KARPAGAM ACADEMY OF HIGHER EDUCATION
(Deemed to be University)
(Established Under Section 3 of UGC Act, 1956)
Pollachi Main Road, Eachanari Post, Coimbatore – 641 021, Tamil Nadu, India.
Phone: 0422- 2980011-14, Fax No: 0422 – 2980022-23 Email:
info@kahedu.edu.in

SYLLABUS

Semester – VI

17CSU601A	**CLOUD COMPUTING**	**4H – 4C**

SCOPE

The main objective of the course is to portray the recent trends in the field of cloud computing and providing exposures to some open source and commercial clouds.

COURSE OBJECTIVES

- Provide a good understanding of the concepts, standards and protocols in Cloud computing
- Understand various basic concepts related to cloud computing technologies
- Understand the architecture and concept of different cloud models: IaaS, PaaS, SaaS
- Understand different cloud programming platforms and tools

COURSE OUTCOMES

- Develop and deploy cloud application using popular cloud platforms,
- Design and develop highly scalable cloud-based applications by creating and configuring virtual machines on the cloud and building private cloud.
- Write comprehensive case studies analyzing and contrasting different cloud computing solutions

UNIT-I

Overview of Computing Paradigm: Recent trends in Computing: Grid Computing, Cluster Computing, Distributed Computing, Utility Computing, Cloud Computing.Introduction to Cloud Computing: Introduction to Cloud Computing, History of Cloud Computing, Cloud service providers, Benefits and limitations of Cloud Computing.

UNIT-II

Cloud Computing Architecture: Comparison with traditional computing architecture (client/server), Services provided at various levels, Service Models- Infrastructure as a Service(IaaS), Platform as a Service(PaaS), Software as a Service(SaaS), How Cloud Computing Works, Deployment Models- Public cloud, Private cloud, Hybrid cloud, Community cloud, Case study of NIST architecture.

UNIT-III

Case Studies: Case study of Service model using Google App Engine, Microsoft Azure, Amazon EC2 , Eucalyptus.

UNIT-IV

Service Management in Cloud Computing: Service Level Agreements (SLAs), Billing & ACSUounting, Comparing Scaling Hardware: Traditional vs. Cloud, Economics of scaling.

UNIT-V

Cloud Security: Infrastructure Security- Network level security, Host level security, Application level security, Data security and Storage- Data privacy and security Issues, Jurisdictional issues raised by Data location, Authentication in cloud computing.

Suggested Readings

1. Barrie Sosinsky. (2010). Cloud Computing Bible. New Delhi: Wiley-India,
2. Rajkumar Buyya., James Broberg., & Andrzej, M. Goscinski Wile. Cloud Computing: Principles and Paradigms.
3. Nikos Antonopoulos., & Lee Gillam. (2012). Cloud Computing: Principles, Systems and Applications. Springer.
4. Ronald, L. Krutz., & Russell Dean Vines. (2010). Cloud Security: A Comprehensive Guide to Secure Cloud Computing.New Delhi: Wiley-India.
5. Gautam Shroff. (2010). Enterprise Cloud Computing Technology Architecture Applications. Adobe Reader ebooks available from eBooks.com.
6. Toby Velte., Anthony Velte., & Robert Elsenpeter.(2010). Cloud Computing, A Practical Approach. McGraw Hills.
7. Dimitris, N. Chorafas. (2010). Cloud Computing Strategies. CRC Press.

References Books

R1: Dr.Kumar Saurabh (2014), Cloud Computing, 2ed edition, Wiley Pvt.Ltd.New Delhi
R2: Judith Hurwitz, Robin Bloor, Marcia karfman, Fern Halper. (2015), Cloud Computing for Dummies,Wiley Publication, New Delhi.

WEB SITES

1. www.wikipedia.org/wiki/Cloud_computing
2. www.ibm.com/cloud-computing/in/en/
3. www.oracle.com/CloudComputing
4. www.microsoft.com/en-us/cloud/default.aspx

KARPAGAM ACADEMY OF HIGHER EDUCATION

(Deemed to be University Established Under Section 3 of UGC Act 1956)

Coimbatore – 641 021.

LECTURE PLAN

DEPARTMENT OF COMPUTER SCIENCE

STAFF NAME: Dr.V.Sangeetha

SUBJECT NAME: CloudComputing SUB.CODE:17CSU601A

SEMESTER: VI CLASS: III B.Sc (CS)

S.No	Lecture Duration Period	Topics to be Covered	Support Material/Page Nos
		UNIT-I	
1.	1	Recent trends in computing: Grid Computing, cluster computing	w1
2.	1	Distributed Computing,	w1
3.	1	Utility Computing	w1
4.	1	Cloud Computing	w1
5.	1	**Introduction to Cloud Computing:** History of cloud Computing	T1 :4-5
6.	1	Cloud service providers	T1:5-10
7.	1	Benefits and limitations of Cloud Computing.	T1:10-19
8.	1	Recapitulation and Discussion of Important Questions	
Total No Of Periods Planned For Unit 1 :8			
		UNIT-II	
1.	1	Comparison with traditional computing architecture (client/server),	T1 :65,w2
2.	1	Service Models – IaaS,	T1:66 -69
3.	1	Service Models - PaaS, SaaS	T1:70 -75
4.	1	Cloud Computing Works	w2
5.	1	Deployment Models - Public& Private cloud	R1:24-27
6.	1	Hybrid cloud, Community Cloud	R1:28-29
7.	1	Case Study of NIST architecture.	R2:121,167

8.	1	Recapitulation and Discussion of Important Questions	
Total No Of Periods Planned For Unit II :8			
		UNIT-III	
1.	1	Case Study-Service model	W3
2.	1	Exploring Google application	T1:151-161
3.	1	Google App Engine	T1:162-173
4.	1	Microsoft Azure	T1:206-228
5.	1	Amazon EC2	T1:180-190
6.	1	Amazon EC2 – Other IaaS	T1:190-202
7.	1	Eucalyptus	R2:115-116
8.	1	Recapitulation and Discussion of Important Questions	
Total No Of Periods Planned For Unit III :8			
		UNIT-IV	
1.	1	Service Management in Cloud computing	T1:232-237 , w4
2.	1	Service Level Agreements	R1:11, 23
3.	1	Billing & Accounting	R2:237-238
4.	1	Billing & Accounting – Economic Model	R2: 248-250
5.	1	Comparing Scaling Hardware	R2:49-50
6.	1	Traditional vs Cloud	R2:50-52
7.	1	Traditional vs Economics of scaling	R2:52-56
8.	1	Recapitulation and Discussion of Important Questions	
Total No Of Periods Planned For Unit IV:8			
		UNIT-V	
1.	1	Cloud security - Network level security	T1: 250-252
2.	1	Host & Application level security,	T1:252-255
3.	1	Data Security and Storage	R1:85-90
4.	1	Authentication in Cloud computing	T1:78, w4
5.	1	Recapitulation and Discussion of important Questions	

6.	1	Discussion of Previous ESE Question Papers.	
7.	1	Discussion of Previous ESE Question Papers.	
8.	1	Discussion of Previous ESE Question Papers.	
Total No of Periods planned for Unit V: 8			
Total Planned Hours:-40			

Suggested Readings

1. Barrie Sosinsky. (2010). Cloud Computing Bible. New Delhi: Wiley-India,
2. Rajkumar Buyya., James Broberg., & Andrzej, M. Goscinski Wile. Cloud Computing: Principles and Paradigms.
3. Nikos Antonopoulos., & Lee Gillam. (2012). Cloud Computing: Principles, Systems and Applications. Springer.
4. Ronald, L. Krutz., & Russell Dean Vines. (2010). Cloud Security: A Comprehensive Guide to Secure Cloud Computing.New Delhi: Wiley-India.
5. Gautam Shroff. (2010). Enterprise Cloud Computing Technology Architecture Applications. Adobe Reader ebooks available from eBooks.com.
6. Toby Velte., Anthony Velte., & Robert Elsenpeter.(2010). Cloud Computing, A Practical Approach. McGraw Hills.
7. Dimitris, N. Chorafas. (2010). Cloud Computing Strategies. CRC Press.

References Books

R1: Dr.Kumar Saurabh (2014), Cloud Computing, 2ed edition, Wiley Pvt.Ltd.New Delhi
R2: Judith Hurwitz, Robin Bloor, Marcia karfman, Fern Halper. (2015), Cloud Computing for Dummies,Wiley Publication, New Delhi.

WEB SITES

1. www.wikipedia.org/wiki/Cloud_computing
2. www.ibm.com/cloud-computing/in/en/
3. www.oracle.com/CloudComputing
4. www.microsoft.com/en-us/cloud/default.aspx

UNIT-I

Overview of Computing Paradigm: Recent trends in Computing: Grid Computing, Cluster Computing, Distributed Computing, Utility Computing, Cloud Computing. **Introduction to Cloud Computing:** Introduction to Cloud Computing, History of Cloud Computing, Cloud service providers, Benefits and limitations of Cloud Computing.

Introduction

Cloud computing is an emerging technology and has many challenges in various aspects of information handling. Cloud Computing provides us means by which we can access the applications as utilities over the internet. It allows us to create, configure, and customize the business applications online. The term Cloud refers to a Network or Internet. In other words, we can say that Cloud is something, which is present at remote location. Cloud can provide services over public and private networks, i.e., WAN, LAN or VPN.

Applications such as e-mail, web conferencing, customer relationship management (CRM) execute on cloud.

Fig 1: Cloud Computing

Cloud computing offers platform independency, as the software is not required to be installed locally on the PC. Hence, the Cloud Computing is making our business applications mobile and collaborative.

OVERVIEW OF COMPUTING PARADIGM:

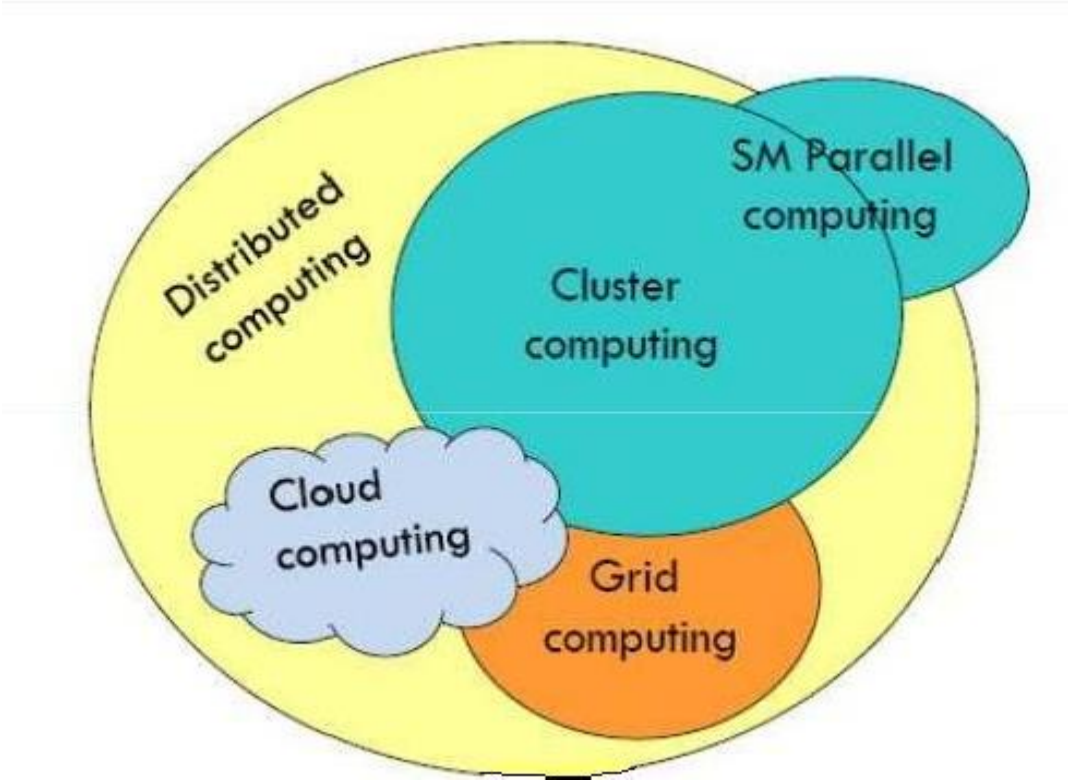

Fig 2: Overview of computing paradigm

What is computing?

- The process of utilizing computer technology to complete a task. Computing may involve computer hardware and/or software, but must involve some form of a computer system.

• **Computing includes -**

- designing,
- developing and building hardware and software systems;
- processing,
- structuring, and managing various kinds of information;
- doing scientific research on and with computers;
- Entertainment media.

1. Grid computing - Grid computing involves connecting geographically remote computers into a single network to create a virtual supercomputer by combining the computational power of all computers on grid.

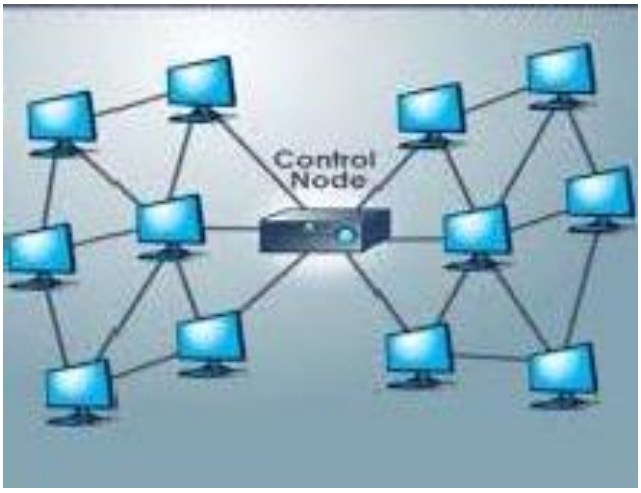

Fig 3: How Grid Computing works

- "the Grid" links together computing resources (PCs, workstations, servers, storage elements) and provides the mechanism needed to access them.
- By implementing our proposed Intranet Grid it is very easy to download multiple files very fast.
- No need to worry about the security as we are authenticating each and every step taking place in our Grid.
- In particular user to access the database. Further implementations could be carried out in the nearest future.

WHY GRID COMPUTING?

- 40% Mainframes are idle
- 90% Unix servers are idle
- 95% PC servers are idle
- 0-15% Mainframes are idle in peak-hour
- 70% PC servers are idle in peak-hour

Advantages

- Can solve larger, more complex problems in a shorter time
- Easier to collaborate with other organizations
- Make better use of existing hardware

Disadvantages

- Grid software and standards are still evolving
- Learning curve to get started
- Non-interactive job submission

Examples:

1. A Scientists studying scientific concepts has the ability to use an entire network of computers in order to analyze data
2. A Business man has the ability to access an entire network of computers in order to forecast the growth of particular stock.

