes the same request needs a different amount of work.



(a) Filestore



(b) Key-value store

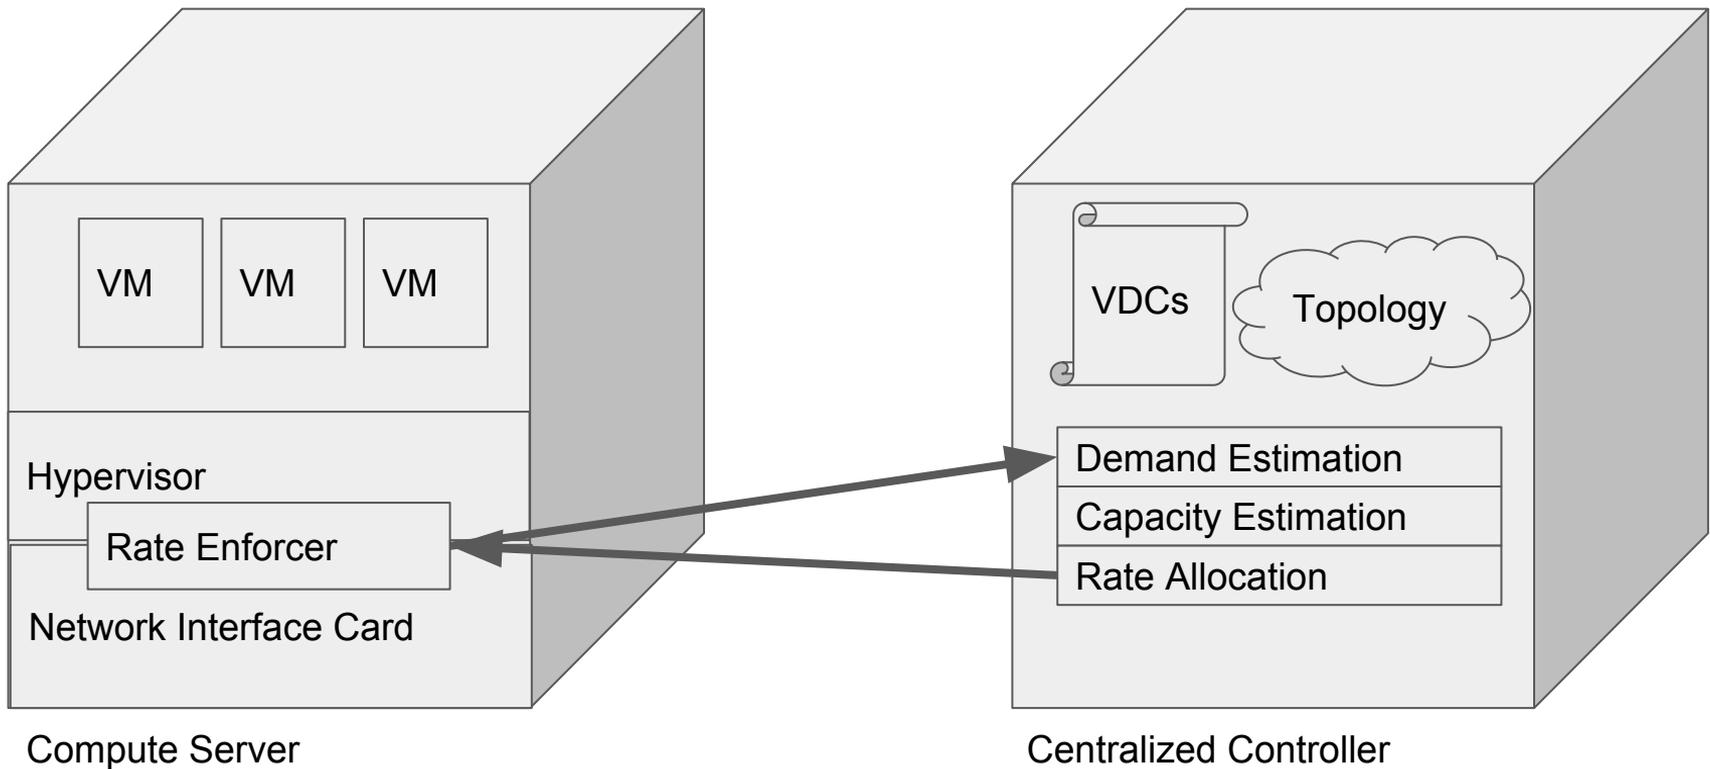


(c) Network

Pulsar

Advantages of Pulsar

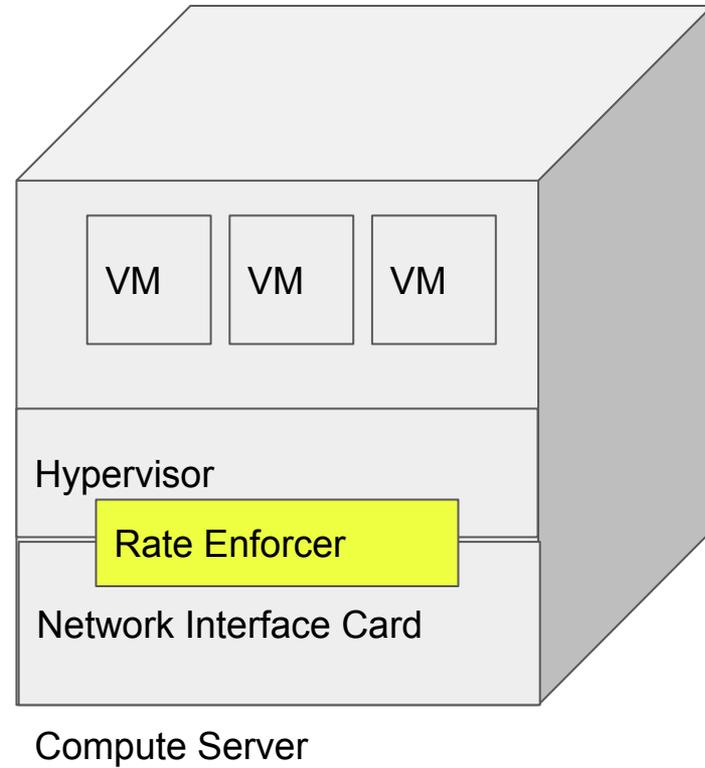
- Easy to deploy.
