

Research on Cloud Computing @ REINS

Haopeng Chen

***RE*liable, *IN*telligent and *SC*alable Systems Group (**REINS**)**

Shanghai Jiao Tong University

Shanghai, China


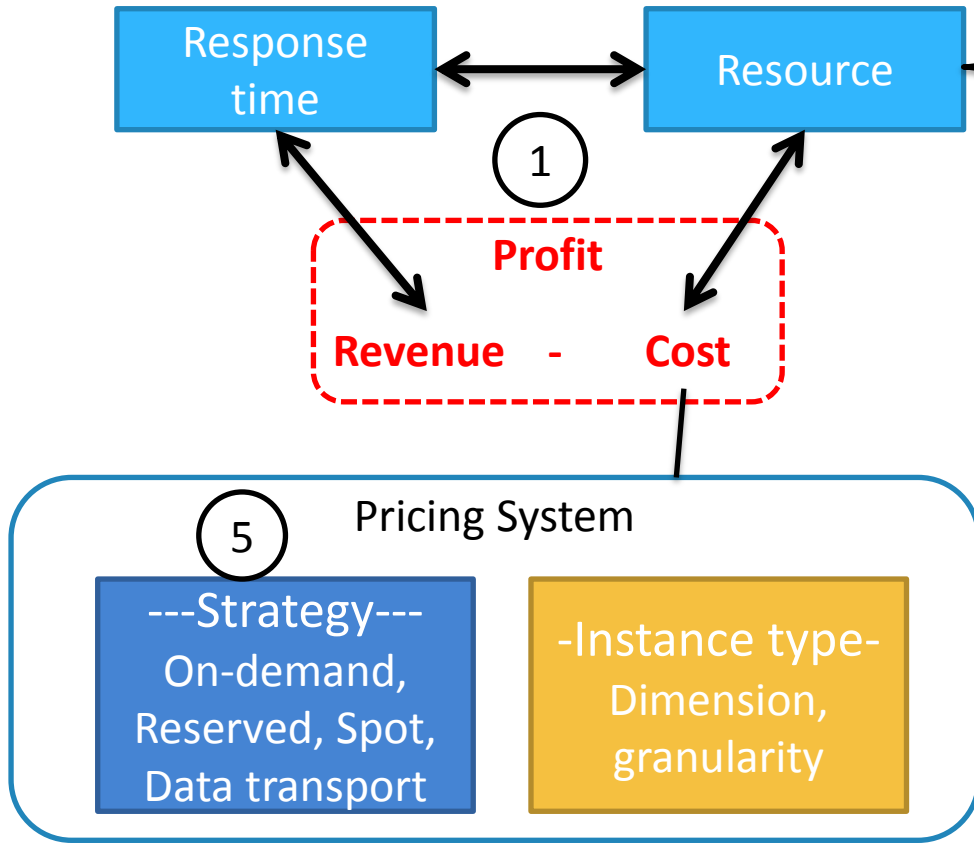
<http://reins.se.sjtu.edu.cn>

e-mail: chen-hp@sjtu.edu.cn

- Cloud resource management
- Cloud federation
- Data storage
- Data processing

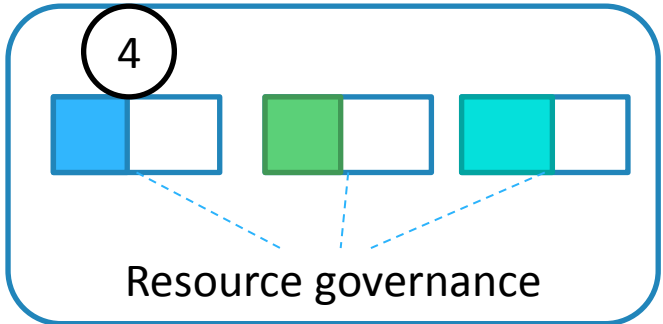
Cloud Resource Management

Commercial View



App	location
Service1	EU
Service2	America-west
...	...