2. Cluster computing

- It is a form of computing in which a group of computers are linked together so they can act like a single entity.
- It is the technique of linking two or more computers into a network (Usually through a local area network) in order to take advantage of the parallel processing power of those computers.

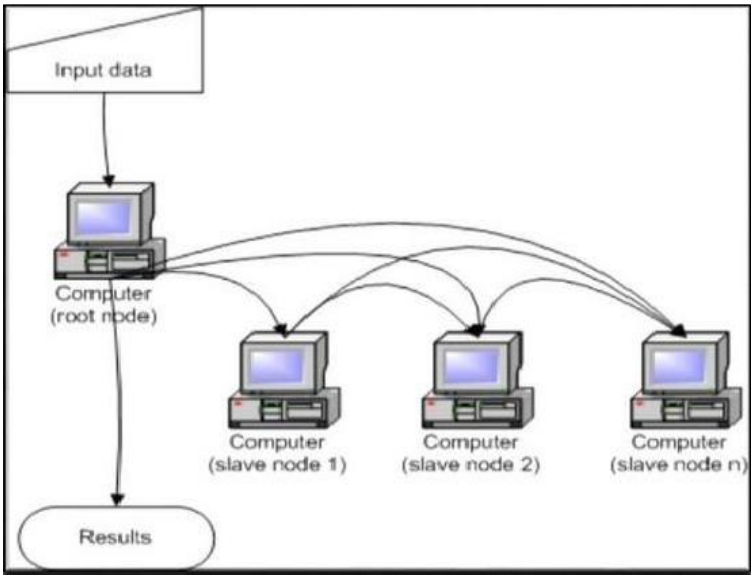

Fig 4: A Simple Cluster Layout

Cluster Application

- Google Search Engine
- Earthquake Simulation Software
- Image Rendering
- Weather Forecasting

Benefits

- High processing power
- Cost Efficient
- High Availability
- Fault Tolerance

Advantages

A computer cluster provides much faster processing speed, larger storage capacity, better data integrity, superior reliability and wider availability of resources.

Disadvantages

Cost is high.....Since clustering needs more servers and hardware to establish one, monitoring and maintenance is hard.

3. Distributed Computing

Distributed computing is the method of processing in which different parts of a program are run simultaneously on two or more computers that are communicating with each other over a network.

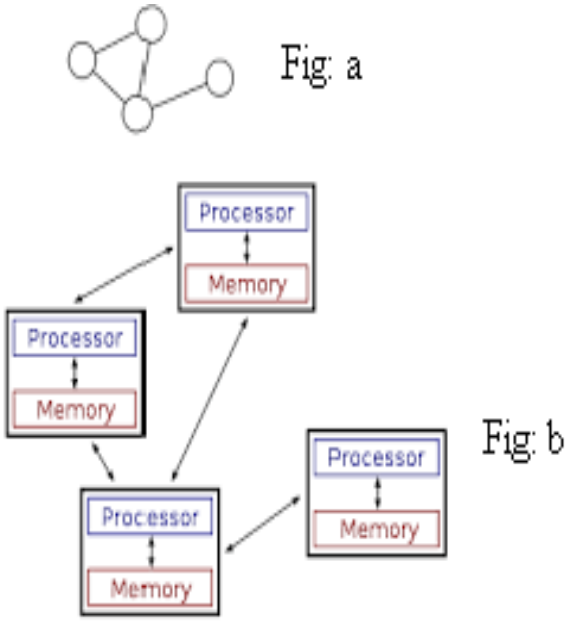

Fig 5: Distributed computing

The figure illustrates the distributed computing systems. Figure (a) is a schematic view of a typical distributed system; the system is represented as a network topology in which each node is a computer and each line connecting the nodes is a communication link. Figure (b) shows the same distributed system in more detail: each computer has its own local memory, and information can be exchanged only by passing messages from one node to another by using the available communication links.

Advantages

- **Inherently Distributed applications**: several applications are inherently distributed in nature and require distributed computing system for their realization

- **Information Sharing among Distributed Users:** In a distributed computing system, information generated by one of the users can be easily and efficiently shared by the users working at other nodes of the system .The use of distributed computing systems by a group of users to work cooperatively is known as computer-supported cooperative working (CSCW), or groupware .
- **Resource Sharing:** Information is not the only thing that can be shared in a distributed computing system. Sharing of software resources such as software libraries and databases as well as hardware resources such as printers, hard disks, and plotters can also be done in a very effective way among all the computers and the users of a single distributed computing system

Disadvantages

- **Complex**: Additional programming required to set up distributed systems
- **Security**: Information passed around the network may be tracked and used for illegal purposes
- **Network dependency**: In case of network failure, the entire system becomes unstable.

4. Utility computing

It is a service provisioning model in which a service provider makes computing resources and infrastructure management available to the customer as needed, and charges them for specific usage rather than a flat rate.

Like other types of on-demand computing (such as grid computing), the utility model seeks to maximize the efficient use of resources and/or minimize associated costs.

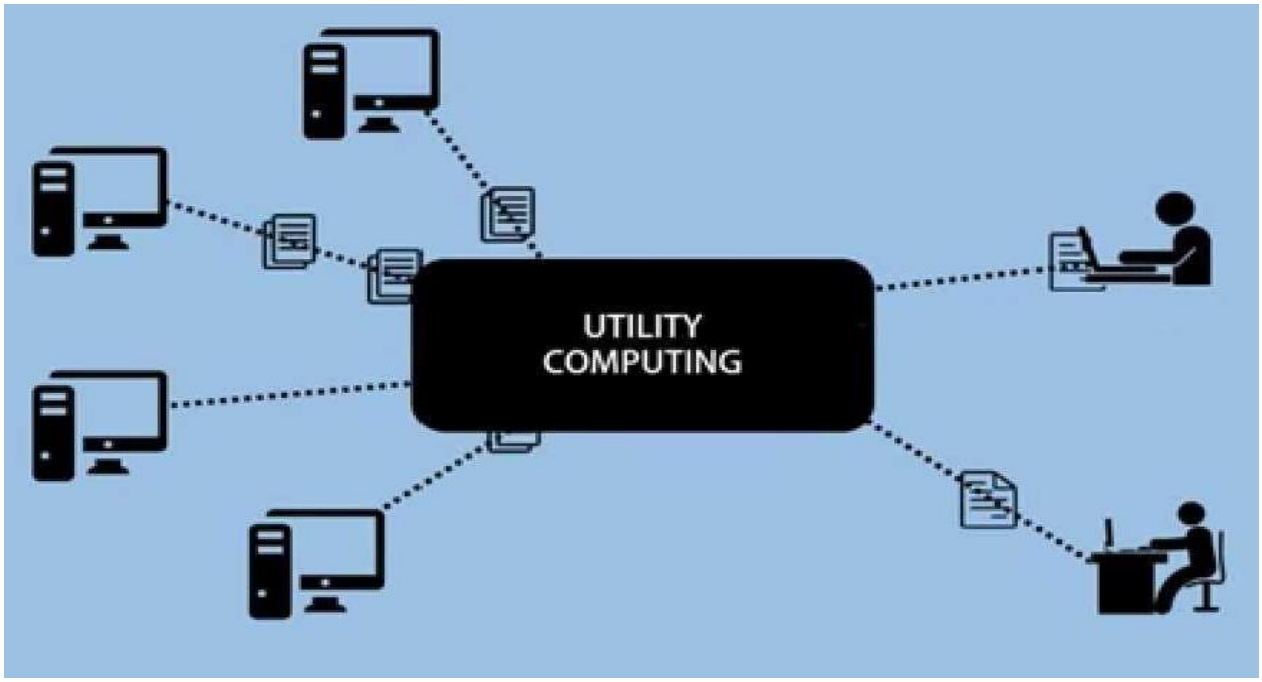

Fig 6: Utility Computing

The Evolution of cloud computing:

With Software as a Service (SaaS) becoming widely accepted, Cloud Computing is becoming increasingly important for small businesses because of its low cost (pay as you go and just for what you need) and agility (rapid ramp up and down).

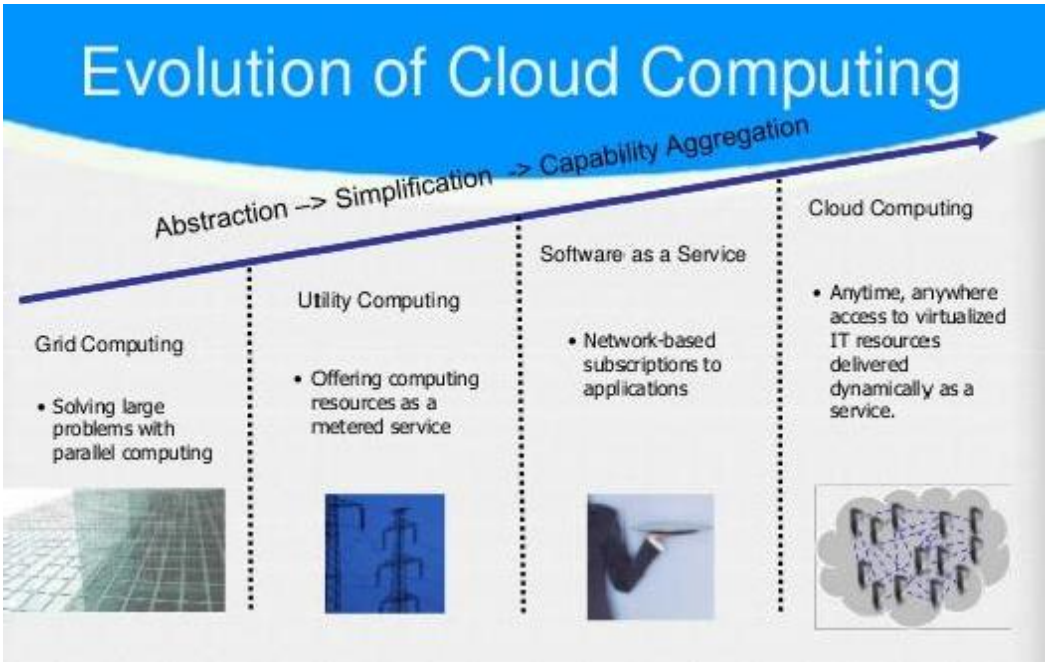

Fig 7: Evolution of Cloud Computing

Business Drivers for Cloud Business growth

1. **Business growth** - is one of the top benefits organisations realise as a result of cloud adoption, with 52% of enterprises reporting increased growth since going cloud (2015 Cloud Enterprise Report).

2. **Efficiency** - is an extremely common cloud driver, with 71% of organisations worldwide ranking it a top area they hope to approve through cloud technology (2015 Cloud Enterprise Report).

3. **Experience** - Next among the business drivers is improving the quality of the customer experience, which 45% of enterprises worldwide rank as a top cloud driver (although that number jumps to 61% looking at only organisations in the UK and Australia).

4. **Assurance** - Finally, there is assurance, which is the idea that data will be more secure in the cloud and the user will attain better uptime because its solutions are maintained by providers that have built their businesses around these competencies

Cloud Computing:

- **In the simplest terms**, cloud computing means storing and accessing data and programs over the Internet instead of your computer's hard drive. The cloud is just a metaphor for the Internet.
- **Definition** - Cloud Computing is the use of hardware and software to deliver a service over a network (typically the Internet). With cloud computing, users can access files and use applications from any device that can access the Internet. An example of a Cloud Computing provider is Google's Gmail. That is Gmail users can access files and applications hosted by Google via the internet from any device.

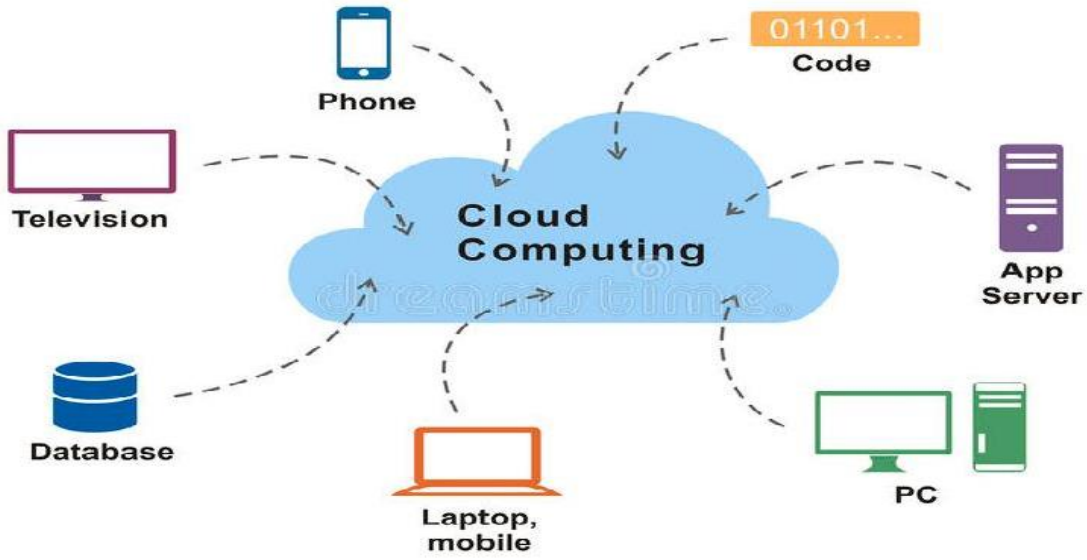

Fig 8: Cloud Computing - Devices

History of Cloud Computing

- **1950s -** In the 50s mainframe computers were huge, occupying entire rooms. Due to the cost of buying and maintaining mainframes, organisations couldn't afford to purchase

one for each user. The solution was "time sharing" in which multiple users shared access to data and CPU time. The term "time sharing" is the premise of cloud computing.