- Doesn't require any modifications to applications, guest OSES, appliances and the network.



Pulsar

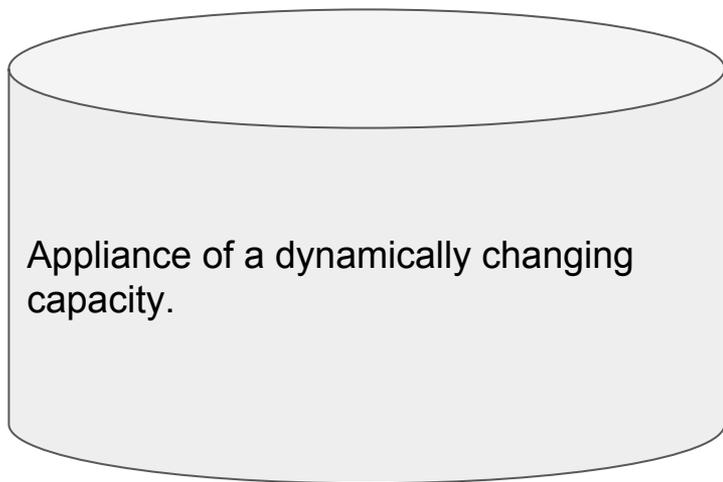
Rate Enforcer

- Collect local VM traffic statistics.
- Enforce tenant allocations.

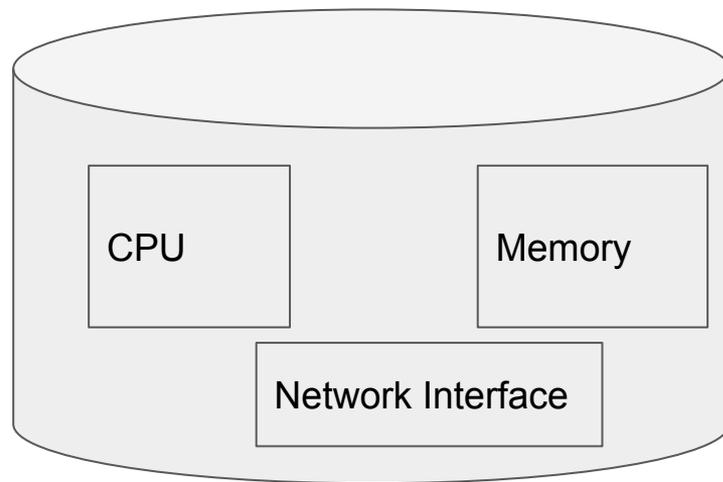


What is a resource to Pulsar?

