

CPET 581 Cloud Computing: Technologies and Enterprise IT Strategies

Lecture 6

Cloud Platform Architecture over Virtualized Data Centers Part -3

Public Cloud Platforms: GAE, AWS, and Azure & Inter-Cloud Resource Management

Text Book: Distributed and Cloud Computing, by K. Hwang, G C. Fox,
and J.J. Dongarra, published Elsevier/Morgan Kaufmann, 2012.

Spring 2015

A Specialty Course for Purdue University's M.S. in Technology Graduate
Program: IT/Advanced Computer App Track

Paul I-Hai Lin, Professor

Dept. of Computer, Electrical and Information Technology
Purdue University Fort Wayne Campus

Prof. Paul Lin

1

Ch. 4 - Topics of Discussion

- Cloud Computing and Service Models
- Data-Center Design and Interconnection Networks
- Architectural Design of Computer and Storage Clouds
- **Public Cloud Platforms: Google App Engine, Amazon Web Services and Microsoft Window Azure**
- Inter-Cloud Resource Management
- Cloud Security and Trust Management

Prof. Paul Lin

2

Five Major Cloud Platforms and Their Service Offering

Table 4.5 Five Major Cloud Platforms and Their Service Offerings [30]

Model	IBM	Amazon	Google	Microsoft	Salesforce
PaaS	BlueCloud, WCA, RC2		App Engine (GAE)	Windows Azure	Force.com
IaaS	Ensembles	AWS		Windows Azure	
SaaS	Lotus Live		Gmail, Docs	.NET service, Dynamic CRM	Online CRM, Gifftag
Virtualization		OS and Xen	Application Container	OS level/ Hypel-V	
Service Offerings	SOA, B2, TSAM, RAD, Web 2.0	EC2, S3, SQS, SimpleDB	GFS, Chubby, BigTable, MapReduce	Live, SQL, Hotmail	Apex, visual force, record security
Security Features	WebSphere2 and PowerVM tuned for protection	PKI, VPN, EBS to recover from failure	Chubby locks for security enforcement	Replicated data, rule-based access control	Admin./record security, uses metadata API
User Interfaces		EC2 command-line tools	Web-based admin. console	Windows Azure portal	
Web API	Yes	Yes	Yes	Yes	Yes
Programming Support	AMI		Python	.NET Framework	

Note: WCA: WebSphere CloudBurst Appliance; RC2: Research Compute Cloud; RAD: Rational Application Developer; SOA: Service-Oriented Architecture; TSAM: Tivoli Service Automation Manager; EC2: Elastic Compute Cloud; S3: Simple Storage Service; SQS: Simple Queue Service; GAE: Google App Engine; AWS: Amazon Web Services; SQL: Structured Query Language; EBS: Elastic Block Store; CRM: Consumer Relationship Management.

Platform as a Service (PaaS): Google App Engine

- Allows users to develop and host web applications in Google datacenters with automatic demand-based scaling
- A free service for a certain limit and requires a Gmail account to access the services. After free limits is exceeded, customers are charged for additional storage, bandwidth and instance hours.
- Programming language supports: Java, Python, and Go
- All billed App Engine applications have a 99.99% uptime SLA.
- The app engine has a few restrictions:
 - Can only execute code called from an HTTP request
 - Java application may only use subset from the JRE standard edition
 - Java application cannot create new threads

Example 4.2 Google App Engine for PaaS Applications

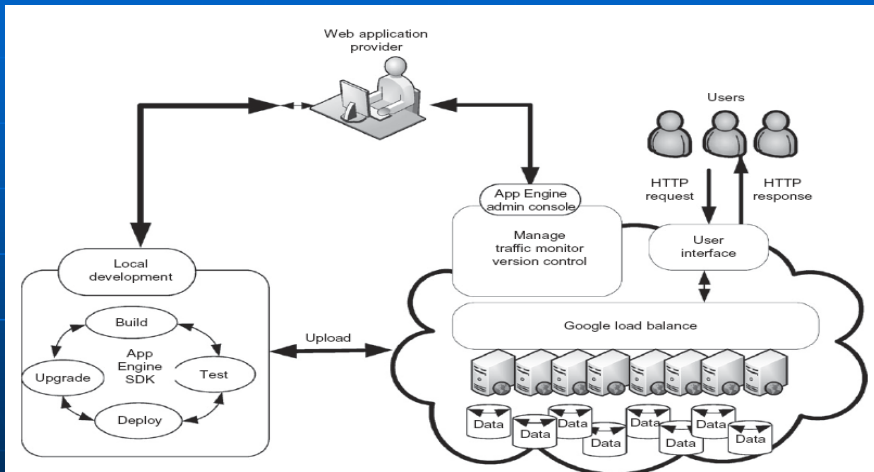


FIGURE 4.7

Google App Engine platform for PaaS operations.