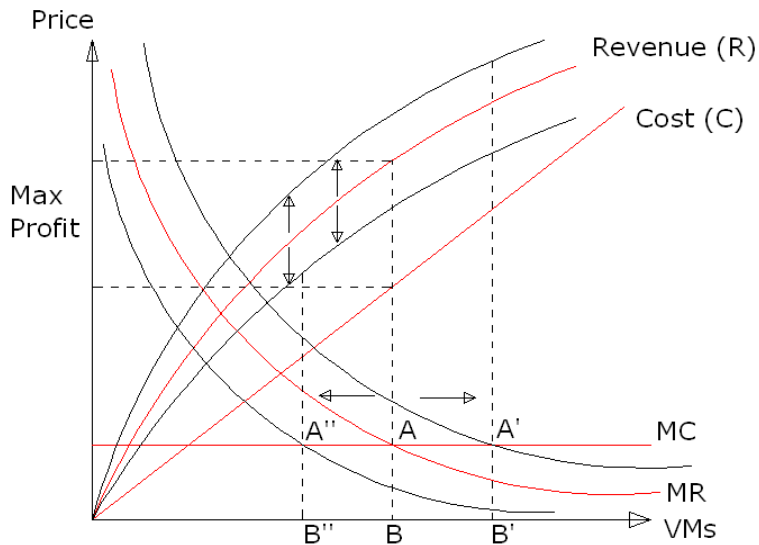
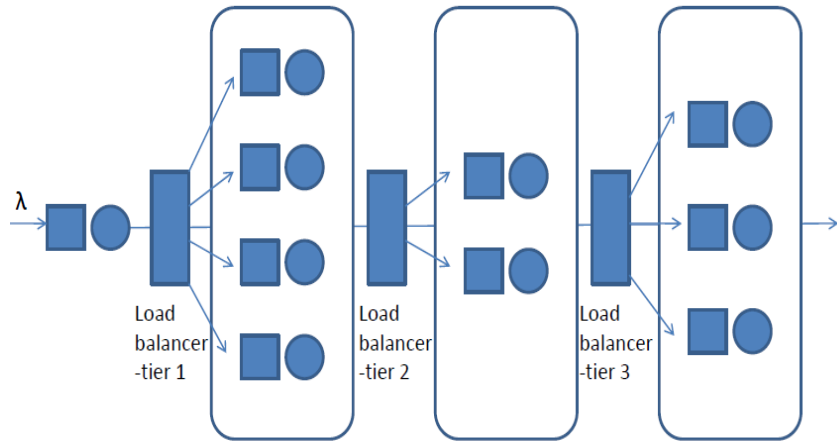
A circled number 3 is located below the table.



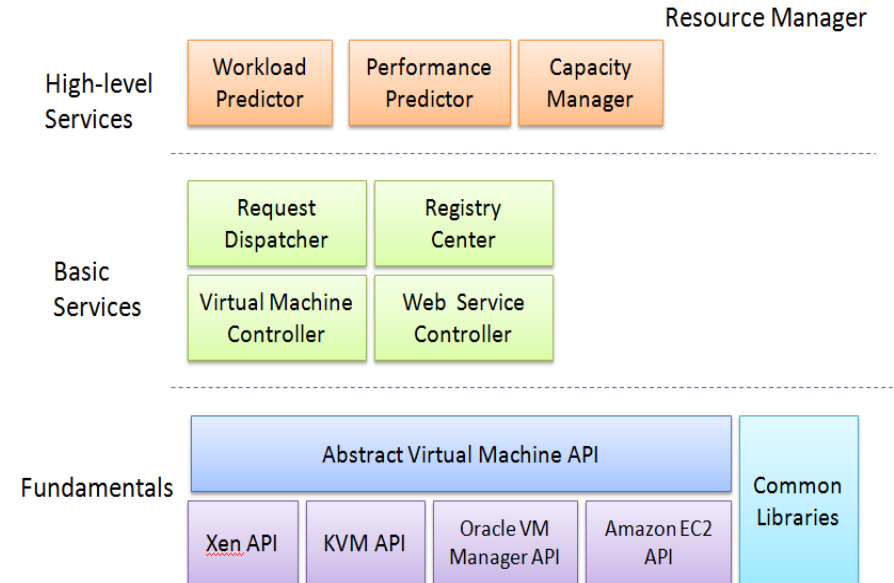
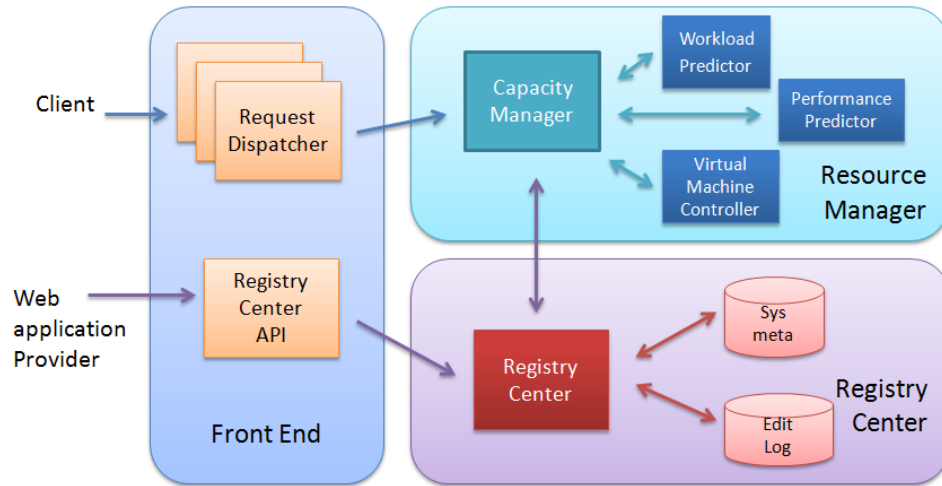
VM life cycle management

