

A Resource and Policy Aware VM Scheduler for Medium-Scale Clouds

Deepal Jayasekara, Mohamed Infaz, Vishmi Money,
Shemil Hashan, and H. M. N. Dilum Bandara

Department of Computer Science and Engineering,
University of Moratuwa

Problem

Medium Scale Enterprise / University Clouds

Varying user requirements



Limited resources



Complex resource allocations



Solution

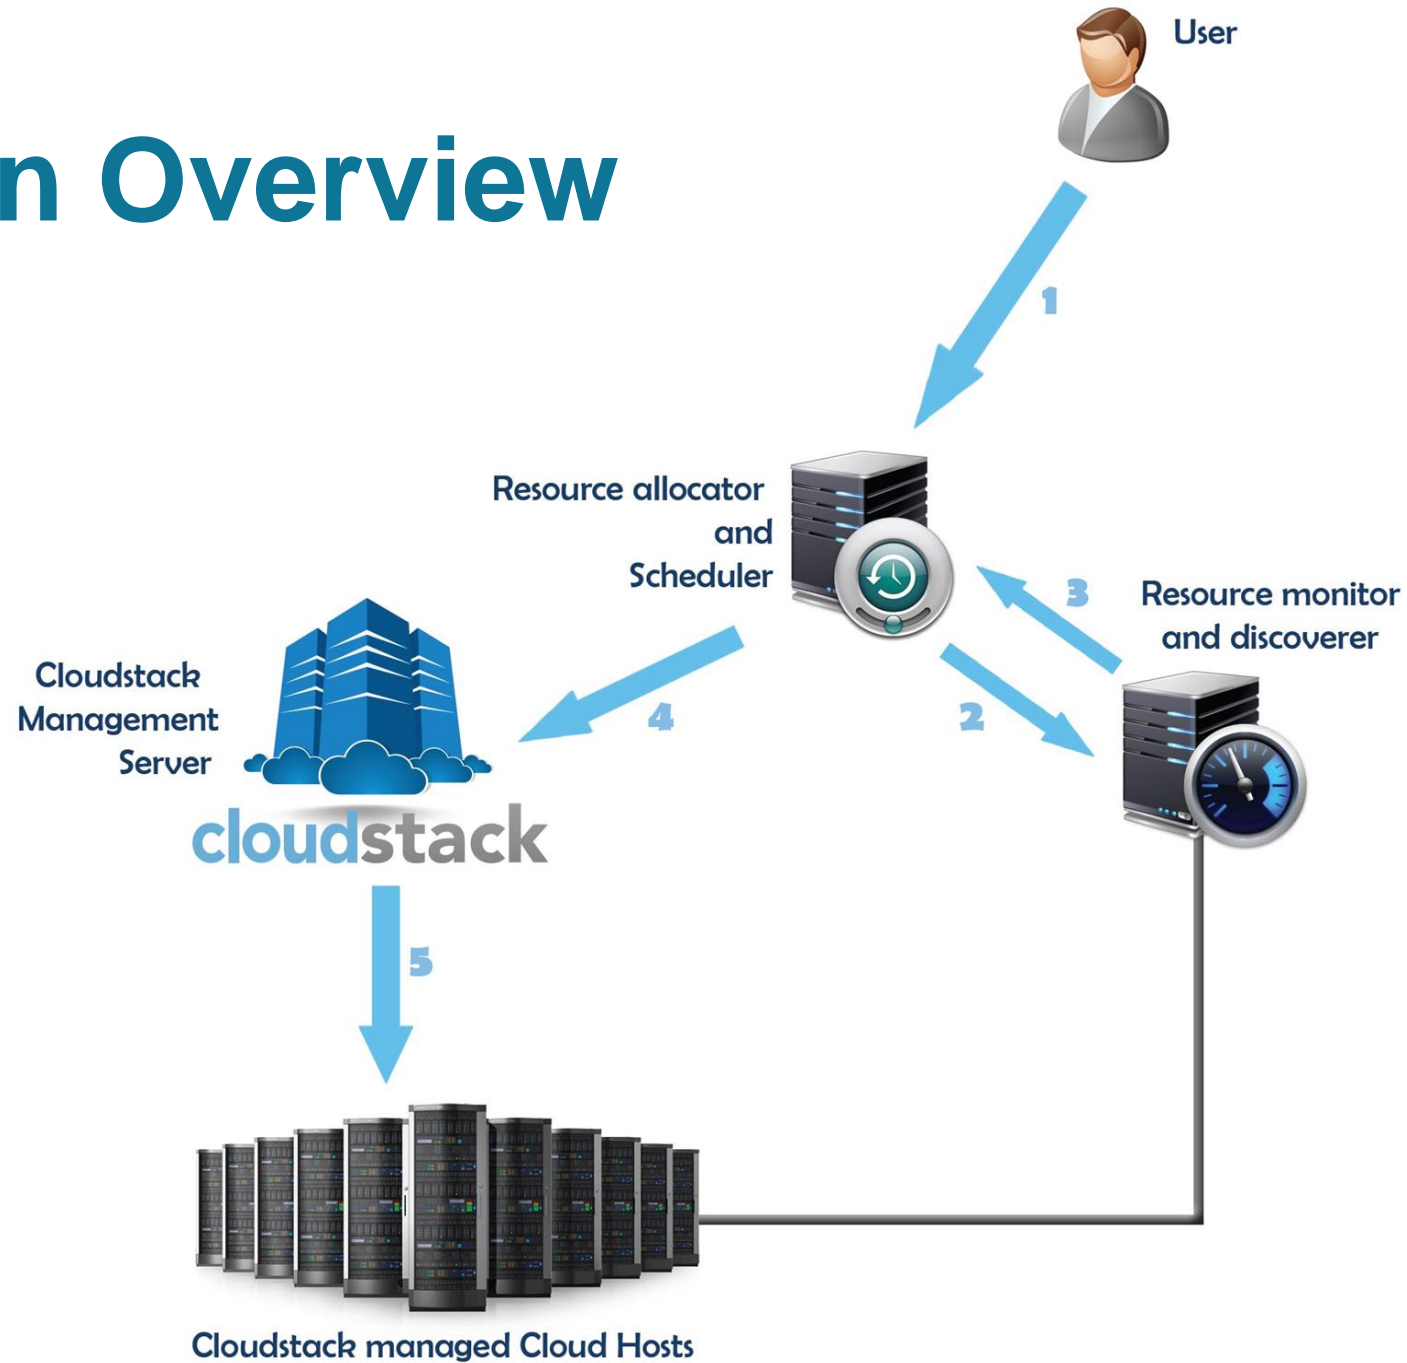
Virtual Machine scheduling mechanism

- which is aware of current resource utilization and availability,
- which is aware of defined policies such as user roles and job priorities,
- which performs complex resource scheduling by extending the functionality of a generic IaaS Cloud Framework

Feature Comparison Among IaaS Platforms

Desirable Policies	Eucalyptus	Openstack	Cloudstack	VMware	Our solution
For Private Clouds	Yes	Yes.	Yes	Yes	Yes
Applicable in small to medium scale	Yes	Yes. But high overhead	Yes	Yes	Yes
Support for user policies	No	No	No	No	Yes
Resource-aware scheduling	Static Resources only	Static & Dynamic Resources	Static Resources only	Static & Dynamic resources. DRS performs initial placement when cluster is specified	Static & Dynamic Resources
Deadline-aware	No	No	No	No	Yes
Migration support	No	Yes	Yes	Yes	Yes

Solution Overview




Evaluation of Technologies

- Resource Monitoring System
 - Nagios [1]
 - Zabbix [2]
- IaaS Platform
 - CloudStack [3]
 - OpenStack [4]
 - Eucalyptus [5]
- Hypervisor
 - Xen [9]
 - KVM [10]
- Programming Language
 - Node.js [8]
- Database
 - MySQL [6]
 - MongoDB [7]