(Courtesy of Yangting Wu, USC)

Functional Modules of Google App Engine (GAE)

- The data store
- The application runtime environment
- The software development kit (SDK)
- The administrative console
- The GAE web service infrastructure

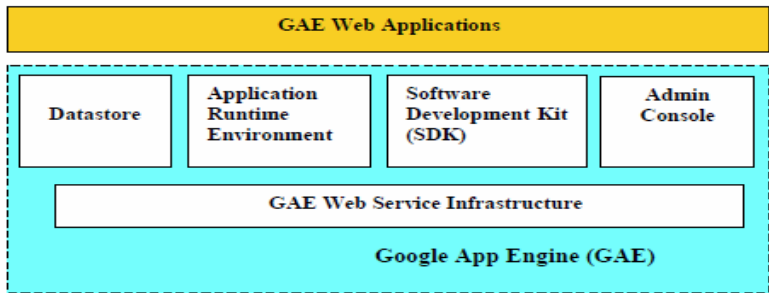


Figure 7.24 Functional components in the Google App Engine (GAE)
 (Courtesy of Google, <http://code.google.com/appengine/>)

Google App Engine (GAE) Architecture: GFS (Google File System), Chubby (distributed locks)

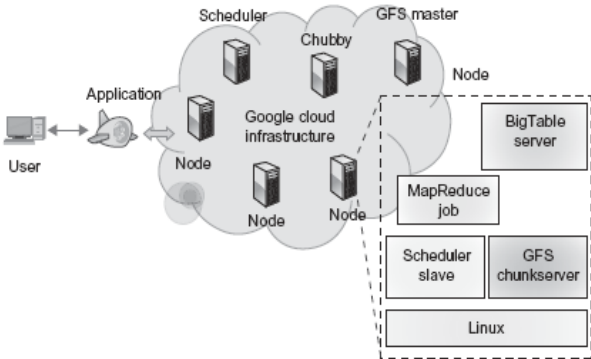


FIGURE 4.20

Google cloud platform and major building blocks, the blocks shown are large clusters of low-cost Servers.
 (Courtesy of Kang Chen, Tsinghua University, China)

Amazon Web Services (AWS) public IaaS services (<http://aws.amazon.com>)

- EC2 (Elastic Compute Cloud)
- S3 (Simple Storage Service)
- EBS (Elastic Block Service)
- Amazon DevPay
- MPI Clusters
- AWS import/export
- Brokering systems
- Small-business companies

Prof. Paul Lin

9

Amazon Web Services (AWS): Key Services .. SQS and SNS (queuing and notification services)




FIGURE 4.21

Amazon cloud computing infrastructure (Key service is identified here; many more are listed in Table 4.5).

(Courtesy of Kang Chen, Tsinghua University, China)

Prof. Paul Lin

10

Amazon Web Services (AWS): Key Services

Compute	Messaging	Storage
Amazon Elastic Compute Cloud (EC2)	Amazon Simple Queue Service (SQS)	Amazon Simple Storage Service (S3)
Amazon Elastic MapReduce	Amazon Simple Notification Service (SNS)	Amazon Elastic Block Storage (EBS)
Auto Scaling		AWS Import/Export
Content Delivery	Monitoring	Support
Amazon CloudFront	Amazon CloudWatch	AWS Premium Support
Database	Networking	Web Traffic
Amazon SimpleDB	Amazon Virtual Private Cloud (VPC)	Alexa Web Information Service
Amazon Relational Database Service (RDS)	Elastic Load Balancing	Alexa Top Sites
E-Commerce	Payments & Billing	Workforce
Amazon Fulfillment Web Service (FWS)	Amazon Flexible Payments Service (FPS)	Amazon Mechanical Turk
	Amazon DevPay	

Microsoft Azure Cloud (PaaS)

- Run its cluster hosted at Microsoft's datacenters that manages computing and storage resources.
- One can download Azure development kit to run a local version of Azure. It allows Azure applications to be developed and debugged on the Windows 7 hosts.
- All cloud services can interact with traditional MS software applications such as Windows Live, Office Live, Exchange Online, etc.
- Azure manages all servers, storage and network resources of the data center.