Performance Model


HTTP servers tier Application servers tier Data Base server tier



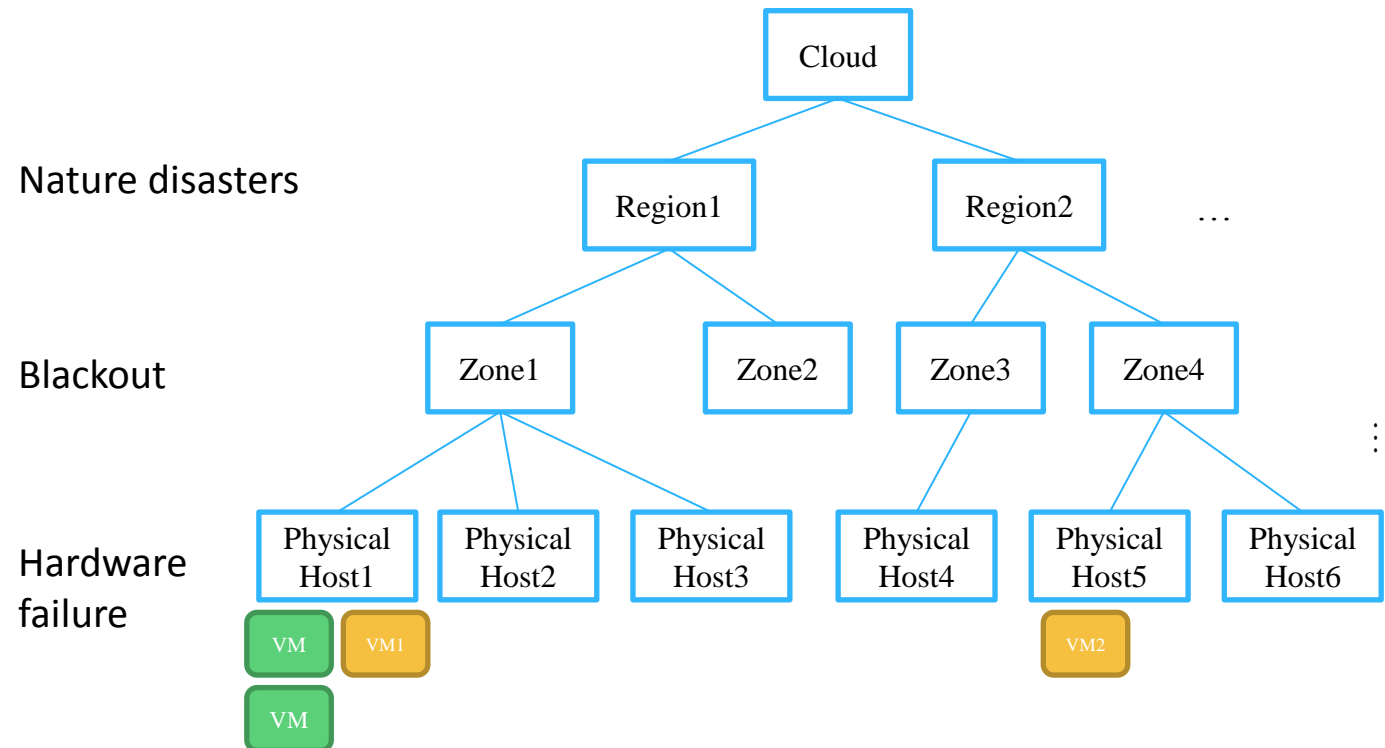
Performance Model



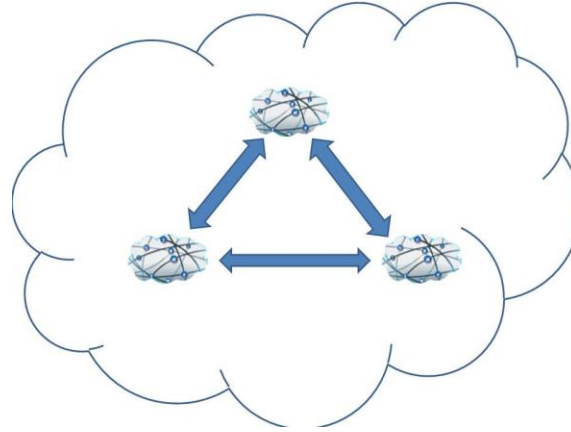
- Motivation
 - Amazon EC2
 - **Region and availability zone**
 - Microsoft Azure
 - ***Fault and Upgrade domains***
 - *guarantee 99.95% availability when you deploy two or more role instances in different domains*



- Approach
 - Hierarchy Tree Structure
 - Conditional-probability Availability Model