Selected Technologies

ZABBIX

cloudstack

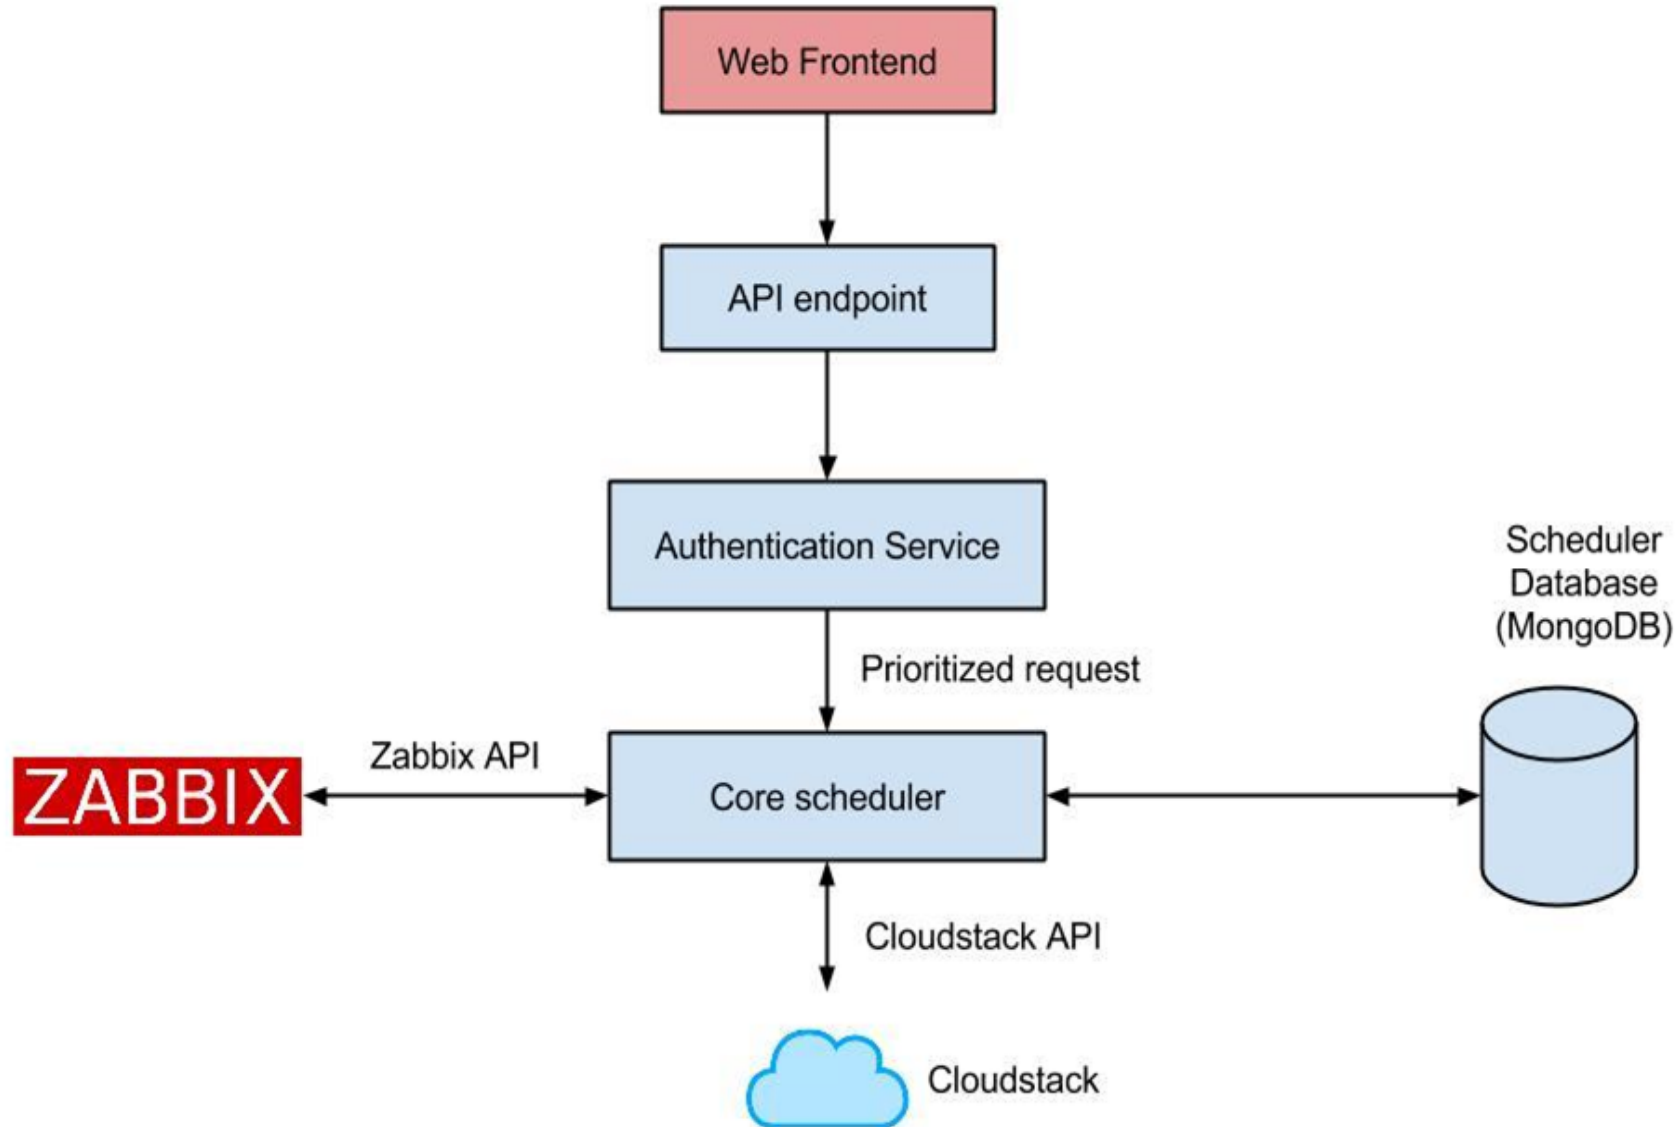


mongoDB

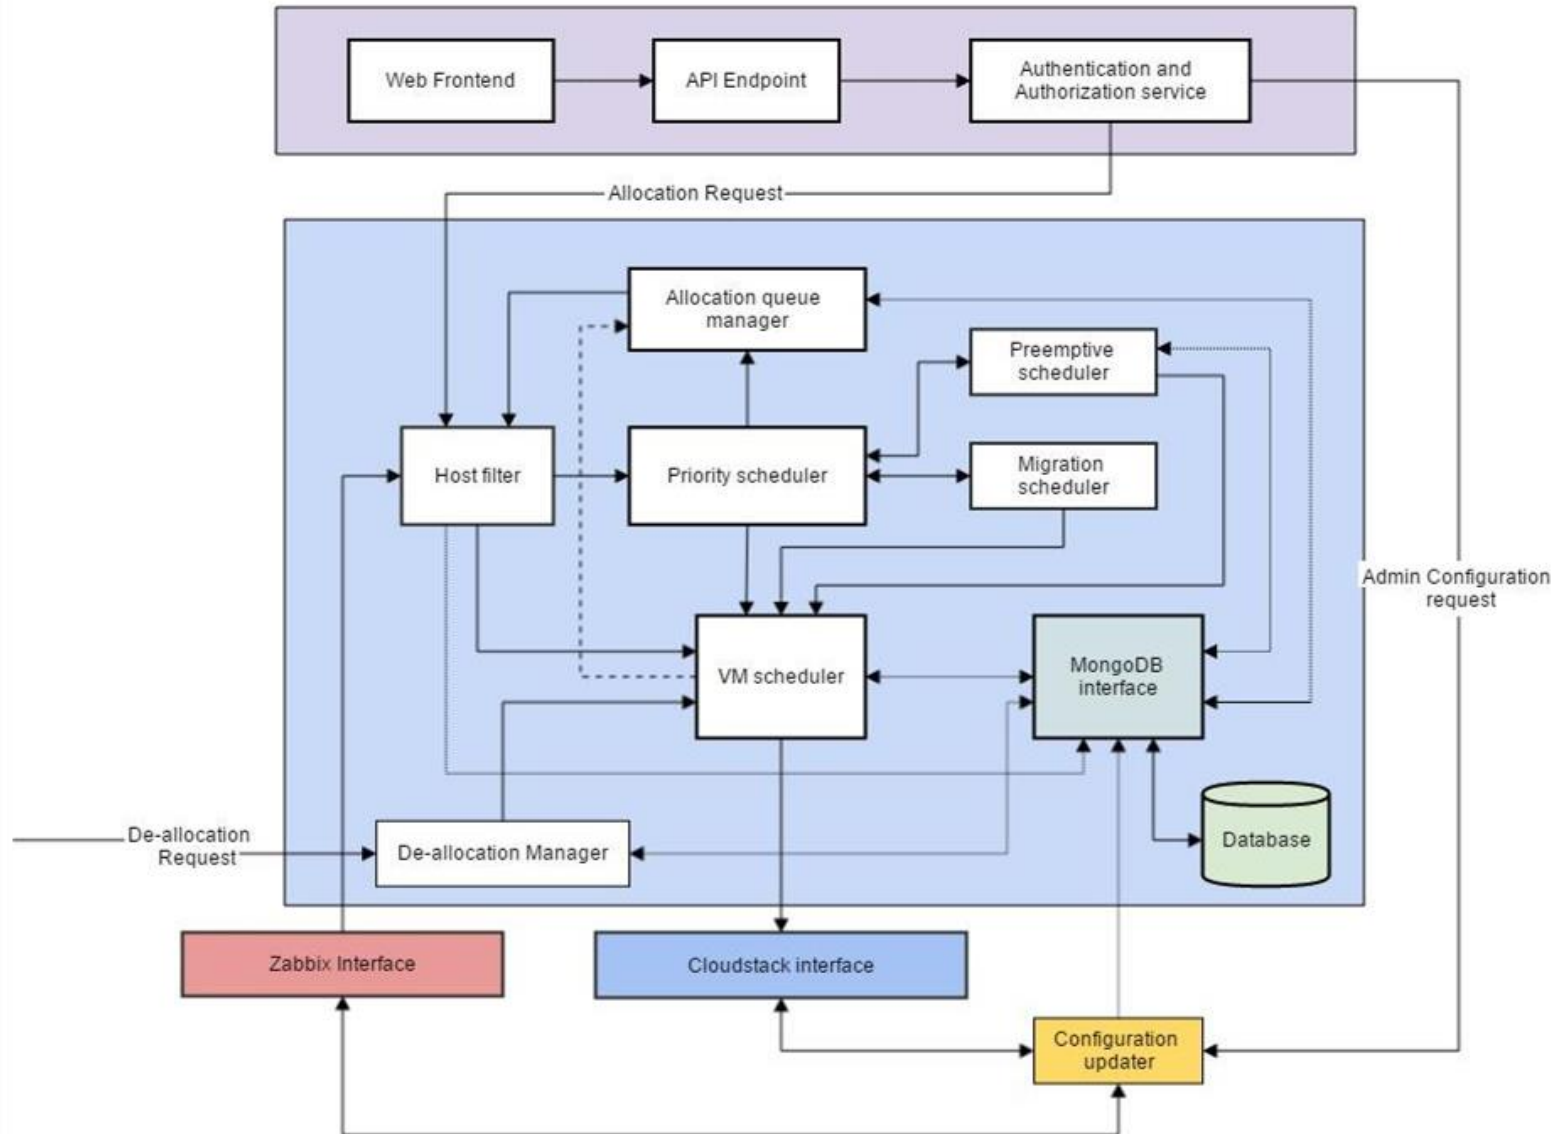
node JS™

KVM