Microsoft Windows Azure Platform

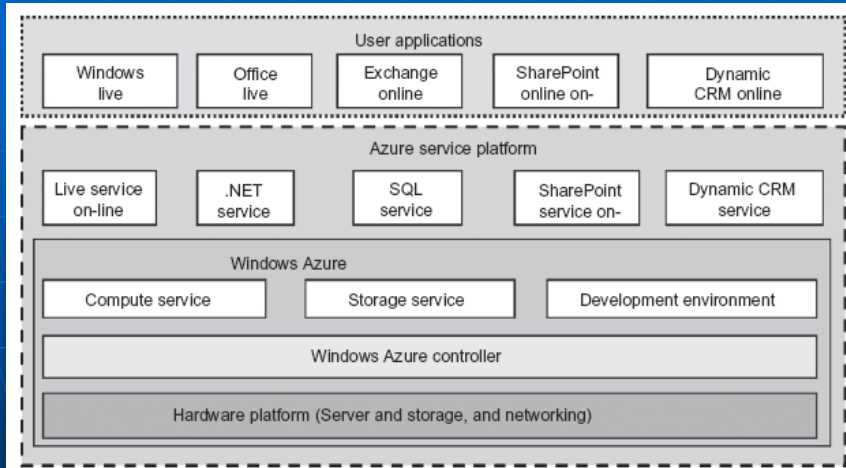


FIGURE 4.22

Microsoft Windows Azure platform for cloud computing.

(Courtesy of Microsoft, 2010, <http://www.microsoft.com/windowsazure>)

4.5 Inter-Cloud Resource Management

4.5.1 Extended Cloud Computing Services

- Fig. 4.23 A stack of six layers of cloud services and their providers: 1) Hardware, 2) Network, 3) Collocation, 4) Infrastructure, 5) Platform, 6) Software applications

Cloud application (SaaS)	Concur, RightNOW, Teleo, Kenexa, Webex, Blackbaud, salesforce.com, Netsuite, Kenexa, etc.			
Cloud software environment (PaaS)	Force.com, App Engine, Facebook, MS Azure, NetSuite, IBM BlueCloud, SGI Cyclone, eBay			
Cloud software infrastructure	Amazon AWS, OpSource Cloud, IBM Ensembles, Rackspace cloud, Windows Azure, HP, Banknorth			
<table border="1" style="display: inline-table; vertical-align: middle;"> <tr> <td>Computational resources (IaaS)</td> <td>Storage (DaaS)</td> <td>Communications (Caas)</td> </tr> </table>	Computational resources (IaaS)	Storage (DaaS)	Communications (Caas)	
Computational resources (IaaS)	Storage (DaaS)	Communications (Caas)		
Collocation cloud services (LaaS)	Savvis, Internap, NTTCommunications, Digital Realty Trust, 365 Main			
Network cloud services (NaaS)	Owest, AT&T, AboveNet			
Hardware/Virtualization cloud services (HaaS)	VMware, Intel, IBM, XenEnterprise			

Amazon's Lesson

- **Down for 3 days since 4/22/2011**
- **1000x of businesses went offline. E.g. Pfizer, Netflix, Quora, Foursquare,Reddit**
- **SLA contract**
 - **99.95% availability (<4.5hour down)**
 - **10% penalty, otherwise**




Copyright © 2012, Elsevier Inc. All rights reserved.

4 - 15

Challenges/Issues in Cloud Computing

Q: Rate the **challenges/issues** ascribed to the 'cloud'/on-demand model
(1=not significant, 5=very significant)



Source: IDC Enterprise Panel, August 2008 n=244

Energy-Efficiency in Google Gmail Services

- <http://www.google.com/green/pdfs/google-green-computing.pdf>
- Clouds win by energy-efficient resource use over datacenters

Business Type	No. of users	No. of servers	IT Power per user	PUE (Power Usage effectiveness)	Total Power per user	Annual Energy per user
Small	50	2	8W	2.5	20W	175 kWh
Medium	500	2	1.8W	1.8	3.2W	28.4 kWh
Large	10000	12	0.54W	1.6	0.9W	7.6 kWh
Gmail (Cloud)	∞	∞	< 0.22W	1.16	< 0.25W	< 2.2 kWh

Copyright © 2012, Elsevier Inc. All rights reserved.

17
4-17

Table 4.7 Cloud Differences in Perspectives of Providers, Vendors, and Users

Cloud Players	IaaS	PaaS	SaaS
IT administrators/cloud providers	Monitor SLAs	Monitor SLAs and enable service platforms	Monitor SLAs and deploy software
Software developers (vendors)	To deploy and store data	Enabling platforms via configurators and APIs	Develop and deploy software
End users or business users	To deploy and store data	To develop and test Web software	Use business software

Table 4.8 Storage Services in Three Cloud Computing Systems

Storage System	Features
GFS: Google File System	Very large sustainable reading and writing bandwidth, mostly continuous accessing instead of random accessing. The programming interface is similar to that of the POSIX file system accessing interface.
HDFS: Hadoop Distributed File System	The open source clone of GFS. Written in Java. The programming interfaces are similar to POSIX but not identical.
Amazon S3 and EBS	S3 is used for retrieving and storing data from/to remote servers. EBS is built on top of S3 for using virtual disks in running EC2 instances.

Conclusion and Summary

Prof. Paul Lin

19