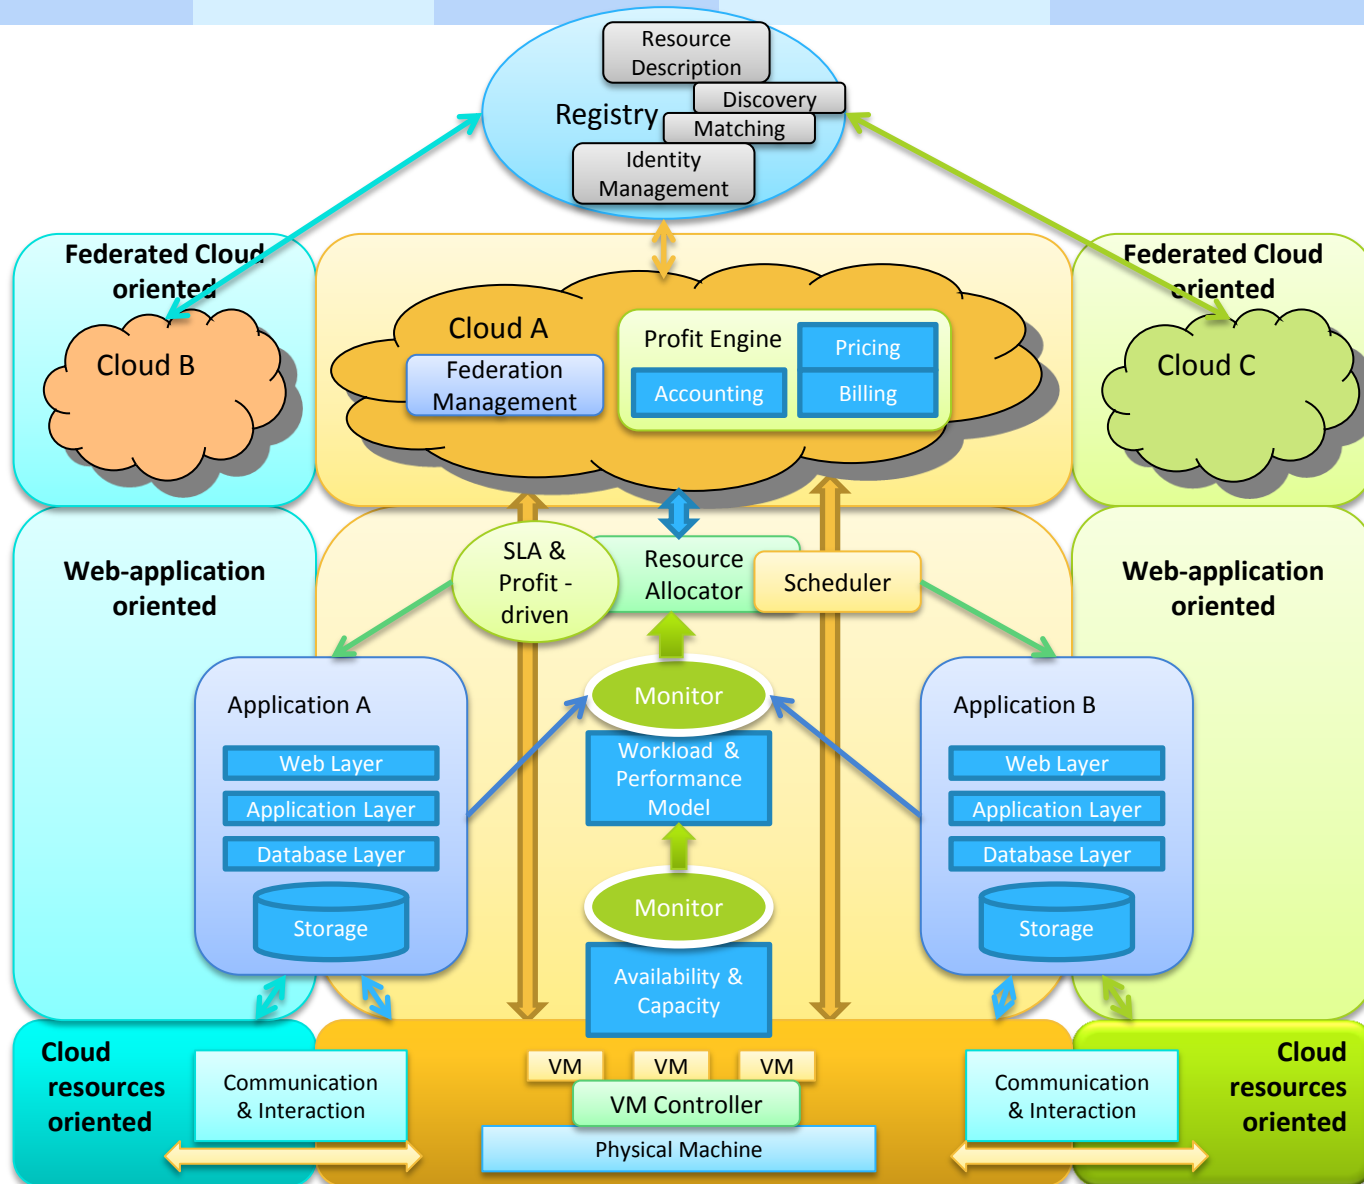


- Message transmission & collaboration among Clouds




- An ideal solution to
 - tasks overload for a single Cloud
 - user's need to transplant their Cloud-apps
 - a new business model for Cloud providers