- A single resource is a **network link** or an **appliance**.
- Pulsar treats appliances as **black boxes**:



A network storage as seen by Pulsar.



What a network storage really is.

What does Pulsar really allocate resources to?

- The most atomic operation is a **flow**.
- A flow encapsulates **all connections between a pair of VMs** that share the same (physical) path.
- A flow is assigned an allocation vector every control interval.
- **allocation vector = local policy at tenant's VDC + provider global policy**

Multi-resource allocation (MRA)

- The goal is to distribute multiple types of resources among clients with various demands.

Allocation	Flow 1	Flow 2
Storage access	5	12
Network	10	0

What will be allocated

=

Demand	Flow 1	Flow 2
Storage access	5	12
Network	10	0

How much a flow is capable of consuming

Weight	Flow 1	Flow 2
Storage access	1.0	3.0
Network	3.0	0.5

Weights form a policy (either local or global)

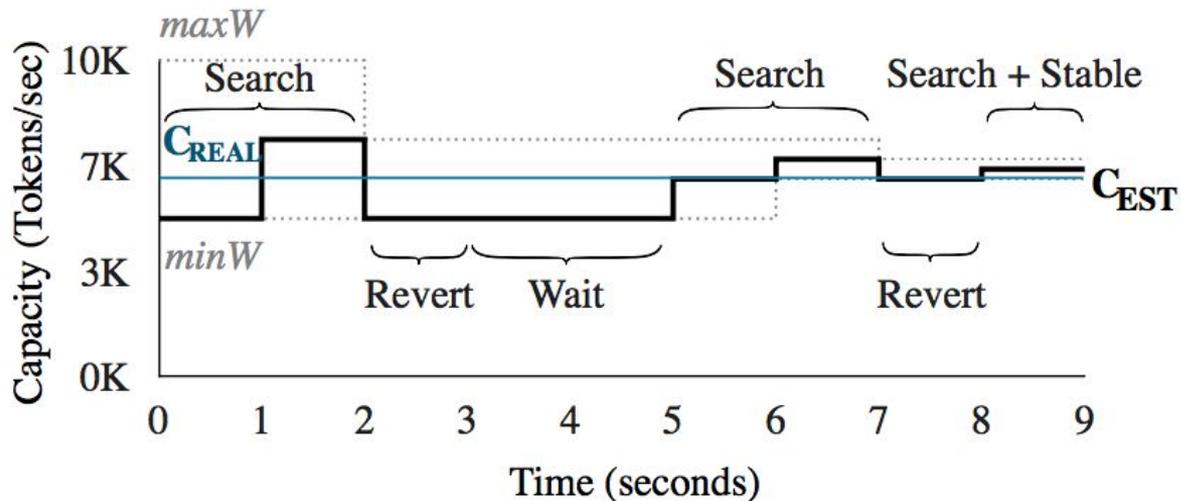
Availability	Resource
Storage access	17
Network	10

Available resources

Capacity estimation algorithm

There are 4 phases of finding the capacity.

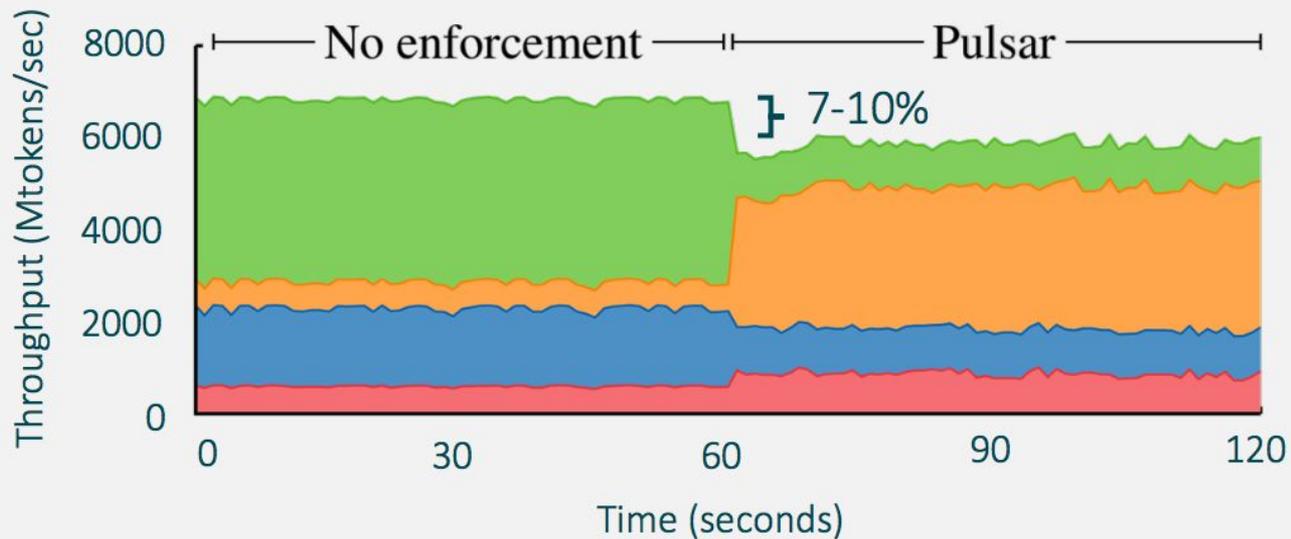
1. **Binary search increase** (try another optimal capacity candidate).
2. **Revert** (if the appliance is overloaded)
3. **Wait** (to handle the outstanding requests)
4. **Stable**. (The value with acceptable overhead has been found).