- **1969** - J.C.R. Licklider developed the ARPANET (Advanced Research Projects Agency Network) – the network that became the basis of the internet. His vision was for everyone on the globe to be interconnected and accessing programs and data at any site, from anywhere.
- **1970s** - IBM released an operating system called VM that allowed admins to have multiple virtual systems, or "Virtual Machines" (VMs) on a single physical node. The VM operating system took the 50s "time sharing" model to the next level and most of the basic functions of any virtualisation software that you see nowadays can be traced back to this early VM operating system.
- **1990s** - Telecommunications companies started offering virtualised private network connections, which meant it was possible to allow for more users through shared access to the same physical infrastructure. This change enabled traffic to be shifted as necessary to allow for better network balance and more control over bandwidth usage. Meanwhile, virtualisation for PC-based systems started in earnest, and as the Internet became more accessible, the next logical step was to take virtualisation online.
- **1997**- The term "cloud computing" is coined by University of Texas professor Ramnath Chellappa in a talk on a "new computing paradigm."
- **2002**- Amazon created Amazon Web Services (AWS), providing an advanced system of cloud services from storage to computation.
- **2006** - Amazon introduced the Elastic Compute Cloud (EC2) as a commercial web service. The EC2 let small companies rent computers on which they could run their own computer applications.
- **2009** - Google and Microsoft entered the playing field. The Google App Engine brought low-cost computing and storage services, and Microsoft followed suit with Windows Azure.
- **2010** - The Oneserve field service management software moves to the cloud

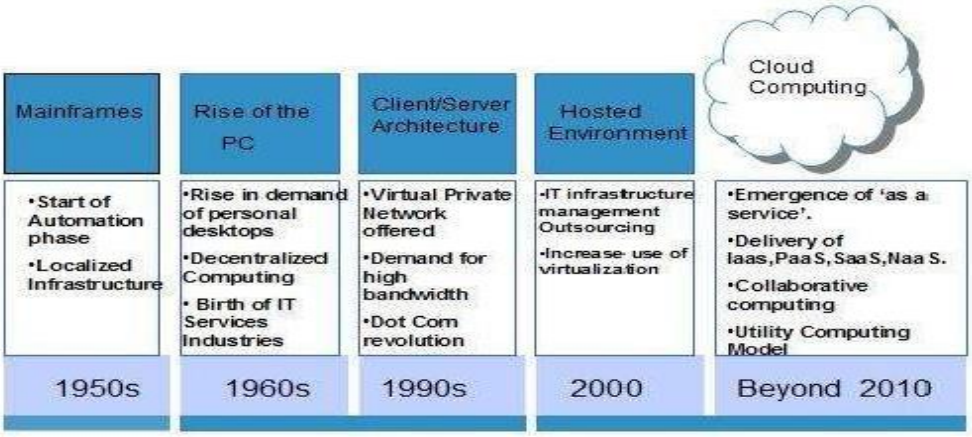

Fig 9: History of Cloud Computing

Challenges in Cloud Computing

Fig 10: Challenges in Cloud Computing

1. Security and Privacy

- Security and privacy are the main challenge in cloud computing.
- These challenges can reduce by using security applications, encrypted file systems, data loss software.

2. Interoperability

- The application on one platform should be able to incorporate services from the other platform. This is known as Interoperability.

- It is becoming possible through web services, but to develop such web services is complex.

3. Portability

- The applications running on one cloud platform can be moved to new cloud platform and it should operate correctly without making any changes in design, coding.
- The portability is not possible, because each of the cloud providers uses different standard languages for their platform.

4. Service Quality

- The Service-Level Agreements (SLAs) of the providers are not enough to guarantee the availability and scalability. The businesses disinclined to switch to cloud without a strong service quality guarantee.

5. Computing Performance

- High network bandwidth is needed for data intensive applications on cloud, this results in high cost.
- In cloud computing, low bandwidth does not meet the desired computing performance.

6. Reliability and Availability

- Most of the businesses are dependent on services provided by third-party, hence it is mandatory for the cloud systems to be reliable and robust.

Features of Cloud Computing

1. **High scalability** - It means on demand provisioning of resources on a large scale without requiring human interaction with each service provider.
2. **High availability and reliability** - Availability of servers is more reliable and high because it minimizes the chances of infrastructure failure.
3. **Agility** - It shares the resources between users and works very quickly.
4. **Multi-sharing** - Multiple user and applications work more efficiently with less cost by sharing common infrastructure using cloud computing.
5. **Maintenance** - Maintenance of cloud computing applications is easier as they are not required to be install on each computer and can also be accessed from various places,

ultimately reducing the cost.

6. **Low cost** - It is cost effective because the company no more needs to set its own infrastructure. It pays according to resources it has consumed.

7. **Services in pay-per-use mode** - APIs(Application Programming Interfaces) are provided to the users for accessing the services on the cloud and pay according to use of the service.

Cloud Service Provider

A cloud service provider, or CSP, is a company that offers some component of cloud computing -- typically infrastructure as a service (IaaS), software as a service (SaaS) or platform as a service (PaaS) to other businesses or individuals.

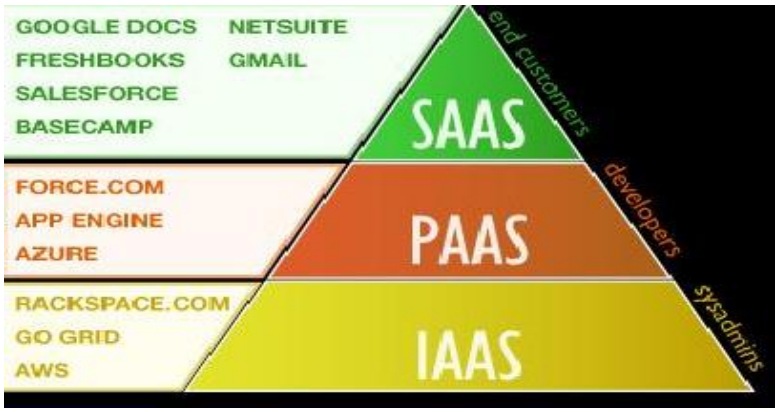

Fig 11: Cloud Service Provider

Examples of cloud service provides are:

1) Amazon Web Service (AWS)
2) Microsoft Azure.
3) Google Cloud Platform.
4) Adobe.
5) VMware.
6) IBM Cloud.
7) Rackspace.
8) Red Hat.

Fig 12: Example of Cloud Services

Cloud Computing Services:

The three major Cloud Computing Offerings are:

- **Software as a Service (SaaS)** - SaaS or software as a service is a software distribution model in which applications are hosted by a vendor or service provider and made available to customers over a network (internet). SaaS is becoming an increasingly prevalent delivery model as underlying technologies that supports Service Oriented Architecture (SOA) or Web Services. Through internet this service is available to users anywhere in the world.

  Traditionaly, software application needed to be purchased upfront &then installed it onto your computer. SaaS users on the other hand, instead of purchasing the software subscribes to it, usually on monthly basis via internet.

  Anyone who needs an access to a particular piece of software can be subscribing as a user, whether it is one or two people or every thousands of employees in a corporation. SaaS is compatible with all internet enabled devices.

  Many important tasks like accounting, sales, invoicing and planning all can be performed using SaaS.

- **Platform as a Service (PaaS)** - Platform as a service, is referred as PaaS, it provides a platform and environment to allow developers to build applications and services. This service is hosted in the cloud and accessed by the users via internet.

To understand in a simple terms, let compare this with painting a picture, where you are provided with paint colors, different paint brushes and paper by your school teacher and you just have to draw a beautiful picture using those tools.

PaaS services are constantly updated & new features added. Software developers, web developers and business can benefit from PaaS. It provides platform to support application development. It includes software support and management services, storage, networking, deploying, testing, collaborating, hosting and maintaining applications.

- **Infrastructure as a Service (IaaS)** - IaaS (Infrastructure As A Service) is one of the fundamental service model of cloud computing alongside PaaS(Platform as a Service). It provides access to computing resources in a virtualized environment "the cloud" on internet. It provides computing infrastructure like virtual server space, network connections, bandwidth, load balancers and IP addresses. The pool of hardware resource is extracted from multiple servers and networks usually distributed across numerous data centers. This provides redundancy and reliability to IaaS.

  IaaS(Infrastructure as a service) is a complete package for computing. For small scale businesses who are looking for cutting cost on IT infrastructure, IaaS is one of the solutions. Annually a lot of money is spent in maintenance and buying new components like hard-drives, network connections, and external storage device etc. which a business owner could have saved for other expenses by using IaaS.

Benefits and limitations of Cloud Computing:

Advantages:

- **Easy implementation** - Cloud hosting allows business to retain the same applications and business processes without having to deal with the backend technicalities. Readily manageable by the Internet, a cloud infrastructure can be accessed by enterprises easily and quickly.
- **Accessibility** - Access your data anywhere, anytime. An Internet cloud infrastructure maximizes enterprise productivity and efficiency by ensuring your application is always

accessible. This allows for easy collaboration and sharing among users in multiple locations. No hardware required. Since everything will be hosted in the cloud, a physical storage center is no longer needed. However, a backup could be worth looking into in the event of a disaster that could leave your company's productivity stagnant.

- **Cost per head** - Overhead technology costs are kept at a minimum with cloud hosting services, enabling businesses to use the extra time and resources for improving the company infrastructure.
- **Flexibility for growth** - The cloud is easily scalable so companies can add or subtract resources based on their needs. As companies grow, their system will grow with them.
- **Efficient recovery** - Cloud computing delivers faster and more accurate retrievals of applications and data. With less downtime, it is the most efficient recovery plan.

Disadvantages:

- **No longer in control** - When moving services to the cloud, you are handing over your data and information. For companies who have an in-house IT staff, they will be unable to handle issues on their own. However, Stratosphere Networks has a 24/7 live help desk that can rectify any problems immediately.
- **May not get all the features** - Not all cloud services are the same. Some cloud providers tend to offer limited versions and enable the most popular features only, so you may not receive every feature or customization you want. Before signing up, make sure you know what your cloud service provider offers.
- **No Redundancy** - A cloud server is not redundant nor is it backed up. As technology may fail here and there, avoid getting burned by purchasing a redundancy plan. Although it is an extra cost, in most cases it will be well worth it.
- **Bandwidth issues** - For ideal performance, clients have to plan accordingly and not pack large amounts of servers and storage devices into a small set of data centers.

Fig 13: Pros /Cons in Cloud Computing

POSSIBLE QUESTIONS

PART A

Q.NO 1 TO 20 (MULTIPLE CHOICE QUESTIONS)

PART B (2 MARKS)

1. Define Cloud Computing
2. Mention any two benefits and limitations in cloud computing.
3. Give few examples for cloud service providers?
4. List out any 5 cloud service providers.
5. Name any two advantages and disadvantage in Grid Computing.
6. List out any two differences between cluster computing and utility computing.
7. Mention any two prons and cons for cloud computing
8. Give a clear structure for overview of cloud computing paradigm
9. What is distributed computing
10. What is evolution of cloud computing.

PART C (6 MARKS)

1. Explain Overview of Computing Paradigm with neat sketch
2. Elucidate History of Computing..
3. Describe the challenges in cloud computing with neat diagram
4. Explain the features of cloud computing.
5. Explain some of the benefits and limitations of cloud computing.

S.NO	Questions
1	_____refers to applications and services that run on a distributed network using virtualized resources and accessed by common Internet protocols and networking standards.
2	The term Cloud refers to _____
3	_____ is a complete operating environment with applications, management, and the user interface
4	The _____is something that you can obtain under contract from your vendor.
5	refers to the components and subcomponents required for Cloud Computing
6	_____ computing connects geographically remote computer into a single network.
7	____ is one of the cloud applications in use.
8	_____ is the technique of linking two or more computers into a network
9	_____ computing process which communicates with each other over a network
10	Google App Engine is an example of _____ services.
11	SQL Azure is an example of _____ services.
12	The Backend platforms are called as_____
13	_____ is taking the physical hardware and going completely virtual

14	Cloud networking is a_____ network
15	In which computing resources the customer can go pay-as-you-go for specific usage.
16	Cloud Providers has _____ main component services in cloud computing

17	The abbreviation for SaaS is _____
18	The abbreviation for IaaS is _____
19	The abbreviation for PaaS is _____
20	_____ abstracts the details of system implementation from users and developers.
21	_____ consists of the particular types of services that you can access on a cloud computing platform.
22	_____is an example of IaaS service providers
23	In which year J.C.R. Licklider developed the ARPANET.
24	All cloud computing applications suffer from the inherent latency that is intrinsic in their _____connectivity
25	Cloud computing is a _____system
26	A single area of concern in cloud computing is_____
27	The use of the word “cloud” makes reference to the____and___essential concepts.
28	_____ provides virtual machines, operating systems, applications, services, development frameworks, transactions, and control structures.
29	Expand EC2
30	_____can be rapidly and elastically provisioned.
31	The abbreviation for AWS is _____
32	_____creates a single point of failure.
33	_____ provides the equivalent of installed applications in the traditional delivery of applications.
34	In which year Amazon introduced the Elastic Compute Cloud (EC2) as a commercial web service.
35	Google App Engine is an example of _____ services.
36	The Backend platforms are called as_____.
37	In which computing paradigm the job submission is non-interactive.

38	. _____is an example of IaaS service providers
39	In case of network failure, the entire system becomes unstable in _____computing
40	_____is one of the large IaaS cloud service providers
41	_____has the least levels of integrated functionality.
42	_____has the most levels of integrated functionality.
43	Expand SLA
44	_____is not a benefit of cloud computing.

opt1	opt2	opt3	opt4	Answer
Cloud Computing	Virtual Computing	Cloud Storage	Cloud Networking	Cloud Computing
data center	data storage	Internet	remote location	Internet
CaaS	PaaS	IaaS	SaaS	SaaS
QoS	QpS	QtS	QaS	QoS
Cloud Computing	Cloud Computing	Cloud Computing	Cloud Based	Cloud Computing Stack
Grid	cluster	Utility	Parallel	Grid
Cloud backup	Cloud storage	Cloud service	Cloud Networking	Cloud backup
Grid	cluster	Utility	Parallel	cluster
Grid	cluster	Utility	Distributed	Distributed
CaaS	PaaS	IaaS	SaaS	PaaS
CaaS	PaaS	IaaS	SaaS	SaaS
Intercloud	Mobile device	Storage	Fat client	Storage
IaaS	PaaS	Daas	SaaS	IaaS

Non Agile	Agile	Latency	Low Latency	Agile
Grid	cluster	Utility	Distributed	Utility
3	2	5	1	3

Software as a Service	Software as a Server	Server as a Service	Structure as a Service	Software as a Service
Infra as a Service	Independent as a Service	Infrastructure as a Service	Infrastructur e service	Infrastructure as a Service
Platform as a Service	Plat as a Service	Platform as a Server	Public as a Service	Platform as a Service
Cloud Computing	Virtual Computing	Cloud Storage	Cloud Networking	Cloud Computing
Development models	Deployment models	Service models	Business models	Service models
Oracle on Demand	GoogleApps	Force.com	Eucalyptus	Eucalyptus
1969	1950s	1970s	1997	1969
MAN	WAN	LAN	LAN & MAN	WAN
stateful	stateup	stateless	statedown	stateless
privacy and network	security and storage	storage and network	privacy and security	privacy and security
Abstraction& Virtualization	Services & applications	Virtualization & Services	Abstraction & applications	Abstraction & Virtualization
IaaS	PaaS	Daas	SaaS	PaaS
Elastic Cloud Compute	Extended Compute	Elastic Compute	Extended Cloud	Elastic Compute Cloud
Data	Network	Information	Resources	Resources
Amazon Web Server	Amazon Web Services	Application Web Services	Amazon Wide	Amazon Web Services
Fat Clients	The Zero Clients	Thick Clients	Cloud Clients	The Zero Clients
IaaS	Daas	SaaS	PaaS	SaaS
1969	2006	2009	2010	2006
CaaS	PaaS	IaaS	SaaS	PaaS
Intercloud	obile device	Storage	Fat client	Storage
Grid	cluster	Utility	Parallel	Grid

