

# High Performance Object Storage Service for Internet Scale Workload

*Qing Zheng, Haopeng Chen*

**Qing Zheng**

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

Shanghai Jiao Tong University

Shanghai, China

e-mail: qzheng2010@hotmail.com

# 2 Kinds of Data

## Structured Data

- Session state
- User profile
- Product catalog
- Order details
- Stock status

## Unstructured Data

- User avatar/files
- App/Sys logs
- Media objects
- DB snapshots
- Email archive

**84%!**

Data Analyze

Data Serving

**2.7 ZB**



How to efficiently manage  
Internet scale  
unstructured data ?

# Answer I : Cloud Computing



*Amazon Simple Storage Service*



Windows Azure

*Windows Azure Blob Storage Service*



*Google Cloud Storage*



*Rackspace Cloud File*

Economic of Scale

Security

Risk Control

No Data Lock-In

---

## Private Storage Cloud



Source software for building private and public clouds.

## Software

OpenStack Software delivers a massively scalable cloud operating system. The five major components are:

Compute	<a href="#">Details &amp; Download</a> ▶
Object Storage	<a href="#">Details &amp; Download</a> ▶
Image Service	<a href="#">Details &amp; Download</a> ▶

New Projects: OpenStack [Dashboard](#) and [Identity](#)

[All Software Projects...](#)

## Community

Join our global community of technologists, developers, researchers, corporations and cloud computing experts.

2642  
PEOPLE

155  
COMPANIES

[Meet Our Community](#)



Nimbus is cloud computing for science.

**Nimbus Platform** is an integrated set of tools that deliver the power and versatility of infrastructure clouds to scientific users. Nimbus Platform allows you to combine Nimbus, OpenStack, Amazon, and other clouds.

[Learn more >](#)

**Nimbus Infrastructure** is an open source EC2/S3-compatible Infrastructure-as-a-Service implementation specifically targeting features of interest to the scientific community such as support for proxy credentials, batch schedulers, best-effort allocations and others.

[Learn more >](#)



## THE OPEN SOURCE CLOUD PLATFORM

<a href="#">Learn</a> about cloud computing	<a href="#">Test Drive</a> eucalyptus	<a href="#">Download</a> the latest eucalyptus	<a href="#">Read</a> eucalyptus documentation	<a href="#">Participate</a> in our community
--	--	---	--	---

### First steps with Eutester

Posted: Monday, 5 Mar 2012 11:00 AM by Graziano Obertelli

In one of my previous post I mention our mantra: "Listen to our community and deliver quality software". To deliver quality software, it is necessary to have the QA process as a first class citizen. Our QA team created quite a spectacular infrastructure to test all sort of combinations and configurations automatically (distro, architectures, versions, hypervisors, networking, images etc ...) to guarantee our users regression-free releases. They also managed to have fun in the process (checkout [Pigeons on a Euca](#)).

[Eutester](#) is the latest brainchild of our QA team, and this blog is about my experience writing a test using it. Eutester is a framework to create automatic tests against a Eucalyptus installations (or any cloud following the AWS API for that matter).

These are the ingredients I needed to bake my first test:

[More...](#)

[Coming Soon: Eucalyptus 3 next](#)

### FASTSTART

Your Eucalyptus Cloud

[Learn how](#)

### Developer Poll

What is the most important feature to you in the 3.0 Roadmap?

- High Availability
- Identity Authorization and Management
- Active Directory/LDAP integration
- Windows Hosting Service
- Boot from EBS
- Resolved Issues

### Navigation

- [Recent posts](#)

### User login

Username: \*

  
  
Password: \*  

[Log in](#)

- [Create new account](#)
- [Request new password](#)

### Active forum topics





What is object storage?

- For people to access and store
  - public or private files
- In object stores
  - one create containers
  - and put objects into these containers for storage.

- Reading file
  - **GET** <http://reins.se.sjtu.edu.cn/objstor/container/object>
  - 200 OK
- Upload file
  - **PUT** <http://reins.se.sjtu.edu.cn/objstor/container/object>
  - 201 Created
- Removing file
  - **DELETE** <http://reins.se.sjtu.edu.cn/objstor/container/object>
  - 204 No Content

- Scalable
- Redundant
- Cost-effectiveness
- Ease-of-use
- Eventual consistent
- Consistent Hashing

- Data backup
- OS template repository
- User data storage
- Database archive
- Log/Email backend



How to evaluate the performance of a cloud object storage service?

- Comparison
  - hardware settings
  - software configurations or implementation
- Optimization
  - system tuning
  - system refactoring or algorithm refining
- Path-finding
  - future improvements
  - workload characterization

- API Standard
  - S3, Cloud Files, ...
- Workload Model
- Simple & Extensible
- System Design
  
- Other issue
  - Ease-of-use (configuration, UI)
  - Benchmark performance & scalability
  - Background data
  - System stability
  - Side effects of operations



- Concurrency pattern
  - worker number
  - container range
- Access pattern
  - object size
  - read / write ratio

- *COSBench*
  - Cloud Object Storage Benchmark
- *Architecture style*
  - master/slave
  - standalone

- COSBench: A benchmark tool for cloud storage services
  - being reviewed



Thank You!