Cloud Federation



- Problems
 - Cycle leased

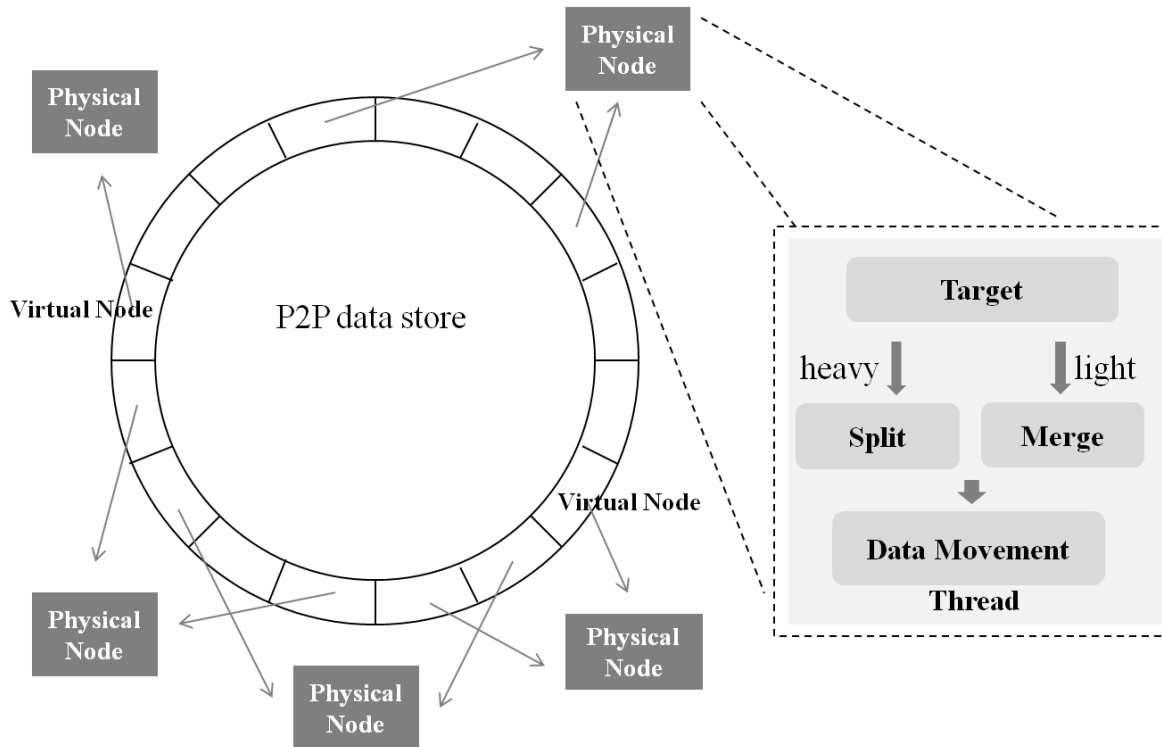


remote running



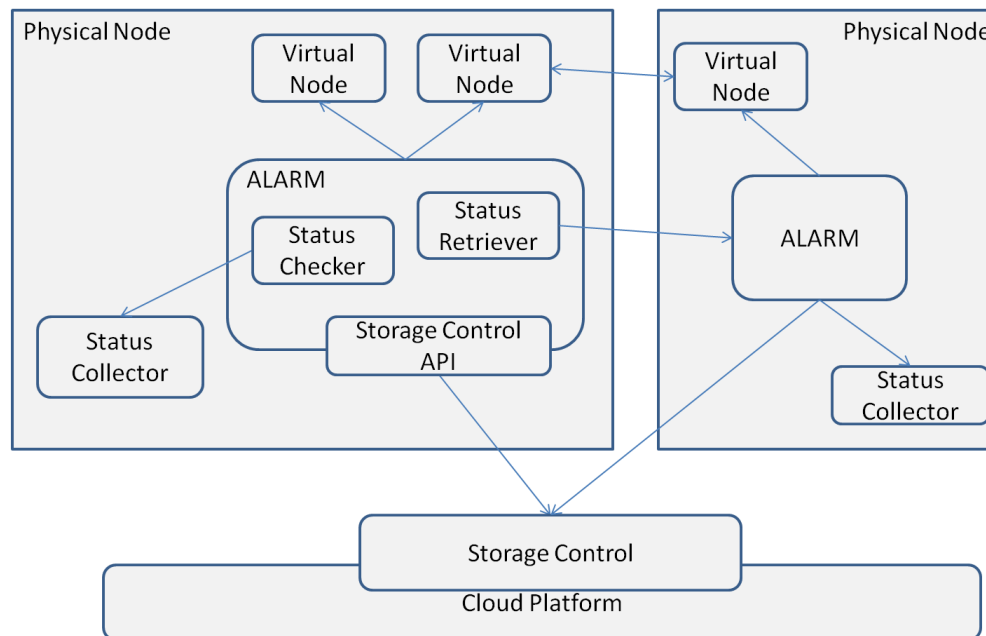
cycle solution

- ALARM: Autonomic Load-Aware Resource Management for P2P Key value Stores in Cloud