Oracle on Demand	GoogleApps	Force.com	Eucalyptus	Eucalyptus
Grid	cluster	Utility	Distributed	Distributed
Rackspace.com	Salesforce.com	GoGrid.com	Openstack.com	Rackspace.com
IaaS	PaaS	Daas	SaaS	IaaS
IaaS	SaaS	Daas	PaaS	SaaS
Strorage Level Agreement	Service Level Agreement	Service Level Applications	Storage Level	Service Level Agreement
Resource pooling	Rapid elasticity	Infinite data	Measured service	Infinite data

KARPAGAM ACADEMY OF	If your application needs large amounts of data transfer, _____ may not be the best model for you.
46	_____cloud is used for healthcare industry
47	_____ is not an opertaion of Quality of Service.
48	A _____ cloud combines multiple clouds are bound together as a unit.
49	Expand VM
50	One of the fundamental components of PaaS middleware is the mapping of _____ onto the cloud infrastructure
51	Cloud computing represents a _____ in the way in which systems are deployed.
52	The scale of cloud computing networks and their ability to provide_____ makes them highly reliable.
53	Cloud computing industry continues to address _____ concerns, if you have an application that works with sensitive data.
54	____is not an architectural standards in Cloud computing.
55	_____is one of the large IaaS cloud service providers
56	IDaaS Stands for_____
57	_____ has a number of operating systems and some enterprise applications that they offer on a rental basis to customers in the form of a number of canned images.
58	The _____ providing applications and enabiling technology, infrastru
59	Expand API .
60	_____is a CPU emulator and virtual machinemonitor

Distributed computing	Load balancing	Virtualization	Cloud computing	Cloud computing
Private cloud	Public cloud	Hybrid cloud	Community cloud	Community cloud
Data replication	Queries	System monitoring	Disaster recovery	Queries
Community cloud	Public cloud	Private cloud	Hybrid cloud	Hybrid cloud
virtual machine	vendor machine	virtual mechanisms	vendor mechanisms	virtual machine
Dynamic applications	Standalone applications	Standard applications	Distributed applications	Distributed applications
Real time applications	Real Paradigm shift	Infinitely Scalable	Measurable Service	Real Paradigm shift
Lower costs	Ease of utilization	Load balancing and failover	Simplified maintenance and upgrade	Load balancing and failover
Security	Privacy	Storage	Bigdata	Security
Grid computing	Distributed computing	Autonomic systems	Standardize d Web services	Distributed computing
Rackspace.com	Salesforce.com	GoGrid.com	Openstack.com	Rackspace.com
Infrastructure as a Service	Independent as a Service	Interdependent as a Service	Identity as a Service	Identity as a Service
Eucalyptus	Amazon	MS Azure	GoGrid	Amazon
vendors	partners	business leaders	GoGrid.com	vendors
Application Programming	App Programming	Application Program	Application Program	Application Programming
Parallels	QEMU	Jumpbox	Vmachines	QEMU

IT Infrastructure

Traditional data centres consist of various pieces of hardware, such as a desktop computer, which are connected to a network via a remote server. This server is typically installed on the premises, and provides all employees using the hardware, access to the business's stored data and applications.

Businesses with this IT model must purchase additional hardware and upgrades in order to scale up their data storage and services to support more users. Mandatory software upgrades are also required with traditional IT infrastructure to ensure fail safe systems are in place to in case a hardware failure occurs. For many businesses with IT data centers, an in-house IT department is needed to install and maintain the hardware.

On the other hand, traditional IT infrastructures are considered to be one of the most secure data hosting solutions and allows you to maintain full control of your company's applications and data on the local server. They are a customized, dedicated system ideal for organisations that need to run many different types of applications.

Cloud Computing Architecture

Cloud Computing architecture comprises of many cloud components, which are loosely coupled. We can broadly divide the cloud architecture into two parts:

- Front End
- Back End

Each of the ends is connected through a network, usually Internet. The following diagram shows the graphical view of cloud computing architecture:

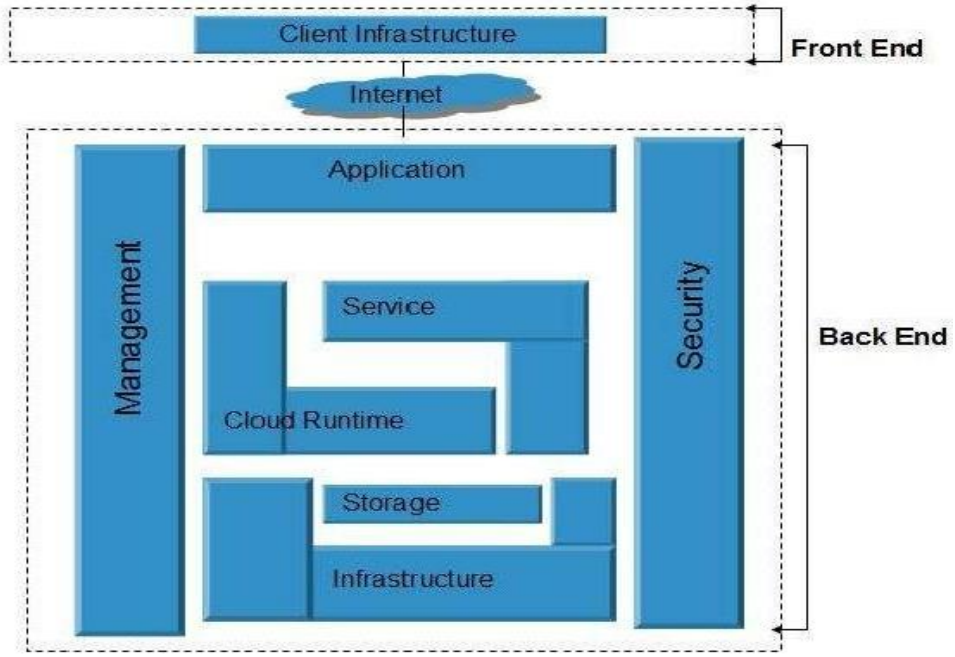

Front End

The front end refers to the client part of cloud computing system. It consists of interfaces and applications that are required to access the cloud computing platforms, Example - Web Browser.

Back End

The back End refers to the cloud itself. It consists of all the resources required to provide cloud computing services. It comprises of huge data storage, virtual machines, security mechanism, services, deployment models, servers, etc.

Cloud Computing vs Traditional IT infrastructure

Cloud computing is far more abstract as a virtual hosting solution. Instead of being accessible via physical hardware, all servers, software and networks are hosted in the cloud, off premises. It's a real-time virtual environment hosted between several different servers at the same time. So rather than investing money into purchasing physical servers in-house, you can rent the data storage space from cloud computing providers on a more cost effective pay-per-use basis.

The main differences between cloud hosting and traditional web hosting are:

Resilience and Elasticity

The information and applications hosted in the cloud are evenly distributed across all the servers, which are connected to work as one. Therefore, if one server fails, no data is lost and downtime is avoided. The cloud also offers more storage space and server resources, including better computing power. This means your software and applications will perform faster.

Traditional IT systems are not so resilient and cannot guarantee a consistently high level of server performance. They have limited capacity and are susceptible to downtime, which can greatly hinder workplace productivity.

Flexibility and Scalability

Cloud hosting offers an enhanced level of flexibility and scalability in comparison to traditional data centres. The on-demand virtual space of cloud computing has unlimited storage space and more server resources. Cloud servers can scale up or down depending on the level of traffic your website receives, and you will have full control to install any software as and when you need to. This provides more flexibility for your business to grow.

With traditional IT infrastructure, you can only use the resources that are already available to you. If you run out of storage space, the only solution is to purchase or rent another server.If you hire more employees, you will need to pay for additional software licences and have these manually uploaded on your office hardware. This can be a costly venture, especially if your business is growing quite rapidly.

Automation

A key difference between cloud computing and traditional IT infrastructure is how they are managed. Cloud hosting is managed by the storage provider who takes care of all the necessary hardware, ensures security measures are in place, and keeps it running smoothly. Traditional data centres require heavy administration in-house, which can be costly and time consuming for your business. Fully trained IT personnel may be needed to ensure regular

monitoring and maintenance of your servers – such as upgrades, configuration problems, threat protection and installations.

Running Costs

Cloud computing is more cost effective than traditional IT infrastructure due to methods of payment for the data storage services. With cloud based services, you only pay for what is used – similarly to how you pay for utilities such as electricity. Furthermore, the decreased likelihood of downtime means improved workplace performance and increased profits in the long run.

With traditional IT infrastructure, you will need to purchase equipment and additional server space upfront to adapt to business growth. If this slows, you will end up paying for resources you don't use. Furthermore, the value of physical servers decreases year on year, so the return on investment of investing money in traditional IT infrastructure is quite low.

Security

Cloud computing is an external form of data storage and software delivery, which can make it seem less secure than local data hosting. Anyone with access to the server can view and use the stored data and applications in the cloud, wherever internet connection is available. Choosing a cloud service provider that is completely transparent in its hosting of cloud platforms and ensures optimum security measures are in place is crucial when transitioning to the cloud.

Public Cloud allows systems and services to be easily accessible to general public. The IT giants such as **Google, Amazon** and **Microsoft** offer cloud services via Internet. The Public Cloud Model is shown in the diagram below.

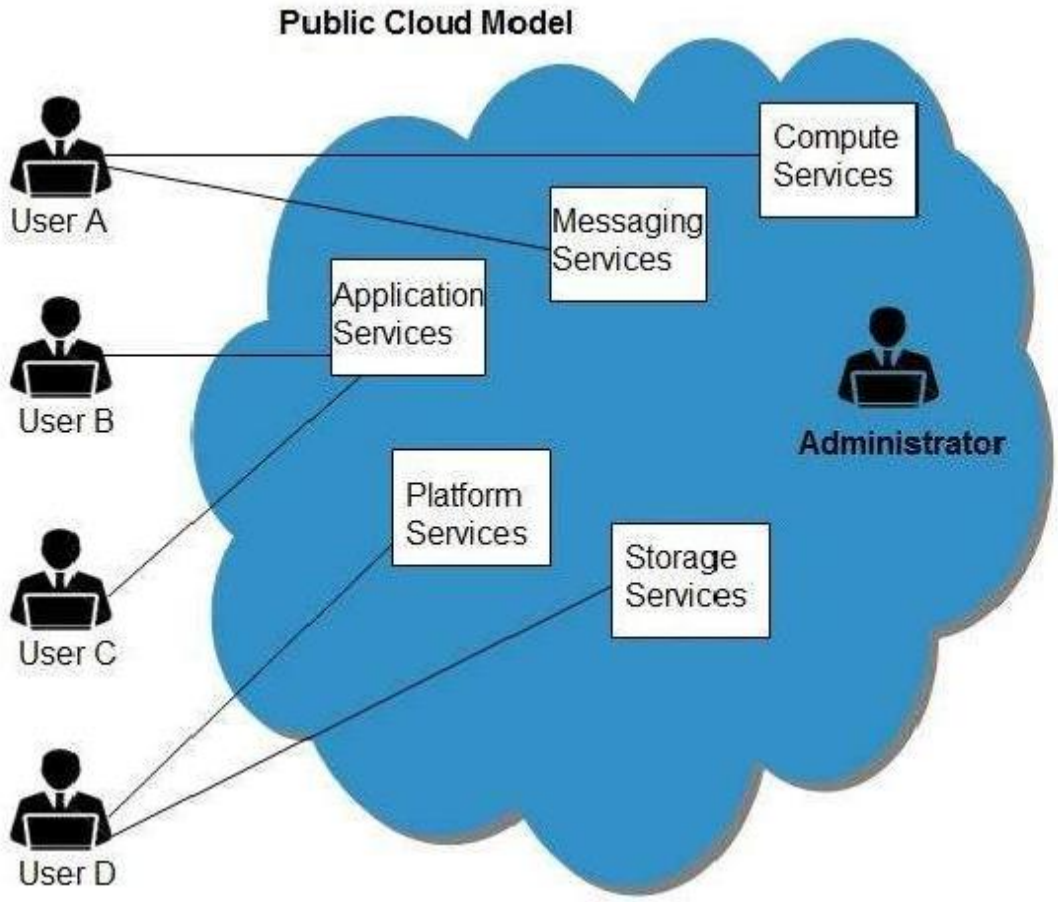

Benefits

There are many benefits of deploying cloud as public cloud model. The following diagram shows some of those benefits:

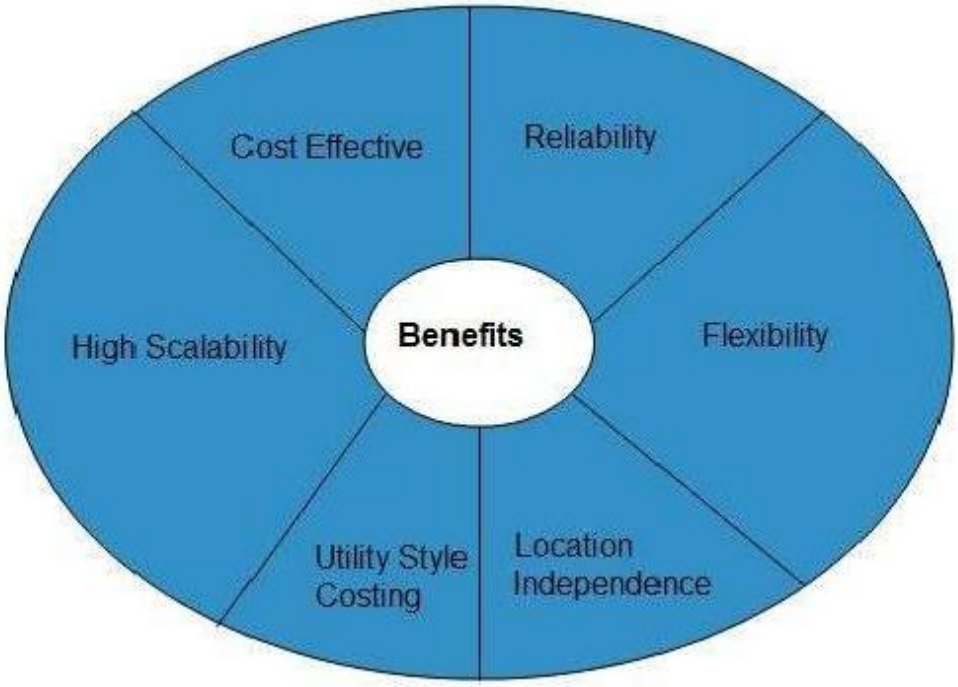

Cost Effective

Since **public cloud** shares same resources with large number of customers it turns out inexpensive.

