
DISTRIBUTED SYSTEMS (COMP9243)

Slide 1


Lecture 10c: Cloud Computing

- ① Cloud Computing

WHAT IS CLOUD COMPUTING?

Slide 2

A style of computing in which dynamically scalable and often virtualized resources are provided as a service over the Internet. (Wikipedia)



Slide 3

Why is it called *Cloud*?

- services provided on virtualised resources
- virtual machines spawned on demand
- location of services no longer certain
- similar to *network cloud*

KEY ASPECTS OF CLOUD COMPUTING

Slide 4

- ① Clusters of Machines
- ② Virtual Machines
- ③ Services (VM management, IP address assignment, naming, storage, etc.)
- ④ Platform: OS & Middleware
- ⑤ Software stack (Web server, DBMS, etc.)
- ⑥ Applications & Web Services

Flavours of Cloud Computing:

- ✗ Application Services
- ✓ Application Platform: app development env e.g. App Engine
- ✓ Platform: OS and services e.g. Azure and .NET
- ✓ Infrastructure: virtual machine, BYO OS e.g. EC2

CLOUD COMPARED TO...

Cloud computing is not:

- The Internet (or Interweb, or Intertubes)
- Cluster computing
- Grid computing
- SOA
- Web 2.0 and Web services
- Software as a Service

Slide 5

Public vs Private Clouds?

Public: open services available to everyone

Private: owned, operated, and available to specific organisation
Is this still cloud computing?

BENEFITS

Flexibility:

- Flexible provisioning
- Add machines on demand
- Add storage on demand

Slide 6

Effort:

- Low barrier to entry
 - Initial effort: no need to spec and set up physical infrastructure
 - Continuing effort: no need to maintain physical infrastructure
-

Cost:

- Initial capital expenditure
- Costs of over-provisioning for scalability
- Pay for what you use

Slide 7

Reliability:

- Redundancy
 - Trust reliability of provider
 - Data backups
 - *What happens when provider goes down?*
 - *What about Security? Privacy?*
-

INFRASTRUCTURE AS A SERVICE

Service provider provides:

- Server and network hardware
- Virtual machines
- IP addresses
- Services to manage VMs (create, start, stop, migrate)
- Optional: network storage

Slide 8

Client provides:

- OS and OS environment
 - Web server, DBMS, etc.
 - Middleware
 - Application software
-

Challenges – Client:

- transparency (naming, redirection)
- scalability: replication and load balancing decisions
- synchronisation and coordination
- security
- fault tolerance
- software maintenance and sys admin

Slide 9

Challenges – Provider:

- load management
 - IP address management, DNS management
 - infrastructure fault tolerance
 - monitoring, logging, billing
 - storage
-

EXAMPLE 1: EC2

Amazon Web Services:

- Elastic Compute Cloud (EC2)
- Simple Storage Solution (S3)
- Simple DB
- Simple Queue Service

Slide 10

Elastic Compute Cloud:

- virtual cores: 1.7 GB memory, 160 GB of instance storage
 - 10 cents/hour
 - launch Amazon Machine Image (AMI) on virtual cores
 - preconfigured or custom images
-

RELIABILITY

Availability Zones:

- 99.95% availability per service region
- Regions: geographically dispersed
- Availability zones: contained in Regions
- Availability zones: insulated from failures in other zones

Slide 11

Elastic IP addresses:

- IP address associated with account
- dynamic remapping to specific instances

Elastic Load Balancing:

- distributes traffic across instances
 - based on DNS domain names
 - monitors 'health' of instances: customisable
 - routes to healthy instances
-

Auto Scaling:

- Automatically scale capacity up or down
- User-defined conditions

Slide 12

Security:

- firewall settings for instances
-

STORAGE

Elastic Block Store:

- Block level storage volumes
- Mounted as block device on an instance
- Replicated in Availability zone
- Only one instance at a time
- Re-attach on instance failure
- 10 cents per GB per month + 10 cents per million IO requests

Slide 13

Snapshots:

- Point in time copy of EBS volume
 - Stored in S3
 - Differential
 - Can be used to bootstrap image
-

PLATFORM AS A SERVICE

Service provider provides:

- Hardware infrastructure
- OS and platform software
- Distributed storage management
- Load balancing, replication, migration
- Management and Monitoring services

Slide 14

Client provides:

- Application
-

Challenges: Client:

- learn new API and environment
- follow API
- optimise to limits of API and platform
- security for own app

Slide 15

Challenges: Provider:

- transparency (naming, redirection)
 - scalability: replication and load balancing decisions
 - synchronisation and coordination
 - security
 - fault tolerance
 - monitoring
 - software maintenance and sys admin
-

EXAMPLE 2: APP ENGINE



Slide 16

- Python development language
 - Python runtime environment
 - storage based on Big Table
 - optimisation via Memcache
 - lots of APIs
 - per use billing
 - transparent scaling
-

Slide 17