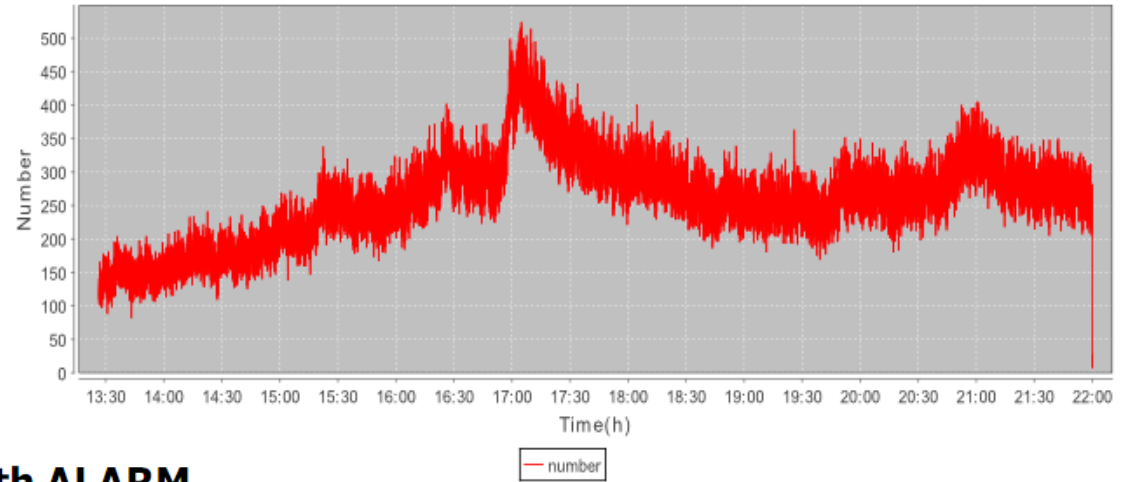


- Physical Node**
- Virtual Node**

- Storage Control : Interact with cloud platform, maintain storage info
- Status Collector : collect the resource utilization on the local node
- Status Checker : execute the target operation
- Status Retriever : retrieve info from other nodes in the system