Reliability

The **public cloud** employs large number of resources from different locations. If any of the resources fails, public cloud can employ another one.

Flexibility

The public cloud can smoothly integrate with private cloud, which gives customers a flexible approach.

Location Independence

Public cloud services are delivered through Internet, ensuring location independence.

Utility Style Costing

Public cloud is also based on **pay-per-use** model and resources are accessible whenever customer needs them.

High Scalability

Cloud resources are made available on demand from a pool of resources, i.e., they can be scaled up or down according the requirement.

Disadvantages

- Here are some disadvantages of public cloud model:

Low Security

- In **public cloud model,** data is hosted off-site and resources are shared publicly, therefore does not ensure higher level of security.

Less Customizable

- It is comparatively less customizable than private cloud.

Private Cloud allows systems and services to be accessible within an organization. The Private Cloud is operated only within a single organization. However, it may be managed internally by the organization itself or by third-party. The private cloud model is shown in the diagram below.

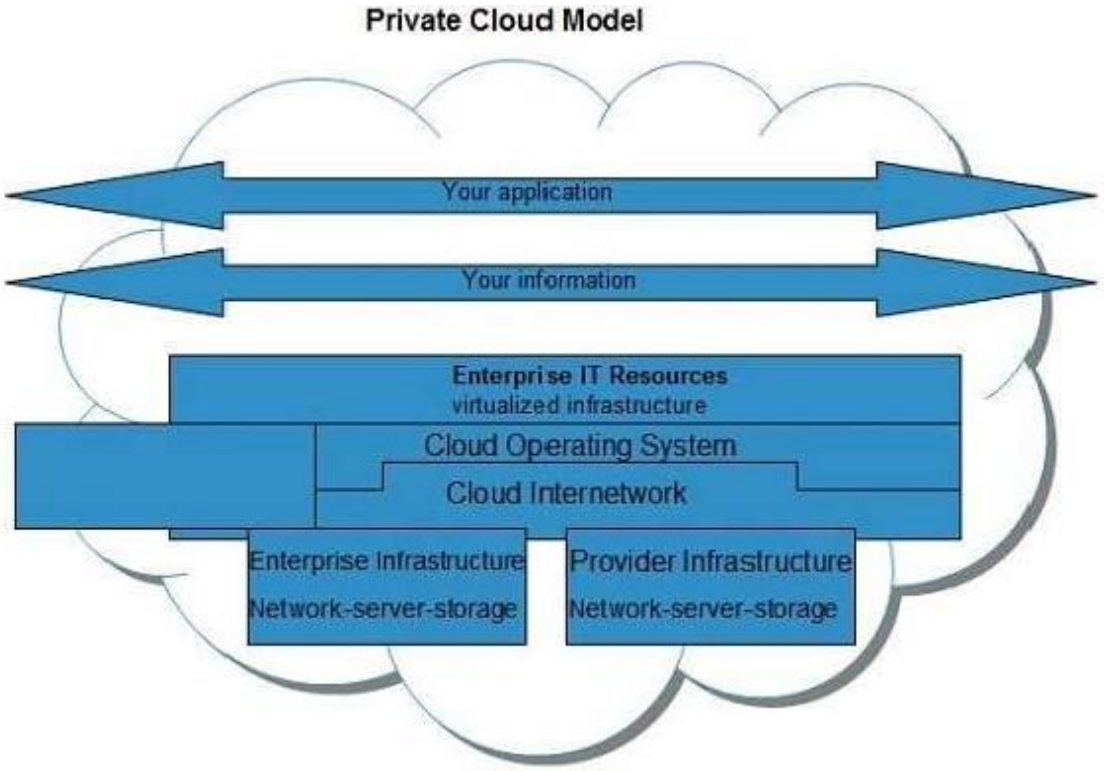

Benefits

There are many benefits of deploying cloud as private cloud model. The following diagram shows some of those benefits:

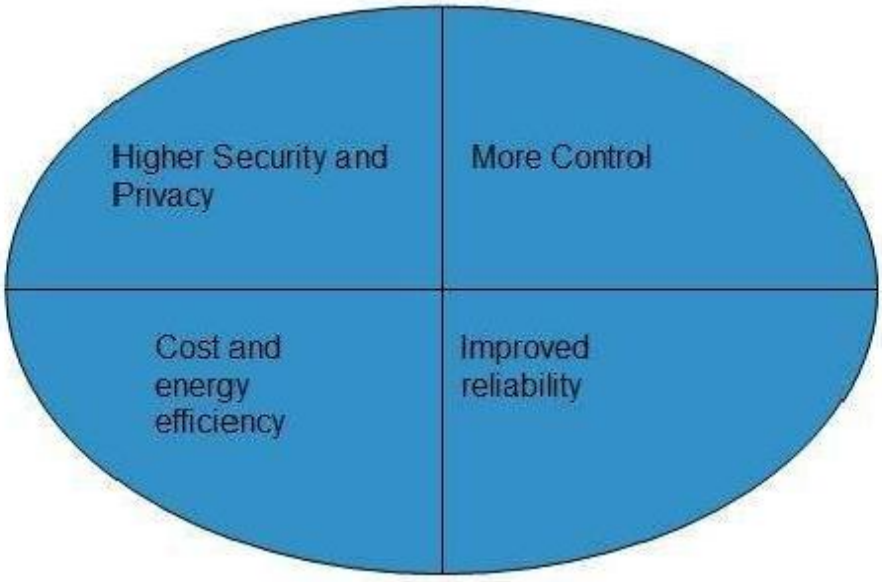

High Security and Privacy

Private cloud operations are not available to general public and resources are shared from distinct pool of resources. Therefore, it ensures high **security** and **privacy.**

More Control

The **private cloud** has more control on its resources and hardware than public cloud because it is accessed only within an organization.

Cost and Energy Efficiency

The **private cloud** resources are not as cost effective as resources in public clouds but they offer more efficiency than public cloud resources.

Disadvantages

Restricted Area of Operation

- The private cloud is only accessible locally and is very difficult to deploy globally.

High Priced

- Purchasing new hardware in order to fulfill the demand is a costly transaction.

Limited Scalability

- The private cloud can be scaled only within capacity of internal hosted resources.

Additional Skills

- In order to maintain cloud deployment, organization requires skilled expertise.

Hybrid Cloud is a mixture of **public** and **private** cloud. Non-critical activities are performed using public cloud while the critical activities are performed using private cloud. The Hybrid Cloud Model is shown in the diagram below.

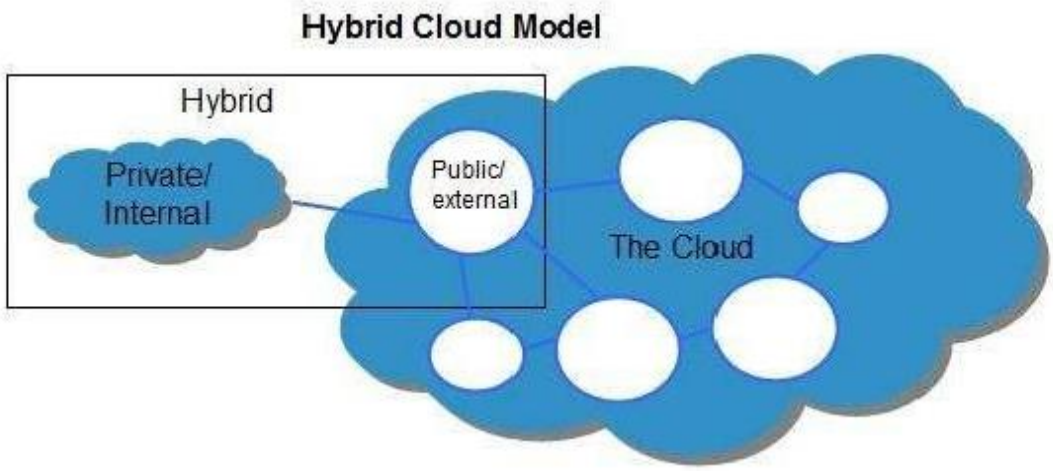

Benefits

There are many benefits of deploying cloud as hybrid cloud model. The following diagram shows some of those benefits:

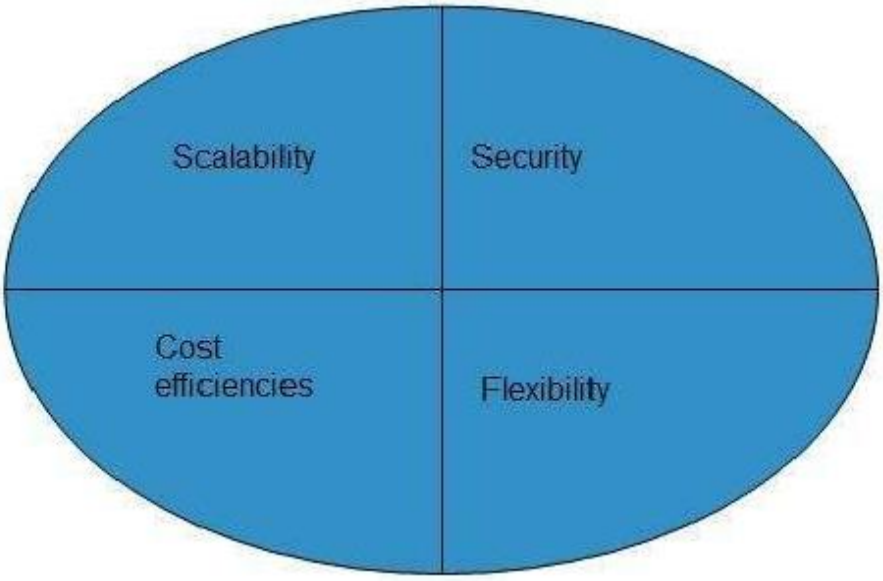

Scalability

- It offers features of both, the public cloud scalability and the private cloud scalability.

Flexibility

- It offers secure resources and scalable public resources.

Cost Efficiency

- Public clouds are more cost effective than private ones. Therefore, hybrid clouds can be cost saving.

Security

- The private cloud in hybrid cloud ensures higher degree of security.

Disadvantages

Networking Issues

- Networking becomes complex due to presence of private and public cloud.

Security Compliance

- It is necessary to ensure that cloud services are compliant with security policies of the organization.

Infrastructure Dependency

- The **hybrid cloud model** is dependent on internal IT infrastructure, therefore it is necessary to ensure redundancy across data centers.

Community Cloud allows system and services to be accessible by group of organizations. It shares the infrastructure between several organizations from a specific community. It may be managed internally by organizations or by the third-party. The Community Cloud Model is shown in the diagram below.

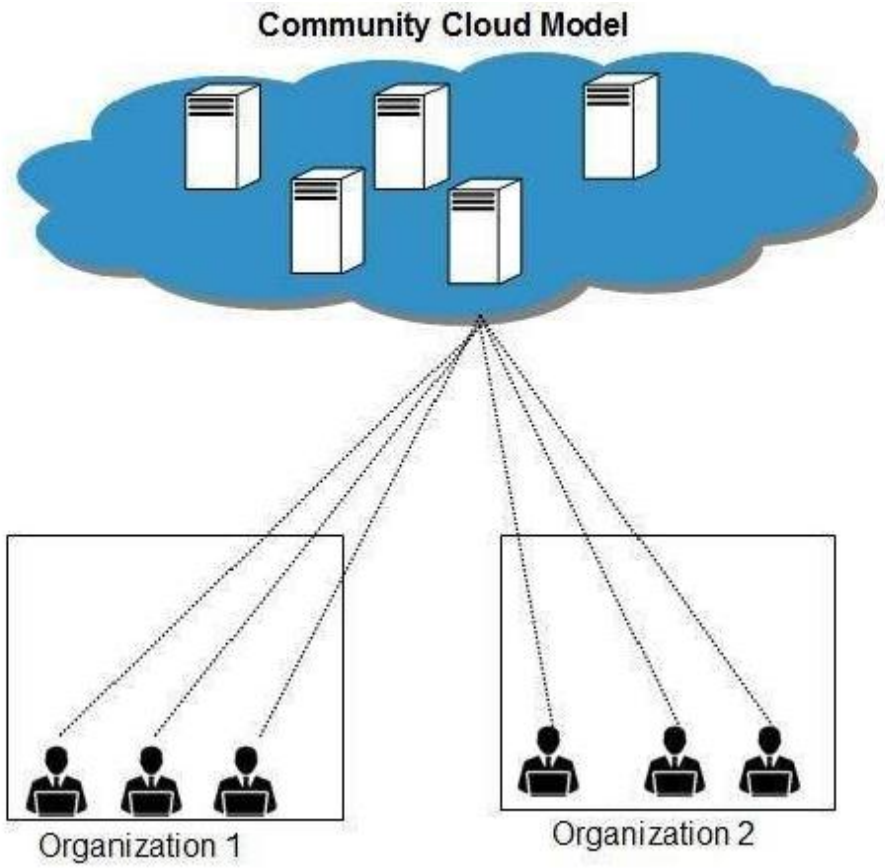

Benefits

There are many benefits of deploying cloud as **community cloud model.**

Cost Effective

- **Community cloud** offers same advantages as that of private cloud at low cost.
- Sharing Among Organizations
- Community cloud provides an infrastructure to share cloud resources and capabilities among several organizations.

Security

- The community cloud is comparatively more secure than the public cloud but less secured than the private cloud.

Issues

- Since all data is located at one place, one must be careful in storing data in community cloud because it might be accessible to others.
- It is also challenging to allocate responsibilities of governance, security and cost among organizations.

NIST ARCHITECTURE

The NIST cloud computing reference architecture is a generic high-level conceptual model that is a powerful tool for discussing the requirements, structures, and operations of cloud computing. The model is not tied to any specific vendor products, services, or reference implementation, nor does it define prescriptive solutions that inhibit innovation. It defines a set

of actors, activities, and functions that can be used in the process of developing cloud computing architectures, and relates to a companion cloud computing taxonomy. It contains a set of views and descriptions that are the basis for discussing the characteristics, uses, and standards for cloud computing.

The NIST cloud computing reference architecture focuses on the requirements of what cloud service provides, not on a design that defines a solution and its implementation. It is intended to facilitate the understanding of the operational intricacies in cloud computing. The reference architecture does not represent the system architecture of a specific cloud computing system; instead, it is a tool for describing, discussing, and developing the system-specific architecture using a common framework of reference..

NIST working group

- Cloud Computing Target Business Use Cases
- Cloud Computing Reference Architecture and Taxonomy
- Cloud Computing Standards Roadmap
- Cloud Computing SAJACC (Standards Acceleration to Jumpstart the Adoption of Cloud Computing)
- Cloud Computing Security

Objectives

The design of the NIST cloud computing reference architecture serves the objectives to: illustrate and understand various cloud services in the context of an overall cloud computing conceptual model; provide technical references to USG agencies and other consumers to understand, discuss, categorize, and compare cloud services; and communicate and analyze

security, interoperability, and portability candidate standards and reference implementations.