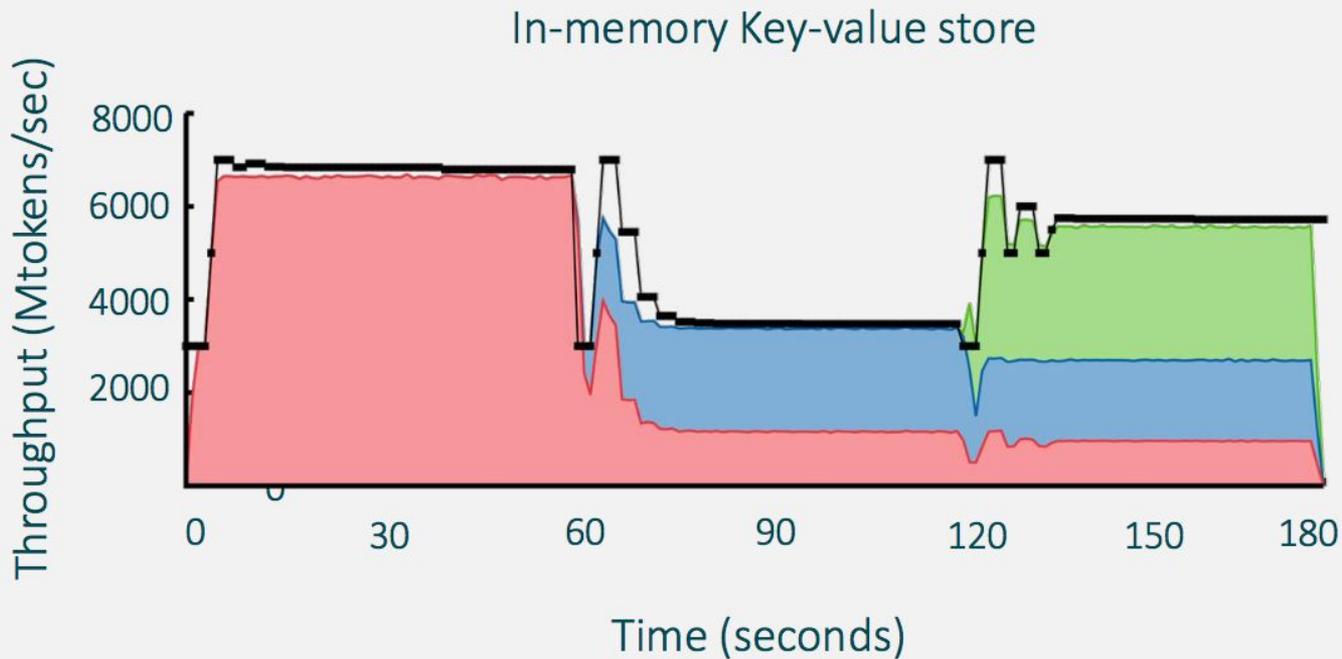
Pulsar evaluation

Pulsar achieves tenants' guarantees

Tenant	A*	B	C*	D
Guarantee	400 M	1600 M	800 M	800 M
VMs	49 (many flows)	48	8 (large IO window)	8



Pulsar estimates capacities and copes with changing workloads



Thank you

The presentation is based on the “End-to-end Performance Isolation Through Virtual Datacenters” paper by Sebastian Angel, Hitesh Ballani, Thomas Karagiannis, Greg O’Shea, and Eno Thereska. Some figures taken from the paper and the presentation by Sebastian Angel.