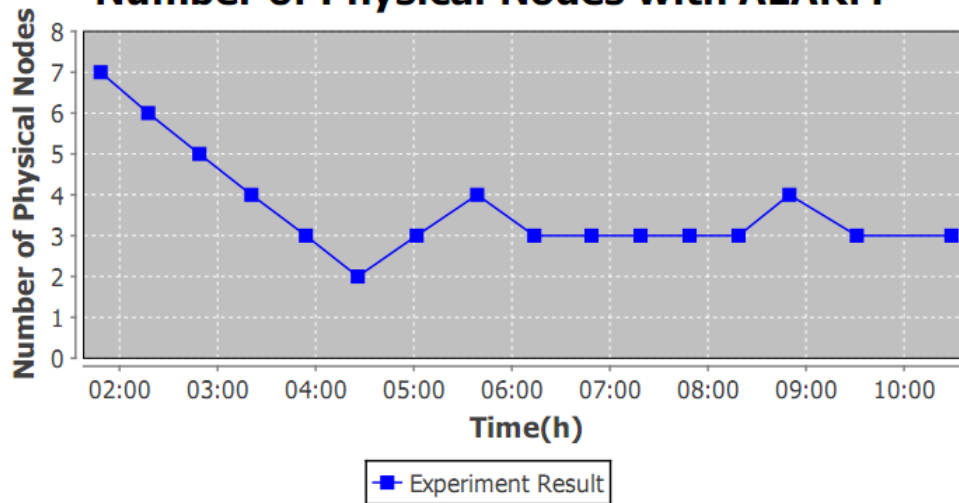


Log Numbers

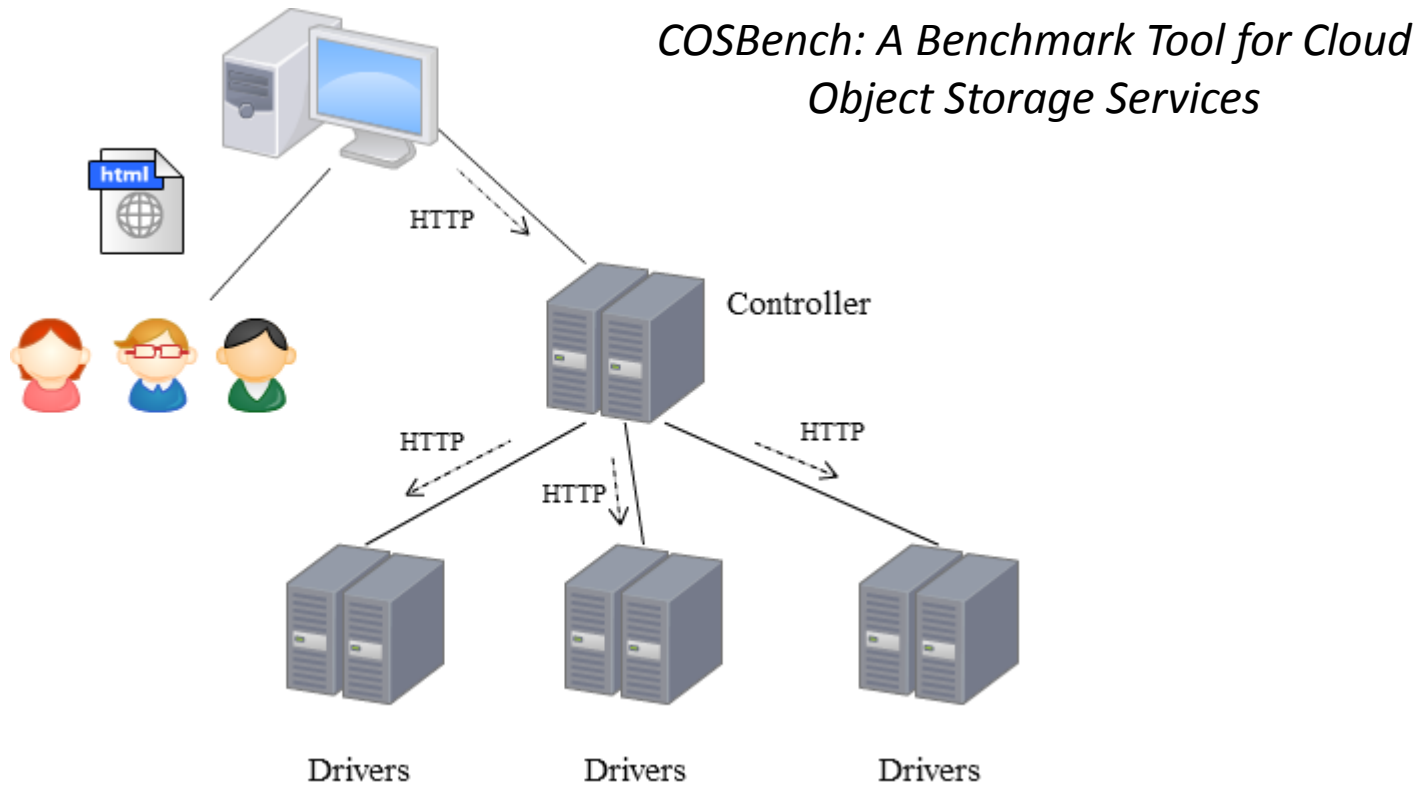


the number of client requests per second

Number of Physical Nodes with ALARM



- Implementing benchmark suite for characterizing the performance of object storage systems