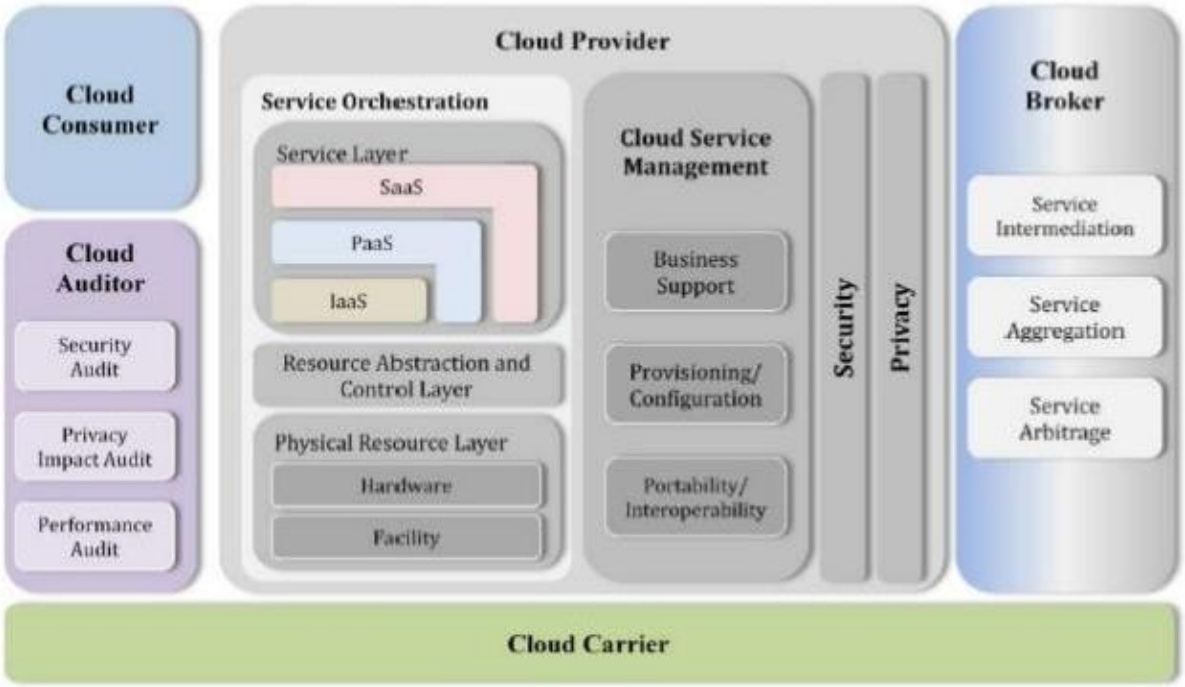

Cloud Computing Reference Architecture

Overview

The Overview of the Reference Architecture describes five major actors with their roles and responsibilities using the newly developing Cloud Computing Taxonomy. **The NIST cloud computing reference architecture defines five major actors: cloud consumer, cloud provider, cloud auditor, cloud broker, and cloud carrier (Cloud Actors).** These core individuals have key roles in the realm of cloud computing. Each actor is an entity (a person or an organization) that participates in a transaction or process and/or performs tasks in cloud computing.

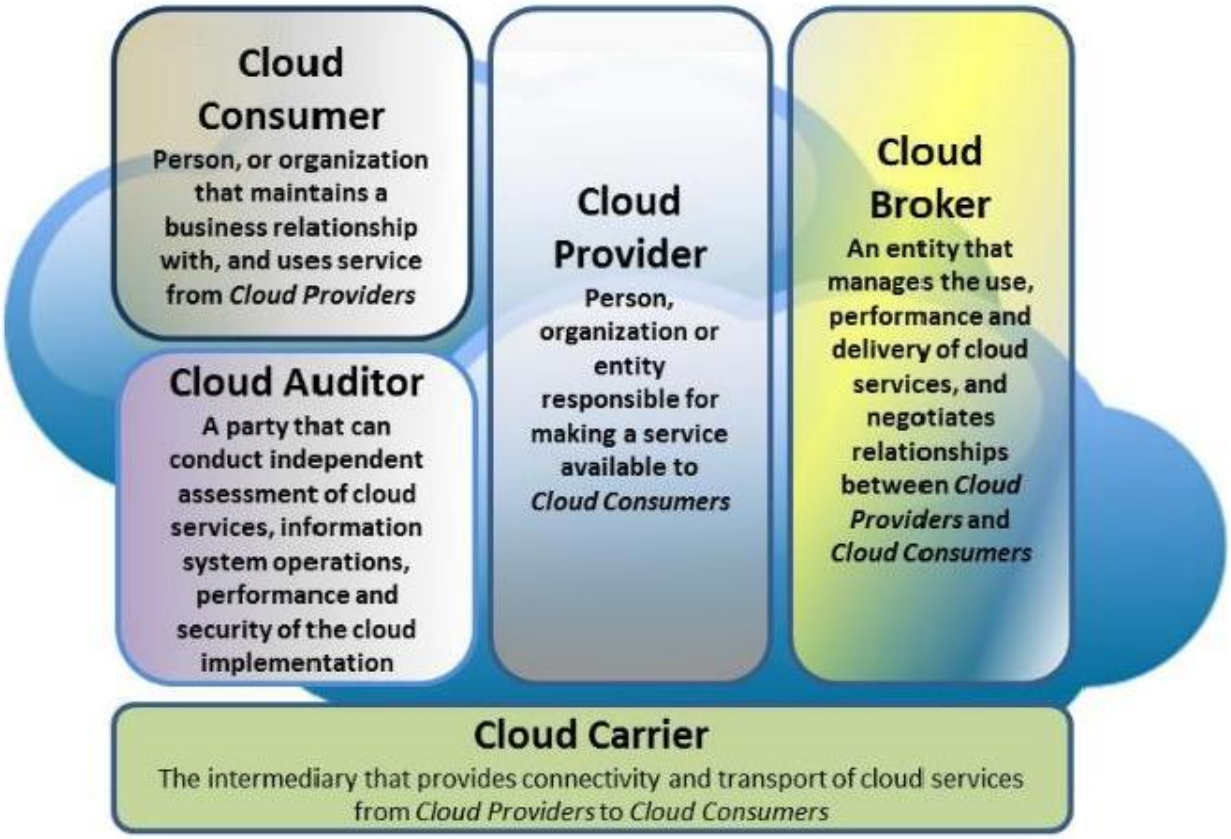

Cloud Actors

An interaction between the Actors in Cloud Computing shows the interactions among the actors in the NIST cloud computing reference architecture. A cloud consumer may request cloud services from a cloud provider directly or via a cloud broker. A cloud auditor conducts independent audits and may contact the others to collect necessary information. The details will be discussed in the following sections and be presented as successive diagrams in increasing levels of detail.

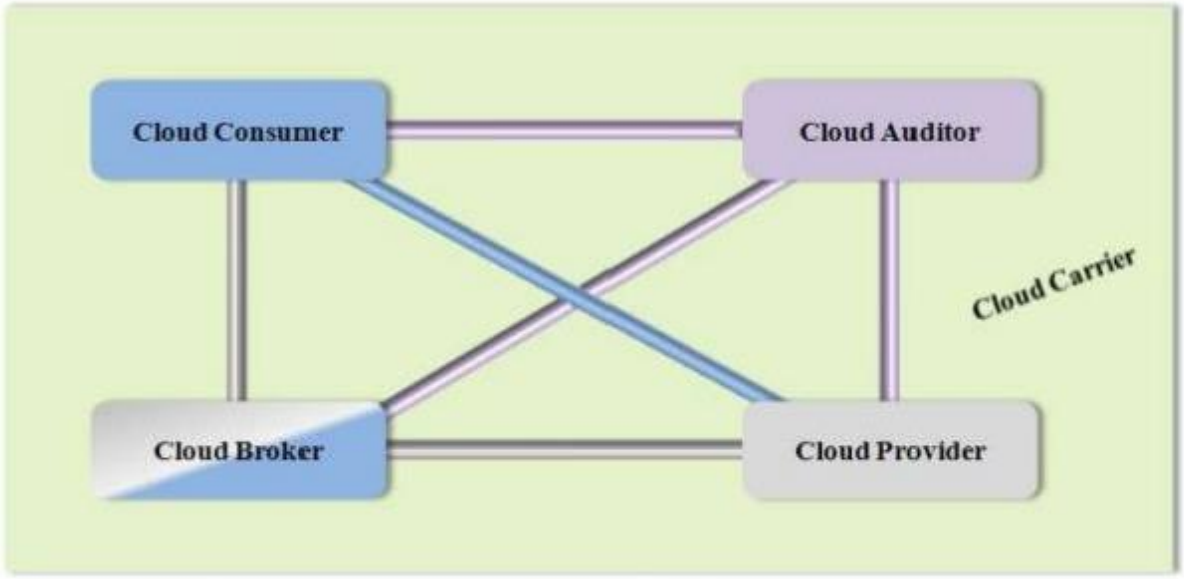

Interactions between the Actors in Cloud Computing

CLOUD CONSUMER

The cloud consumer is the ultimate stakeholder that the cloud computing service is created to support. A cloud consumer represents a person or organization that maintains a business relationship with, and uses the service from, a cloud provider. A cloud consumer browses the service catalog from a cloud provider, requests the appropriate service, sets up service contracts with the cloud provider, and uses the service. The cloud consumer may be billed for the service provisioned, and needs to arrange payments accordingly. Depending on the services requested, the activities and usage scenarios can be different among cloud consumers.

Service Models	Consumer Activities	Provider Activities
SaaS	Uses application/service for Business process operations.	Installs, manages, maintains, and supports the software application on a cloud Infrastructure.

PaaS	Develops, tests, deploys, and manages applications hosted in a Cloud system.	Provisions and manages cloud infrastructure and middleware for the platform consumers; provides development, deployment, and Administration tools to platform consumers.
IaaS	Creates/installs, manages, and monitors services for IT infrastructure operations	Provisions and manages the physical processing, storage, networking, and the hosting environment and cloud Infrastructure for IaaS consumers.

Cloud Consumer and Cloud Provider

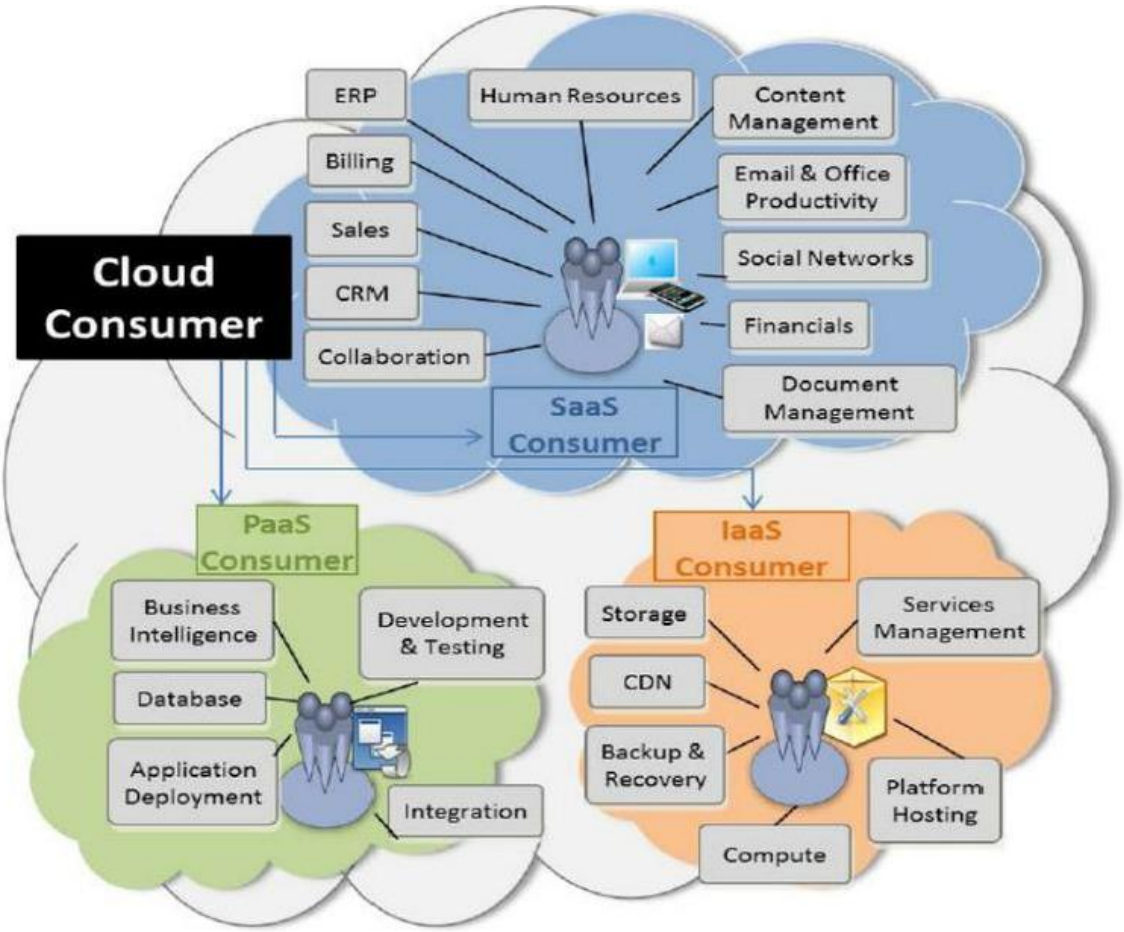

Examples of Services Available to a Cloud Consumer

SaaS applications are usually deployed as hosted services and are accessed via a network

connecting SaaS consumers and providers. The SaaS consumers can be organizations that provide their members with access to software applications, end users who directly use software applications, or software application administrators who configure applications for end users. SaaS consumers access and use applications on demand, and can be billed on the number of consumers or the amount of consumed services. The latter can be measured in terms of the time in use, the network bandwidth consumed, or the amount/duration of data stored.

For PaaS, cloud consumers employ the tools and execution resources provided by cloud providers for the purpose of developing, testing, deploying, and managing applications hosted in a cloud system. PaaS consumers can be application developers who design and implement application software, application testers who run and test applications in various cloud systems, application deployers who publish applications into a cloud system, and application administrators who configure and monitor application performance on a platform. PaaS consumers can be billed by the number of consumers, the type of resources consumed by the platform, or the duration of platform usage.

For IaaS, consumers are provisioned with the capabilities to access virtual computers, network accessible storage, network infrastructure components, and other fundamental computing resources,on which consumers can deploy and run arbitrary software. IaaS consumers can be system developers, system administrators, and information technology (IT) managers who are interested in creating, installing, managing and monitoring services for IT infrastructure operations. IaaS consumers are provisioned with the capabilities to access these computing resources, and are billed for the amount of resources consumed.