High-level Architecture



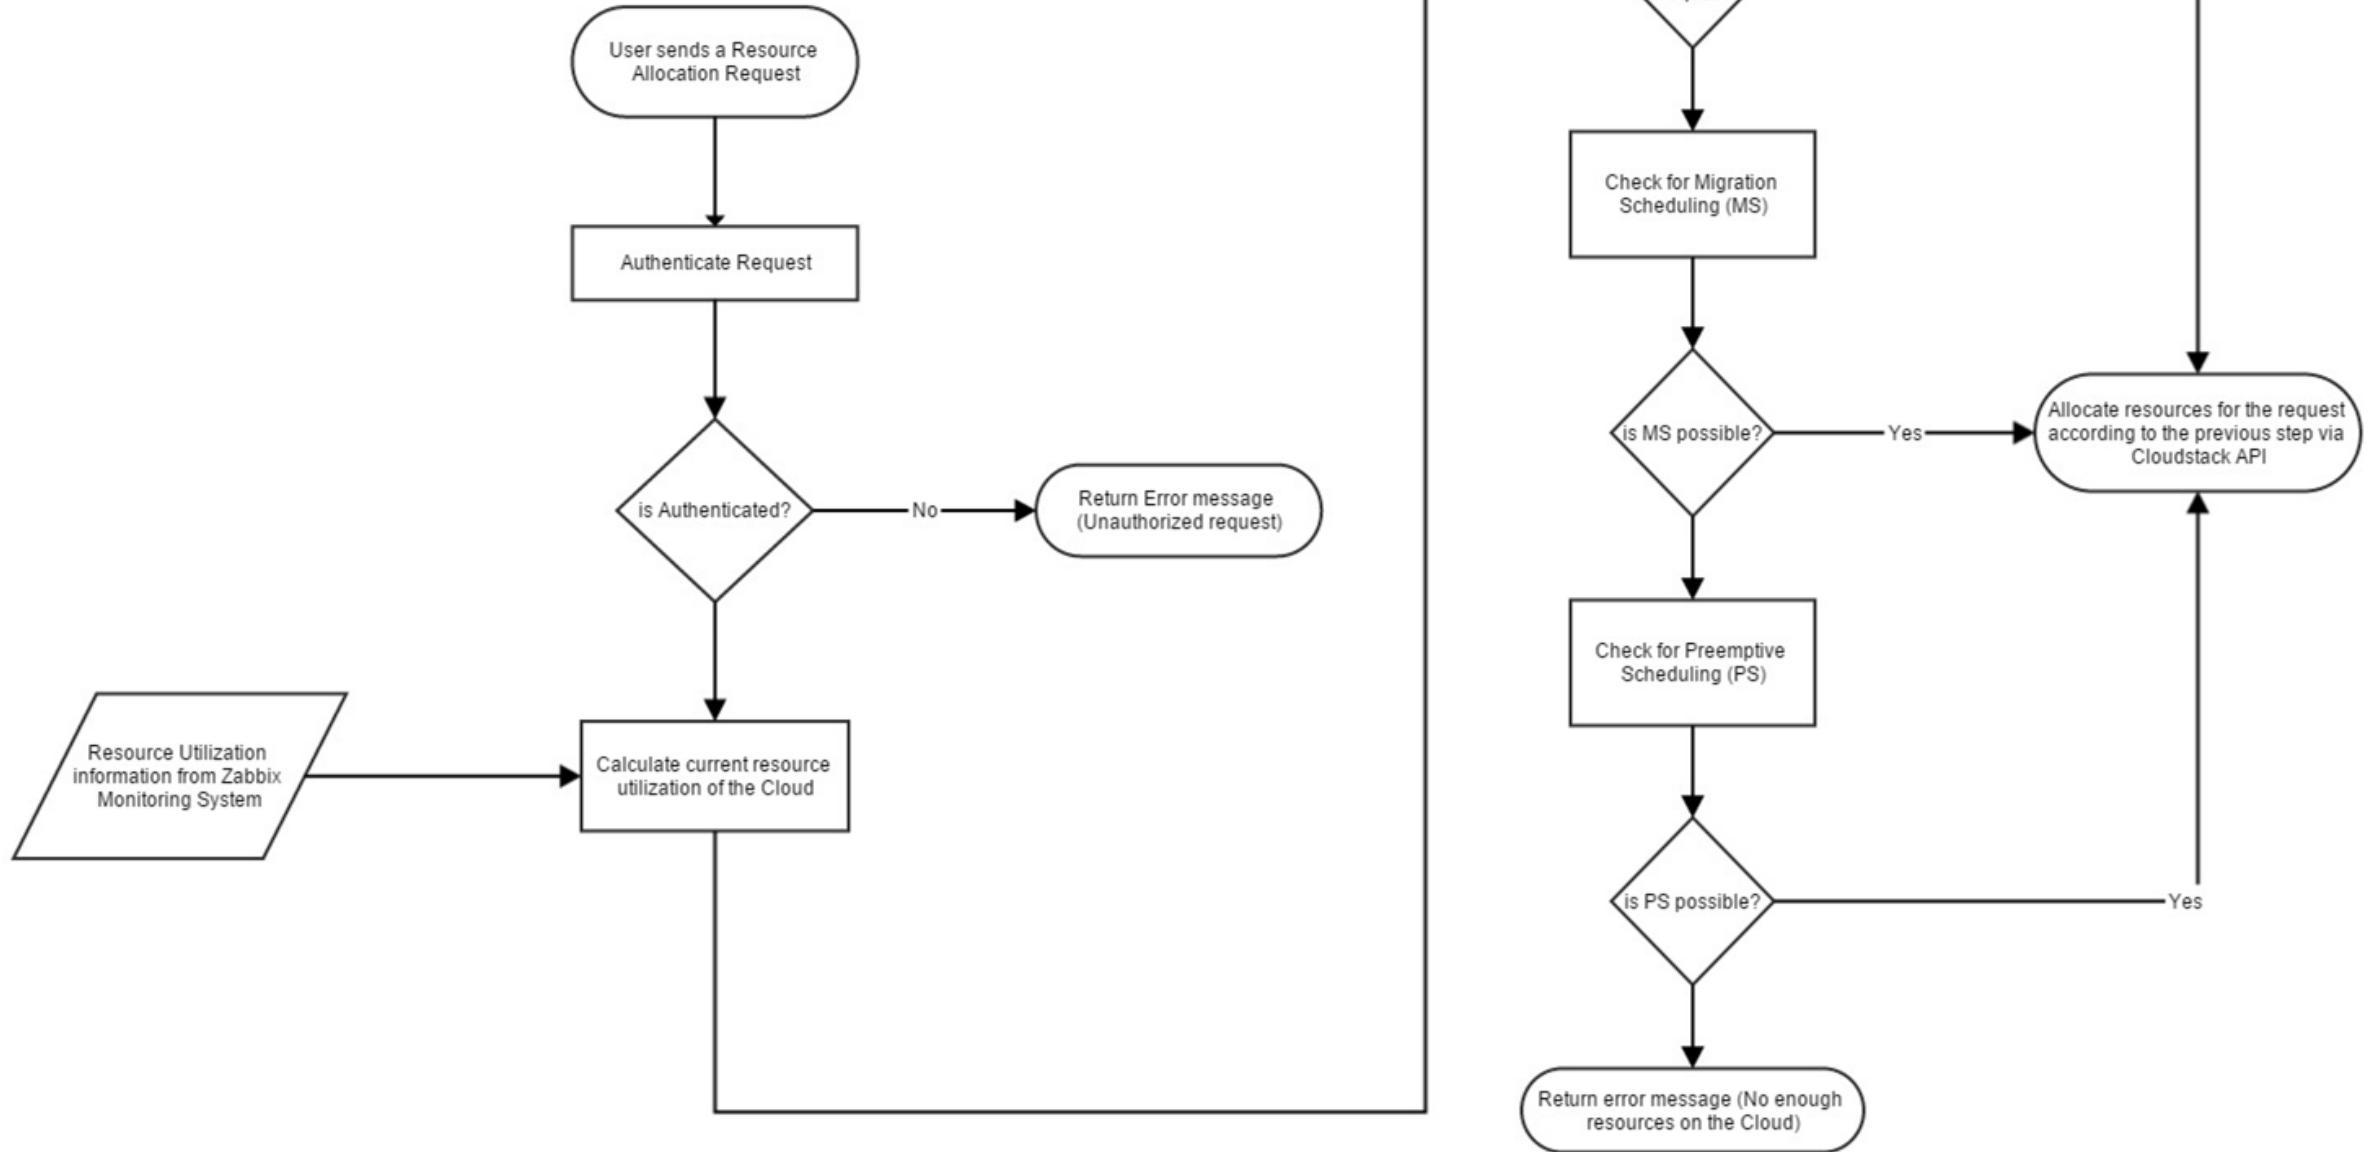
Core Scheduler Architecture



Resource Requests

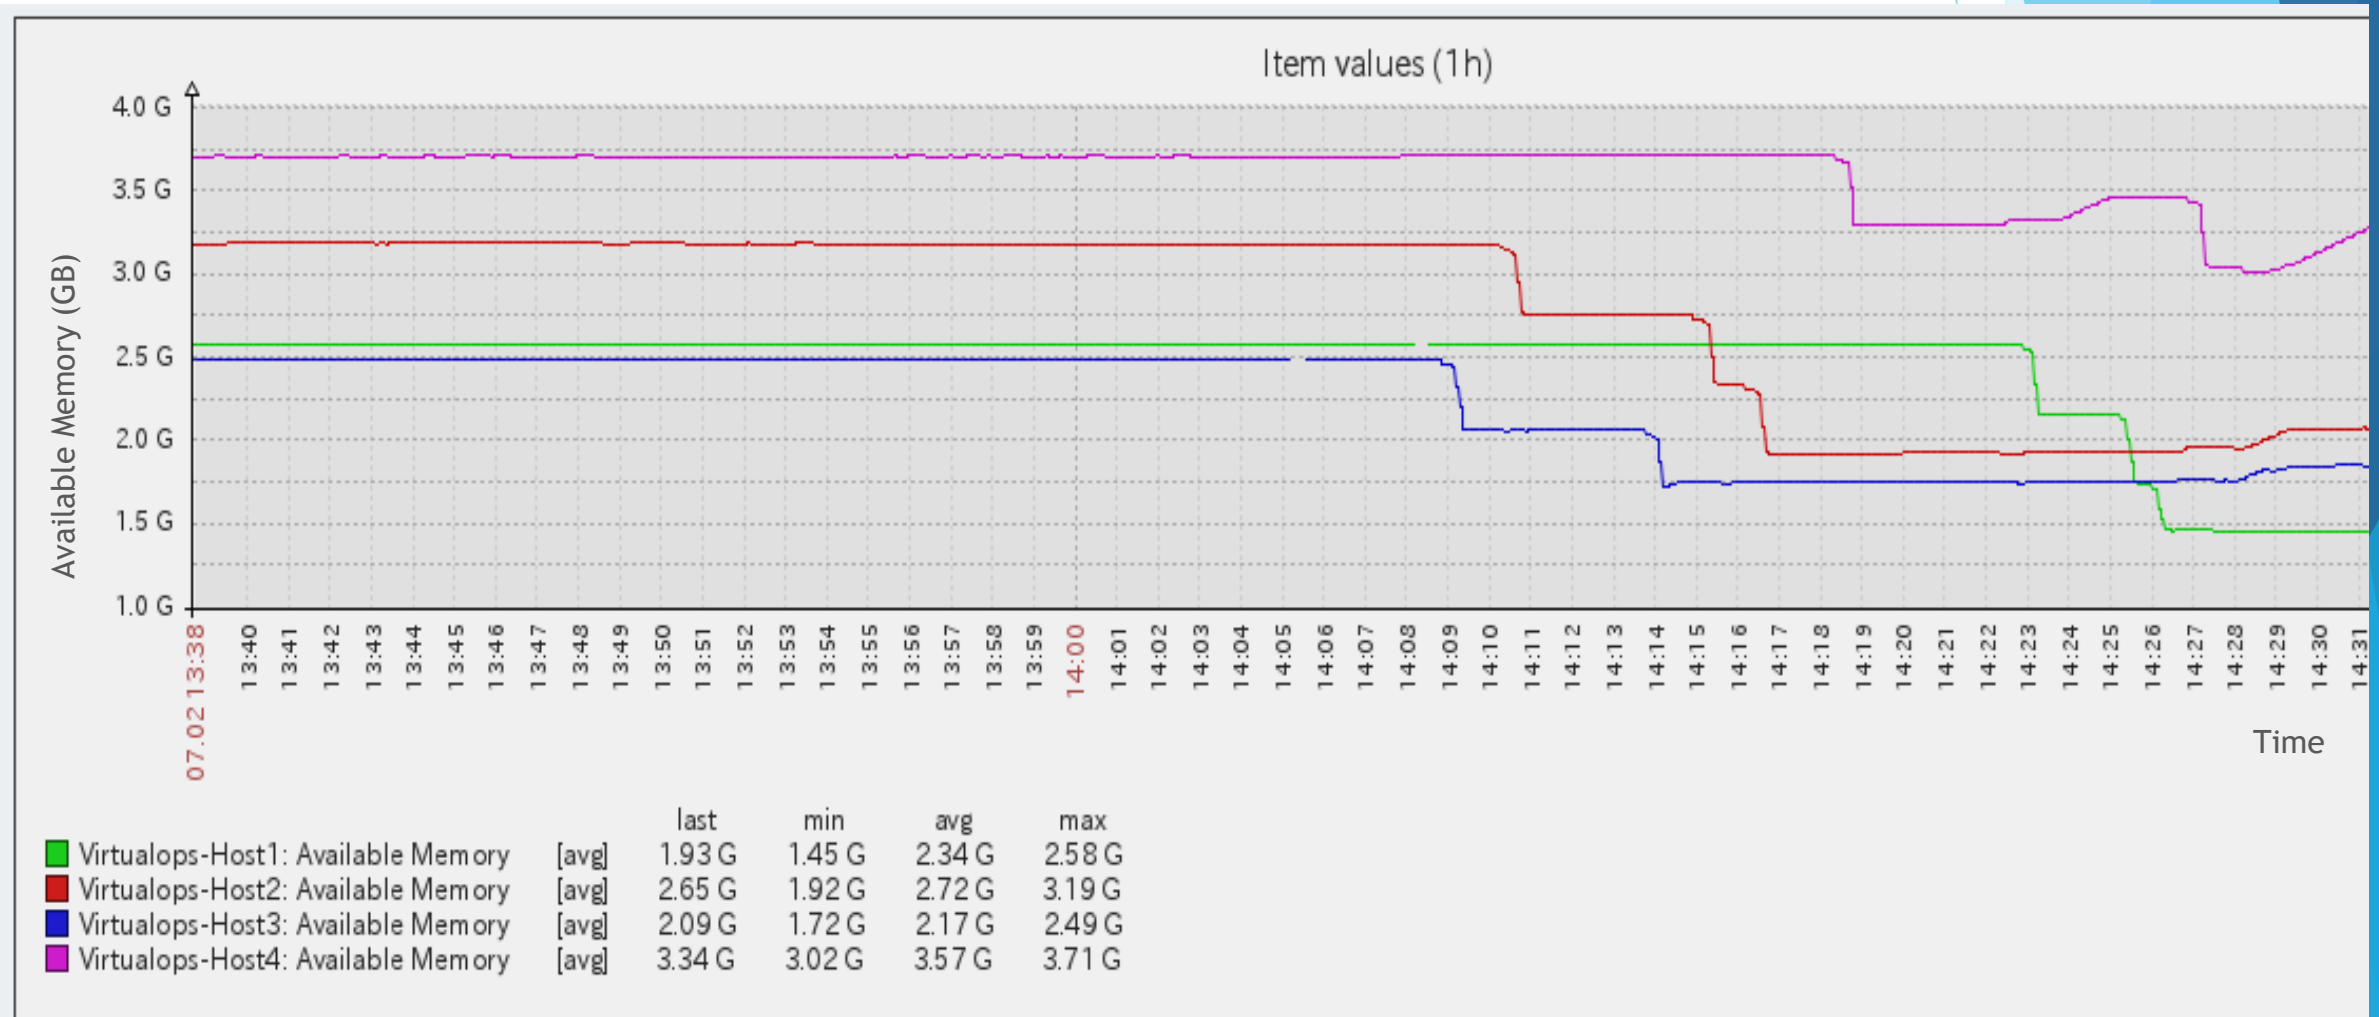
```
<resource_request>
  <group>
    <vm_count>10</vm_count >
    <image>
      <type>iso</type>
      <id>iso1</id>
    </image>
    <cpu>
      <cores>1</cores>
      <frequency>1</frequency>
      <unit>GHz</unit>
    </cpu>
    <min_memory>
      <size>2</size>
      <unit>GB</unit>
    </min_memory>
    <min_storage>
      <primary>5</primary>
      <unit>GB</unit>
    </min_storage>
    <priority>3</priority>
  </group>
</resource_request>
```