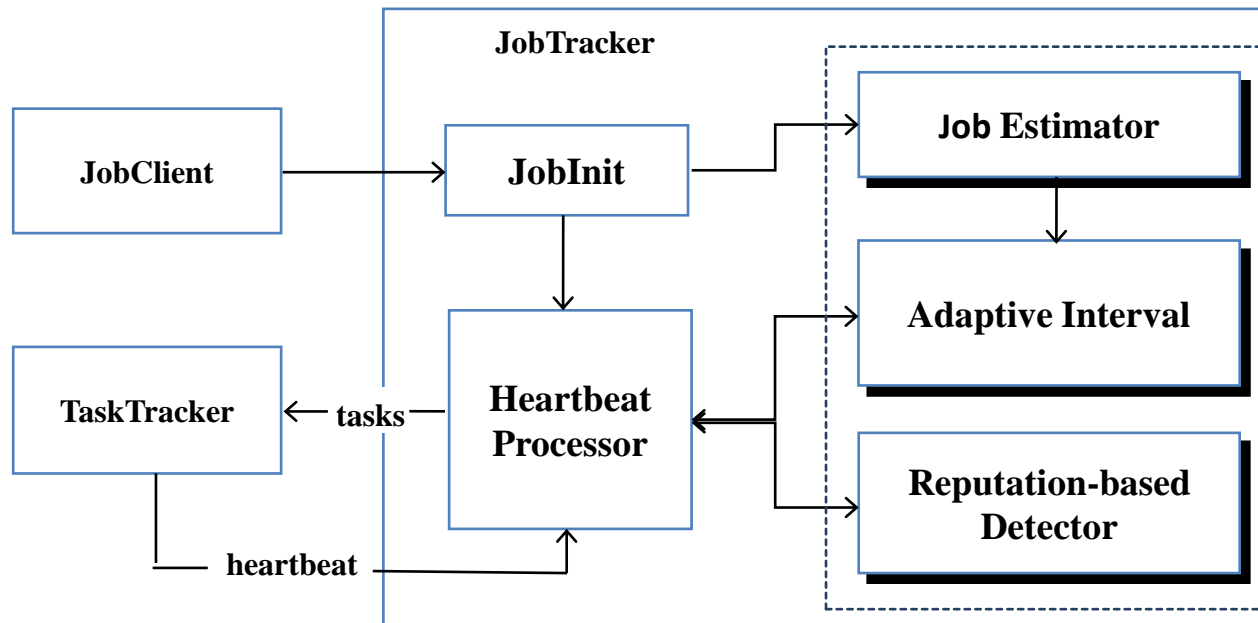


- Define workload model that resembles the usage patterns of a variety of real world systems/applications
 - Read intensive
 - Write intensive



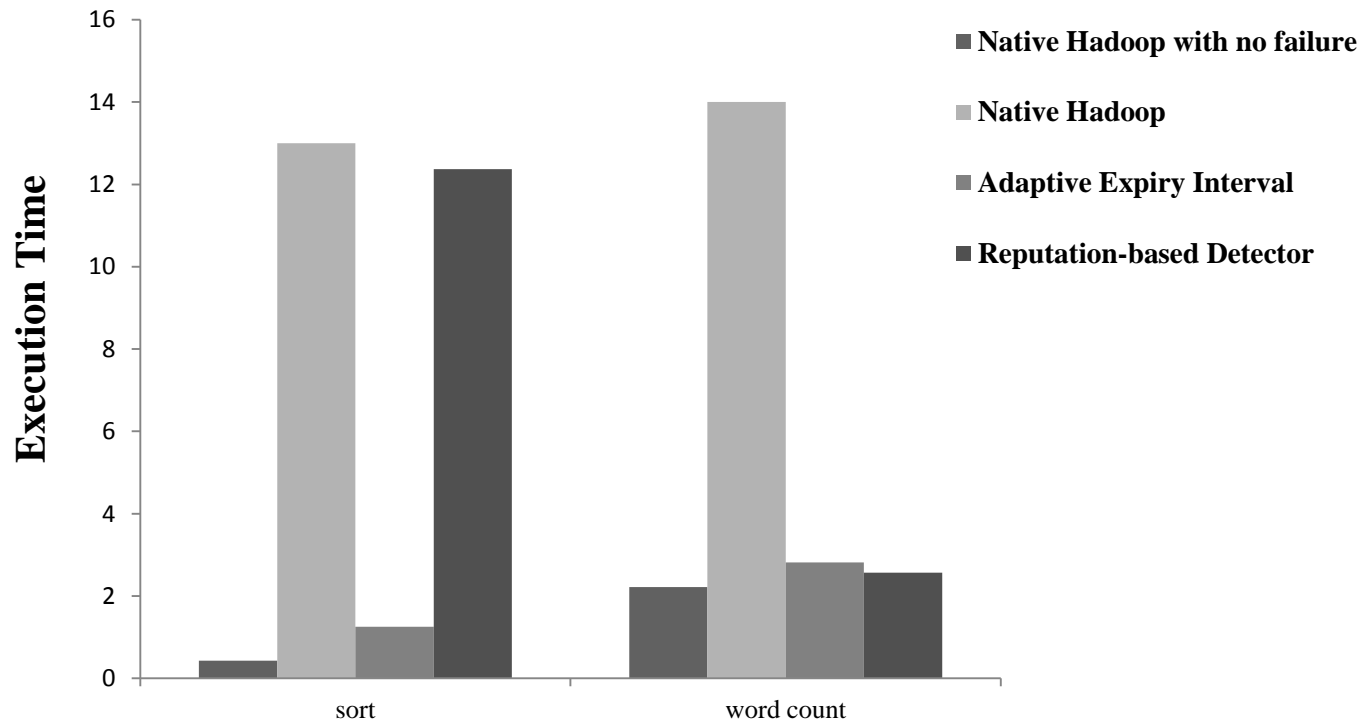
- Delayed detection of the failure worker
 - increase execution time
 - mis-blacklist
 - unnecessary backup tasks
- Goal
 - effectively detect the failure in the MapReduce cluster
 - **Adaptive failure detection via heartbeat under Hadoop**

- Adaptive Interval
- Reputation-based Detector



- Results

- Adaptive interval is better for small jobs
- Reputation-based detector is better for large jobs



- X10 in cloud
- Mongo DB in Data balancing
- Lucene and Solr
-



Thank You!