Cloud Provider

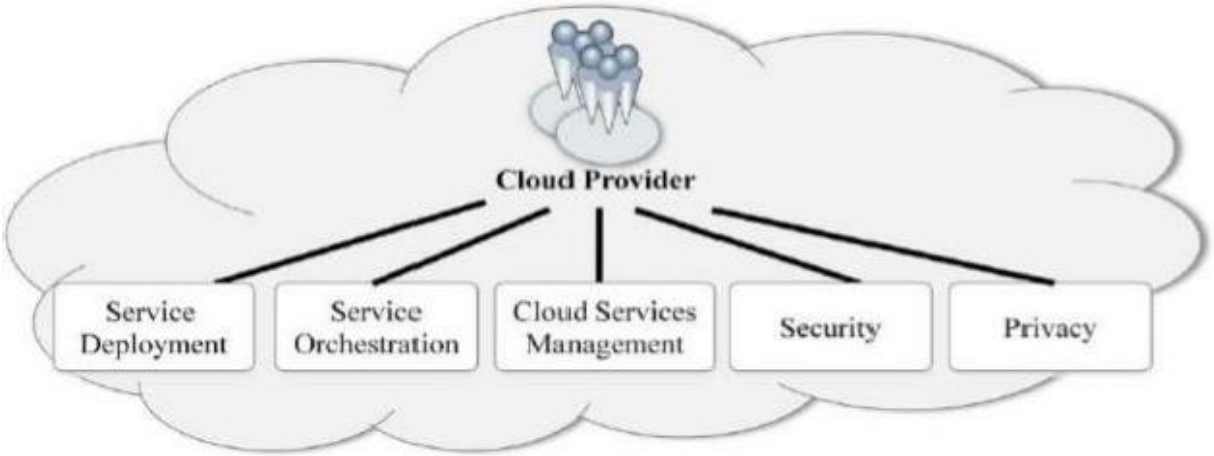

Cloud Provider - Major Activities

A cloud provider can be a person, an organization, or an entity responsible for making a service available to cloud consumers. A cloud provider builds the requested software/platform/ infrastructure services, manages the technical infrastructure required for providing the services, provisions the services at agreed-upon service levels, and protects the security and privacy of the services. Cloud Provider: Major Activities, cloud providers undertake different tasks for the provisioning of the various service models.

For SaaS, the cloud provider deploys, configures, maintains, and updates the operation of the software applications on a cloud infrastructure so that the services are provisioned at the expected service levels to cloud consumers. The provider of SaaS assumes most of the responsibilities in managing and controlling the applications and the infrastructure, while the cloud consumers have limited administrative control of the applications.

For PaaS, the cloud provider manages the cloud infrastructure for the platform, and provisions tools and execution resources for the platform consumers to develop, test, deploy, and administer applications. Consumers have control over the applications and possibly the hosting environment settings, but cannot access the infrastructure underlying the platform including network, servers, operating systems, or storage.

For IaaS, the cloud provider provisions the physical processing, storage, networking, and other fundamental computing resources, as well as manages the hosting environment and cloud infrastructure for IaaS consumers. Cloud consumers deploy and run applications, have more control over the hosting environment and operating systems, but do not manage or control the underlying cloud infrastructure (e.g., the physical servers, network, storage, hypervisors, etc.).

The activities of cloud providers can be discussed in greater detail from the perspectives of Service Deployment, Service Orchestration, Cloud Service Management, Security and Privacy.

As identified in the NIST cloud computing definition, a cloud infrastructure may be operated in one of the following deployment models: public cloud, private cloud, community cloud, or hybrid cloud.

A public cloud is one in which the cloud infrastructure and computing resources are made available to the general public over a public network. A public cloud is owned by an organization selling cloud services and serves a diverse pool of clients.

For private clouds, the cloud infrastructure is operated exclusively for a single organization. A private cloud gives the organization exclusive access to and usage of the infrastructure and computational resources. It may be managed either by the organization or by a third party, and may be implemented at the organization's premise (i.e., on-site private clouds) or outsourced to a hosting company (i.e., outsourced private clouds).

Similar to private clouds, a community cloud may be managed by the organizations or by a third party, and may be implemented at the customer's location (i.e., on-site community cloud) or outsourced to a hosting company (i.e., outsourced community cloud). However, a community cloud serves a set of organizations that have common security, privacy, and compliance considerations, rather than serving a single organization as does a private cloud.

A hybrid cloud is a composition of two or more cloud deployment models (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability. As discussed in this section, both private clouds and community clouds can be either implemented on-site or outsourced to a third party. Therefore, each constituent cloud of a hybrid cloud can be one of the five variants.

SERVICE ORCHESTRATION

Service orchestration refers to the arrangement, coordination, and management of cloud infrastructure to provide the optimizing capabilities of cloud services, as a cost-effective way of managing IT resources, as dictated by strategic business requirements.

The top layer is the service layer, where a cloud provider defines and provisions each of the three service models.

The middle layer is the resource abstraction and control layer. This layer contains the system components that a cloud provider uses to provide and manage access to the physical computing resources through software abstraction.

The lowest layer in the framework is the physical resource layer, which includes all the physical computing resources. This layer includes hardware resources, such as computers (CPU and memory), networks (routers, firewalls, switches, network links, and interfaces), storage components (hard disks), and other physical computing infrastructure elements

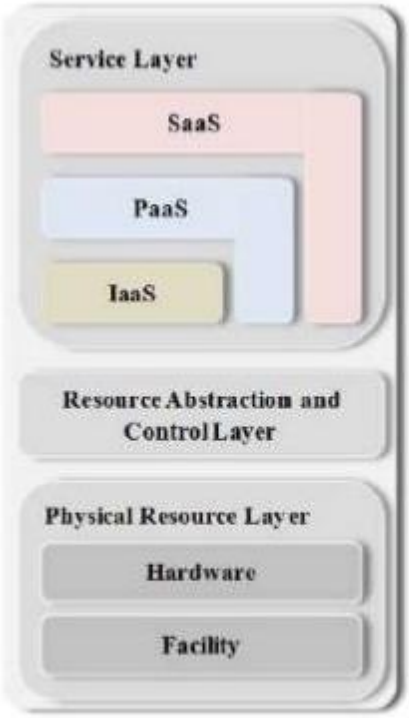

Cloud Provider - Service Orchestration

Cloud Service Management - includes all of the service-related functions that are necessary for the management and operation of those services required by or proposed to cloud consumers. cloud service management can be described from the perspective of business support, provisioning and configuration, and from the perspective of portability and interoperability requirements.

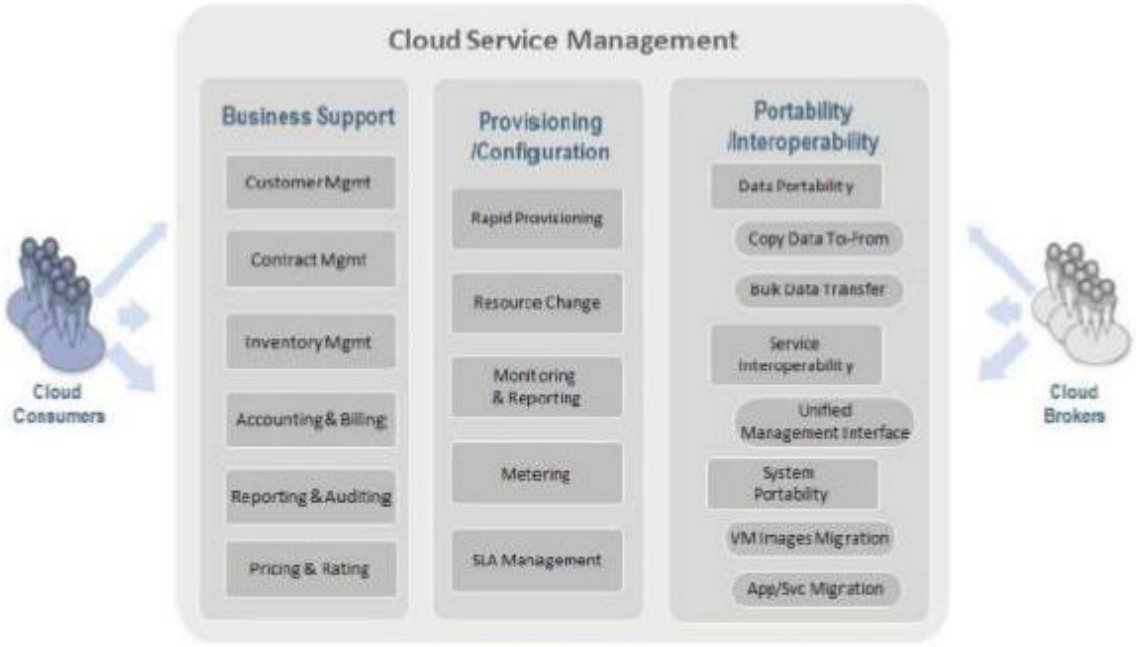

Cloud Provider - Cloud Service Management

Cloud Taxonomy

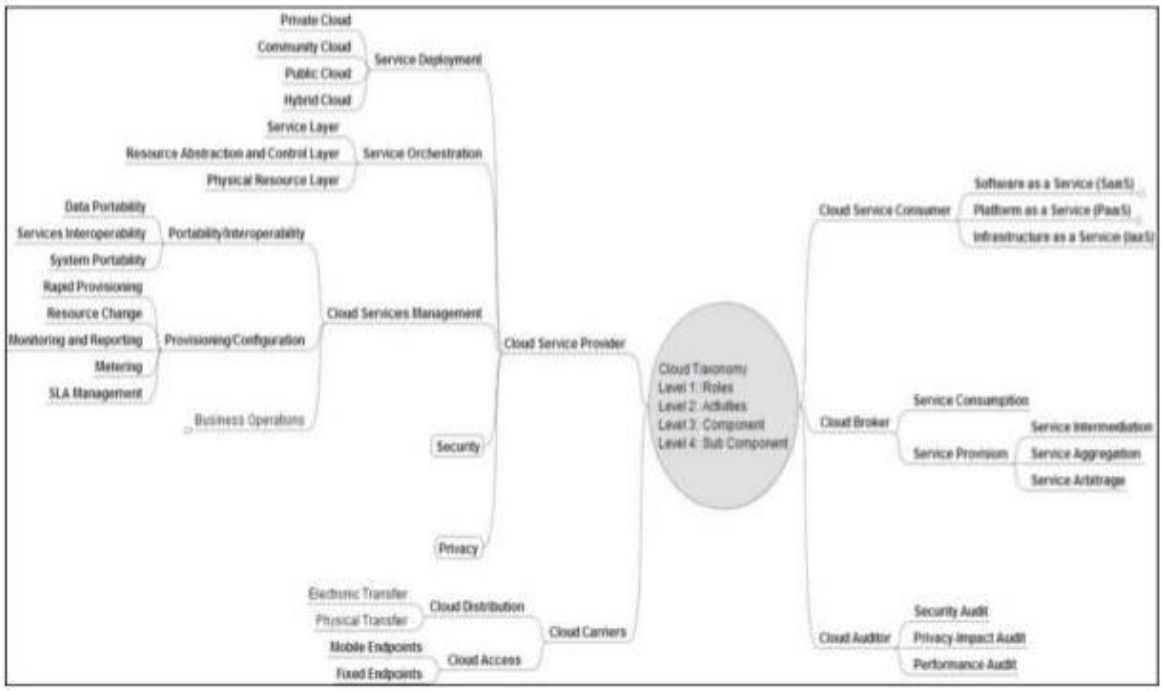

Cloud Taxonomy

POSSIBLE QUESTIONS

PART A

Q.NO 1 TO 20 (MULTIPLE CHOICE QUESTIONS)

PART B (2 MARKS)

1. State the difference between a private cloud and public cloud
2. Compare the characteristics of PaaS and SaaS.
3. What do you mean by SaaS, IaaS, PaaS?
4. Define IaaS.
5. Generalize on PaaS and SaaS..
6. What is Cloud Taxonomy?
7. What is Cloud Service Management in NIST?
8. What does Service Orchestration refer in NIST?
9. List out Cloud Provider major activities.
10. Define NIST Architecture.
11. List out the interactions actors between in cloud computing.

PART C (6 MARKS)

1. Define cloud computing. Enlist and explain different service models.
2. Describe four cloud deployment models with neat diagram and example.
3. Explain the NIST cloud computing reference architecture?
4. Explain in detail software-as- a –service. What are different advantages and disadvantages of SAAS
5. Explain in detail, different implementation level of virtualization?