Workflow

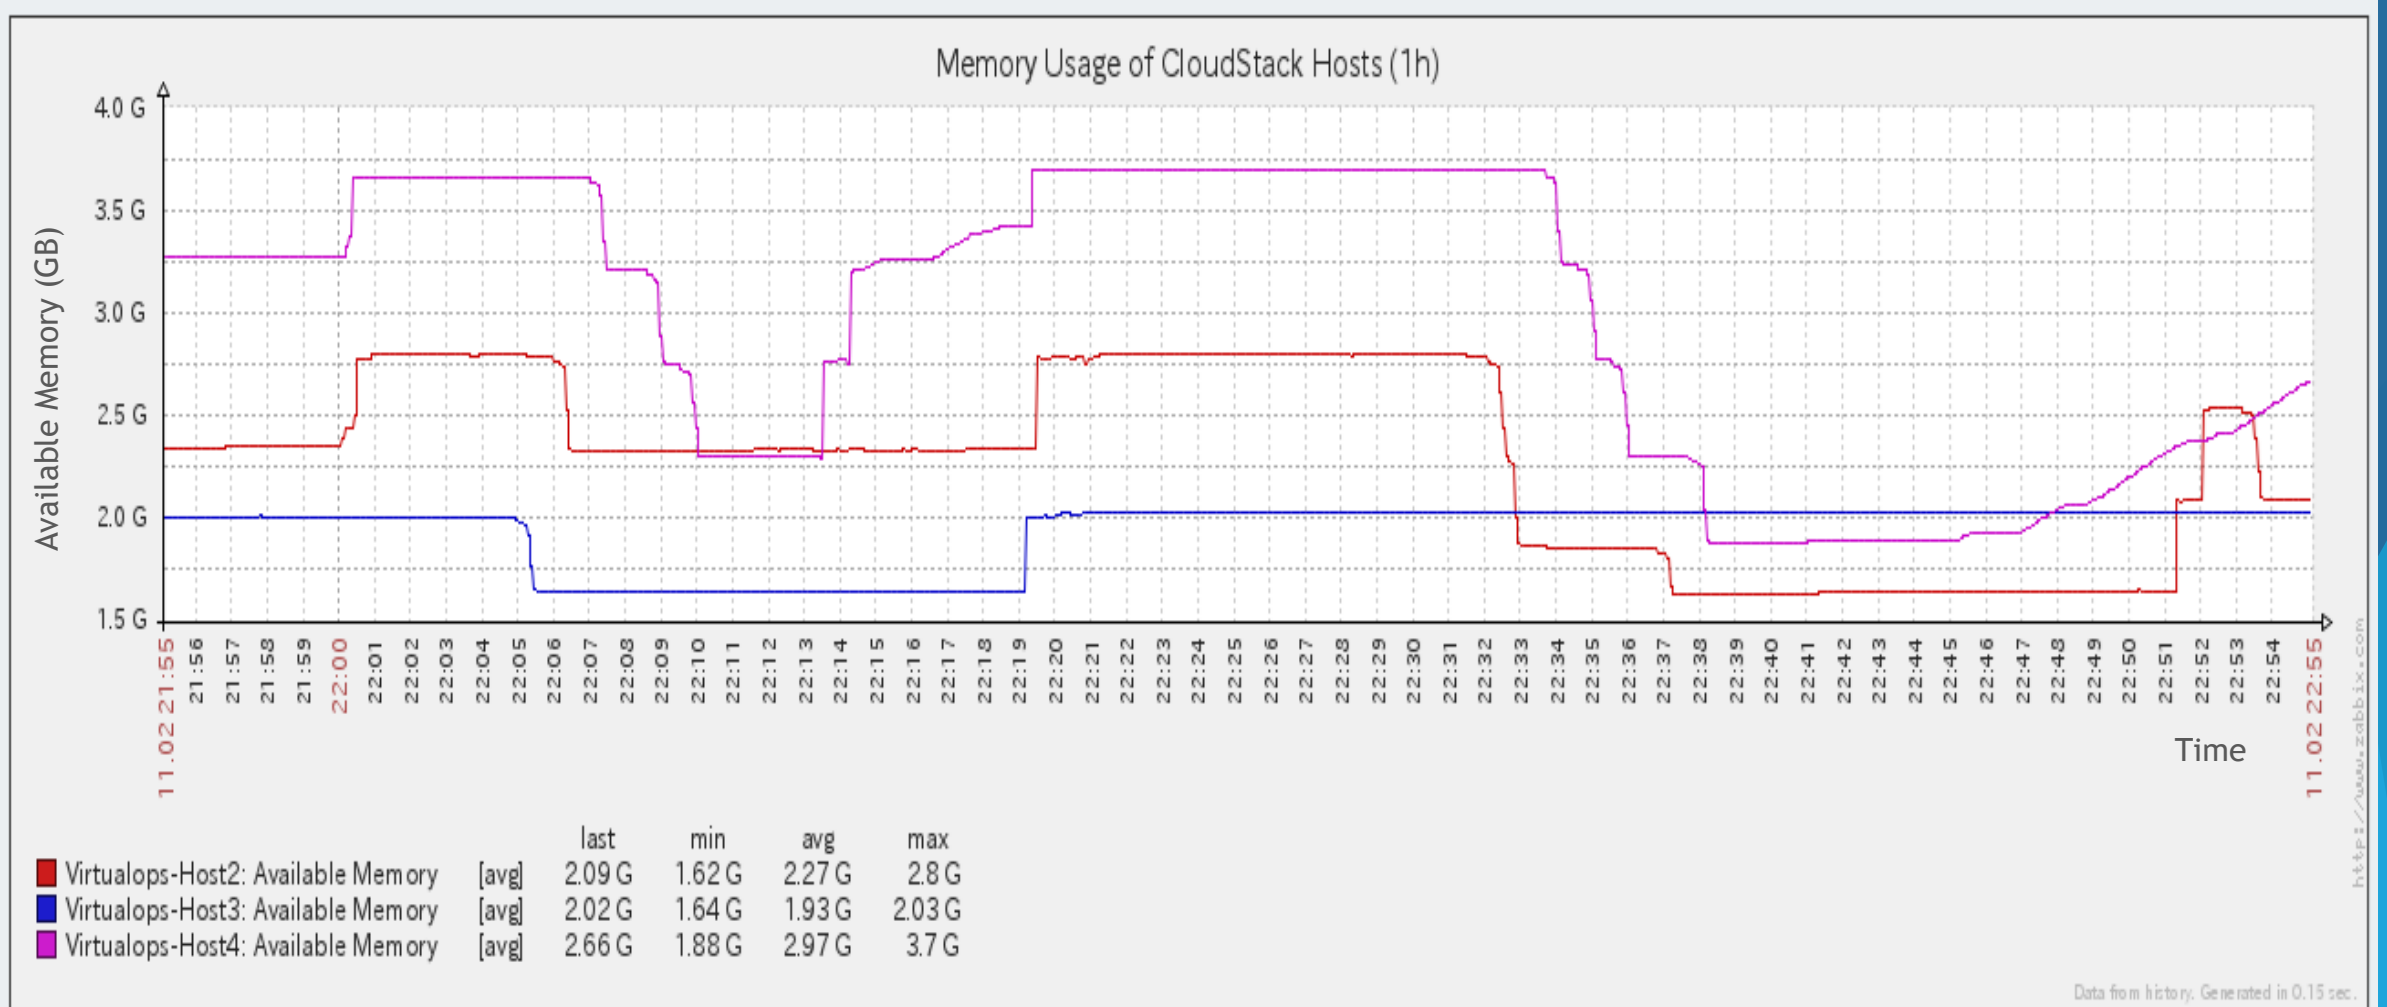


Performance Analysis

Resource Aware Virtual Machine Scheduling



Preemptive Scheduling



Challenges/Limitations

- Lack of support for KVM VM memory snapshots in CloudStack
- Issues in setting up cloud infrastructure
- Inconsistent Documentation
- Issues in latest CloudStack deployment

Future Work

- Support for Advanced Reservation
- VM Cloning
- Migration support for Virtual Machines with local storage
- Monitoring of VM resource utilization

References

- [1] Nagios.org,. 'Nagios - The Industry Standard In IT Infrastructure Monitoring'. N.p., 2015. Web. 11 Feb. 2015.
- [2] SIA, Zabbix. 'Homepage Of Zabbix :: An Enterprise-Class Open Source Distributed Monitoring Solution'. *Zabbix.com*. N.p., 2015. Web. 11 Feb. 2015.
- [3] Apache Cloudstack,. 'Apache Cloudstack'. N.p., 2015. Web. 11 Feb. 2015.
- [4] Openstack.org,. 'Home » Openstack Open Source Cloud Computing Software'. N.p., 2015. Web. 11 Feb. 2015.
- [5] Eucalyptus,. 'Eucalyptus: Open Source Private Cloud Software'. N.p., 2015. Web. 11 Feb. 2015.
- [6] Mysql.com,. 'Mysql :: The World's Most Popular Open Source Database'. N.p., 2015. Web. 11 Feb. 2015.
- [7] MongoDB.org,. 'MongoDB'. N.p., 2015. Web. 11 Feb. 2015.
- [8] Nodejs.org,. 'Node.js'. N.p., 2015. Web. 11 Feb. 2015.
- [9] User, Super. 'Hypervisor'. *Xenproject.org*. N.p., 2015. Web. 11 Feb. 2015.
- [10] Linux-kvm.org,. 'Main Page - KVM'. N.p., 2015. Web. 11 Feb. 2015.

Thank You!