S.NO	Questions	opt1
1	The _____infrastructure is operated for the exclusive use of an organization.	Public Cloud
2	_____ refers to the location and management of the cloud's infrastructure.	Service models
3	A_____ cloud is one where the cloud has been organized to serve a common function or purpose.	Hybrid cloud
4	A _____ cloud combines multiple clouds are bound together as a unit.	Community cloud
5	_____constitute the first expression of cloud computing	Community Cloud
6	A fundamental characteristic of public clouds is _____	Security
7	_____ is most commonly implemented in PaaS solutions that support hybrid clouds.	Dynamic provisioning
8	_____ are distributed systems created by integrating the services of different clouds to address the specific needs of an industry, a community, or a business sector	Community cloud
9	From an architectural point of view, a _____ is most likely implemented over multiple administrative domains	Community cloud
10	A cloud computing deployment lets someone else manage your computing infrastructure while you manage your business is called _____.	Outsourced IT Staffing
11	____is one of the services that are heavily deployed on cloud comput	VoIP
12	_____ are open systems in which fair competition between different solutions can happen	Community cloud
13	Science clouds are an interesting example of _____	Public cloud
14	_____ are appealing and provide a viable option to cut IT costs and reduce capital expenses.	Private cloud
15	Customer information protection is an aspect of _____	Private cloud

16	In most cases the _____ option prevails because of the existing IT infrastructure.	Public cloud
17	From an architectural point of view, a_____ is most likely implemented over multiple administrative domains.	Community cloud
18	_____cloud is used for healthcare industry.	Private cloud
19	Institutions such as government and military agencies will not consider as an option for processing or storing their sensitive data.	Public cloud
20	_____ is the inability to scale on demand and to efficiently address peak loads	Public cloud
21	address scalability issues by leveraging external resources for exceeding capacity demand	Public cloud
22	_____ share common concerns such as their mission, policies, security, regulatory compliance needs, and so on.	Public cloud
23	_____ are open systems in which fair competition between different solutions can happen.	Public cloud
24	____represents the ability for a cloud service to be widely accessible	Multitenancy
25	We can broadly partition cloud computing into____ layers that form a cloud computing ecosystem.	3
26	The _____ forms the basis for Software as a Service.	Application layer
27	_____ is a cloud computing service model in which hardware is virtualized in the cloud.	Development as a Service
28	The fundamental unit of virtualized client in an IaaS deployment is called a _____	Workload
29	A_____ would reserve a machine equivalent required to run each of these workloads.	Client
30	A group of users within a particular instance is called a _____	Aggregation
31	___are the cloud computing equivalent of compute islands.	Aggregation
32	The one example that is most quoted as a PaaS offering is _____	Force.com
33	The ability to provide storage on demand from a storage pool is referred to as _____	Static provisioning
34	Open file backup systems are _____	Less Expensive
35	Continuous Data Protection (CDP) also known as_____	Porting
36	An _____ allows a system to do what is referred to as a bare metal restore.	Image backup

37	_____is an example of software that supplies Image backup.	Carbonite
38	_____is an example of software that supplies Point-in-time backup.	Carbonite
39	____ virtualizes storage into storage clouds.	SystemGRID
40	____ is a direct competitor to Amazon's S3 service.	Nirvanix
41	Most of the user-based applications that work with cloud storage are of type.	Unmanaged storage
42	_____ offer faster data transfers, but impose additional overhead on clients.	File storage devices
43	FedCloud is used for_____	Banking industry
44	In IaaS, the virtualized resources are mapped to _____	Real systems
45	The work done in IaaS can be measured by the number of _____	Data Per Minute
46	In cloud computing, a provisioned server called an ____ is reserved by a customer.	Input
47	From an architectural standpoint, _____ in an IaaS infrastructure is assigned its own private network.	Server
48	_____ limits broadcast and multicast traffic because Data Link Layer in networking is not supported.	Amazon Web Service's routing
49	Consider a transactional eCommerce system, for which a typical stack contains components.	4
50	Pods are managed by a _____	Google App Engine
51	___ are processing domains that are sealed off from the outside.	Silos
52	_____ can be based on specific types of development languages, application frameworks, or other constructs.	Infrastructure
53	The ____ is responsible for all the operational aspects of the service, for maintenance, and for managing the product(s) lifecycle.	Hardware
54	A developer might write an application in a programming language like ____ using the Google API.	C++
55	Gmail is an offering of _____	DaaS
56	The_____ is available over the Internet globally through a browser on demand.	Software
57	_____ applications feature automated upgrades, updates, and	SaaS

57	patch management and much faster rollout of changes.	SaaS
58	_____ supports multiple users and provides a shared data model through a single-instance, multi-tenancy model.	DaaS
59	_____ is at the heart of the Internet as a service that provides identity authorization and lookup.	Distributed Name Service
60	_____ service currently in beta allows developers to store their data in Google's cloud storage infrastructure.	Atmos

opt2	opt3	opt4	Answer
Private Cloud	Community Cloud	Hybrid Cloud	Private Cloud
Deployment models	Development models	Business models	Deployment models
Community cloud	Private cloud	Public cloud	Community cloud
Public cloud	Private cloud	Hybrid cloud	Hybrid cloud
Private Cloud	Public Cloud	Hybrid Cloud	Public cloud
High bandwidth	QoS	multi tenancy	multi tenancy
Provisioning	Distributed Mapping	Mapping	Dynamic provisioning
Private cloud	Public cloud	Hybrid cloud	Community cloud
Private cloud	Public cloud	Hybrid cloud	Community cloud
Outsourced IT management	QoS	Outsourced IT deployment	Outsourced IT management
IPoV	TCP	UDP	VoIP
Public cloud	Private cloud	Hybrid cloud	Community cloud
Community cloud	Private cloud	Hybrid cloud	Community cloud

Community cloud	Public cloud	Hybrid cloud	Public cloud
Community cloud	Public cloud	Hybrid cloud	Private cloud

Community cloud	Private cloud	Hybrid cloud	Private cloud
Public cloud	Hybrid cloud	Private cloud	Community cloud
Public cloud	Hybrid cloud	Community cloud	Community cloud
Hybrid cloud	Private cloud	Community cloud	Public cloud
Hybrid cloud	Private cloud	Community cloud	Private cloud
Hybrid cloud	Private cloud	Community cloud	Hybrid cloud
Hybrid cloud	Community cloud	Private cloud	Community cloud
Hybrid cloud	Community cloud	Private cloud	Community cloud
n-Demand Usag	Ubiquitous Access	Resiliency	Ubiquitous Access
4	5	6	4
Datalink layer	Platform layer	Network layer	Application layer
Platform as a Service	Infrastructure as a Service	Software as a Service	Infrastructure as a Service
Scheduling	Infrastructure	Operating System	Workload
Server	Network	Host	Client
Silos	Pod	Network	Pod
Silos	Pod	Network	Silos
Amazon Web Services	Google's App Engine platform	Quickbase	Google's App Engine platform
Dynamic provisioning	Thick provisioning	Thin provisioning	Thin provisioning
Expensive	Exchangable	Non Exchangable	Expensive
Imaging	Mirroring	Manipulating	Mirroring
Incremental backup	Differential backup	Point-in-time backup	Image backup

Ghost	Apple's Time Machine	SQL Server	Ghost
Ghost	Apple's Time Machine	SQL Server	Carbonite
StorageGRID	DataGRID	GoGRID	SystemGRID
Iron Mountain	Rackspace	EMC Atmos	Rackspace
Managed storage	Web storage	Rack storage	Unmanaged storage
Block storage devices	Network Attached Storage device	Web storage devices	Block storage devices
Medical industry	Government	Merchant transactions	Government
Virtual systems	Sophisticated systems	Dynamic systems	Real systems
Process Per Minute	Transactions Per Minute	Clients Per Minute	Transactions Per Minute
Instance	Application	Output	Instance
Client	New User	Host	Client
Google's App Engine routing	Quickbase routing	Rackspace routing	Amazon Web Service's routing
3	5	6	5
Cloud Storage System	Amazon Web Services	Cloud Control System	Cloud Control System
Aggregation	Pod	Network	Silos
Softwares	Platforms	Compliance	Platforms
Vendor	Software	Infrastructure	Vendor
C	Python	Pascal	Python
SaaS	PaaS	CaaS	PaaS
Hardware	Platform	Compliance	Software
DaaS	PaaS	CaaS	SaaS

DaaS	PaaS	CaaS	SaaS
PaaS	SaaS	IaaS	SaaS
Distributed Name System	Domain Name Service	Domain Name Solutions	Domain Name Service
Platypus	Eucalyptus	Nirvanix	Platypus

UNIT -III

Case Study on Open Source and Commercial Clouds: Microsoft Azure- Amazon EC2-Google Web services – Open Nebula. - Eucalyptus

Case Study on Open Source and Commercial Clouds:

Microsoft Azure

Microsoft Azure is an ever-expanding set of cloud services to help your organisation meet your business challenges. It is the freedom to build, manage and deploy applications on a massive, global network using your favorite tools and frameworks.

The Essentials

Amazon's AWS has a range of offerings that fall under IaaS, and each of these is categorized into four classes:

- content delivery and storage,
- compute,
- networking, and
- database.

No matter which IaaS offering you get, you will be using Amazon's identity and security services such as AWS CloudHSM's key storage service and Amazon's own Active Directory. Not only that, but AWS offerings also have a range of management tools that users can use, including AWS Config, AWS Cloudtrail, and Cloudwatch.

Azure, on the other hand, also has four classes of offerings:

- Data management and databases,
- compute,
- networking, and
- performance.

Security and management tools include Active Directory Federation Services, Azure Active Directory, Multi-Factor Auth, among others, as well as a range of integrations for Azure monitoring and performance tweaks.

Azure has multiple app deployment options for developers. Including App Services, Cloud Services, Service Fabric, Container Service, Functions, Batch, WebJobs and more. No matter what type of application you are developing, Microsoft has great tools in place to help deploy and scale it.

AWS offers similar solutions with Container Service, Elastic Beanstalk, Lambda, and Batch. AWS does not have as many options or features on the app hosting side. Microsoft has flexed their knowledge of developer tools to have a little bit of an advantage for hosting cloud apps.

Containers seem to be the preferred mechanism to deploy apps in the future, especially for open source applications. Look for more and more advancements in hosting containerized apps in the cloud. **Hybrid clouds are easier with Azure**, partly because Microsoft has foreseen the need for hybrid clouds early on. Azure offers substantial support for hybrid clouds, where you can use your onsite servers to run your applications on the Azure Stack. You can even set your computer resources to tap cloud-based resources when necessary.

Amazon EC2

Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.

Features of Amazon EC2

Amazon EC2 provides the following features:

- Virtual computing environments, known as *instances*
- Preconfigured templates for your instances, known as *Amazon Machine Images (AMIs)*, that package the bits you need for your server (including the operating system and additional software)

- Various configurations of CPU, memory, storage, and networking capacity for your instances, known as *instance types*
- Secure login information for your instances using *key pairs* (AWS stores the public key, and you store the private key in a secure place)
- Storage volumes for temporary data that's deleted when you stop or terminate your instance, known as *instance store volumes*
- Persistent storage volumes for your data using Amazon Elastic Block Store (Amazon EBS), known as *Amazon EBS volumes*
- Multiple physical locations for your resources, such as instances and Amazon EBS volumes, known as *regions* and*Availability Zones*
- A firewall that enables you to specify the protocols, ports, and source IP ranges that can reach your instances using*security groups*
- Static IPv4 addresses for dynamic cloud computing, known as *Elastic IP addresses*
- Metadata, known as *tags*, that you can create and assign to your Amazon EC2 resources
- Virtual networks you can create that are logically isolated from the rest of the AWS cloud, and that you can optionally connect to your own network, known as *virtual private clouds* (VPCs)

Using Google Web Services

Google is the prototypical cloud computing services company, and it supports some of the largest Web sites and services in the world. Google uses automated technology to index the Web. It makes its search service available to users as a standard search engine and to developers as a collection of special search tools limited to various areas of content. The application of Google's searches to content aggregation has led to enormous societal changes and to a growing trend of disintermediation.

The most important commercial part of Google's activities is its targeting advertising business: AdWords and AdSense. Google has developed a range of services including Google

Analytics that supports its targeted advertising business.

Google applications are cloud-based applications. The range of application types offered by Google spans a variety of types: productivity applications, mobile applications, media delivery, social interactions, and many more.

Exploring Google Applications

The bulk of Google's income comes from the sales of target advertising based on information that Google gathers from your activities associated with your Google account or through cookies placed on your system using its AdWords system. The company is highly profitable, and that has allowed Google to create a huge infrastructure as well as launch many free cloud-based applications and services that this chapter details. These applications are offered mostly on a free usage model that represents Google's Software as a Service portfolio. A business model that offers cloud-based services for free that are ―good enough‖ is very compelling. While Google is slowly growing a subscription business selling these applications to enterprises, its revenue represents only a small but growing part of Google's current income.

Google's cloud computing services falls under two umbrellas. The first and best-known offerings are an extensive set of very popular applications that Google offers to the general public. These applications include Google Docs, Google Health, Picasa, Google Mail, Google Earth, and many more.

Google's cloud-based applications have put many other vendors' products—such as office suites, mapping applications, image-management programs, and many other categories of traditional shrink-wrapped software—under considerable pressure.

The second of Google's cloud offerings is its Platform as a Service developer tools. In April 2008, Google introduced a development platform for hosted Web applications using Google's

infrastructure called the Google App Engine (GAE). The goal of GAE is to allow developers to create and deploy Web applications without worrying about managing the infrastructure necessary to have their applications run. GAE applications may be written using many high-level programming languages (most prominently Java and Python) and the Google App Engine Framework, which lowers the amount of development effort required to get an application up and running. Goggle also allows a certain free level of service so that the application must exceed a certain level of processor load, storage usage, and network bandwidth (Input/Output) before charges are assessed.

More Google Products equals fewer commercial products.

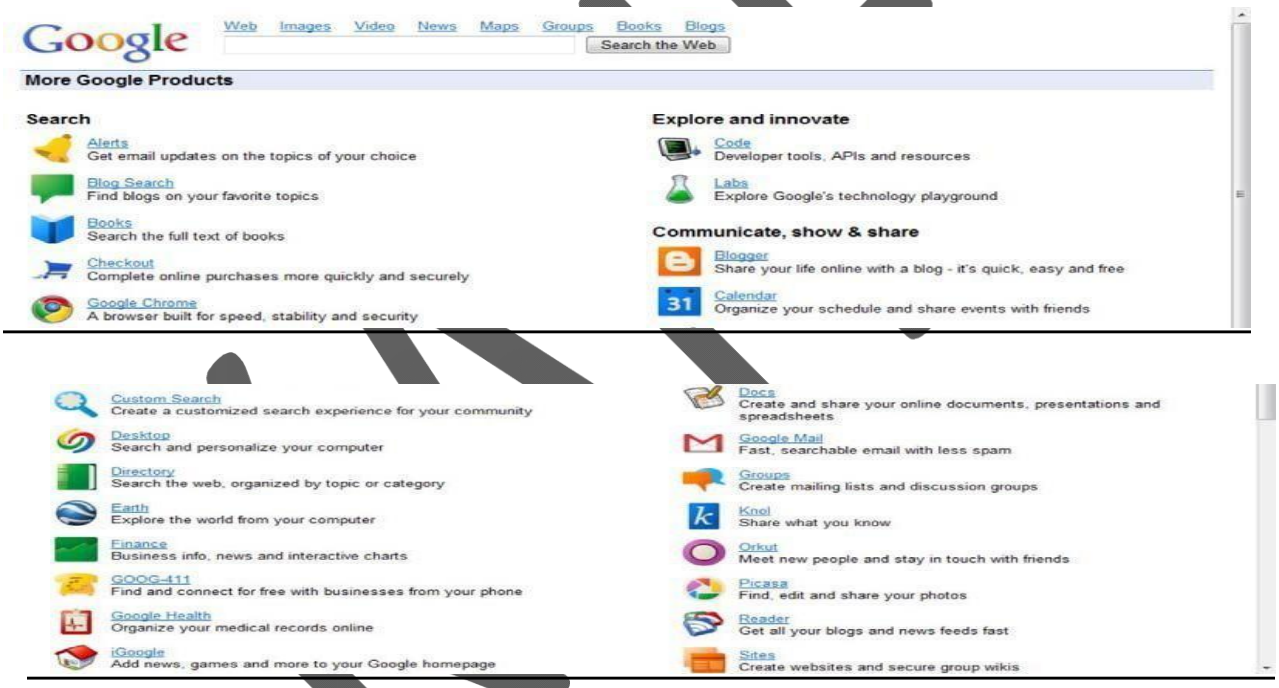

Surveying the Google Application Portfolio

It is fair to say that nearly all the products in Google's application and service portfolio are cloud computing services in that they all rely on systems staged worldwide on Google's one million plus servers in nearly 30 datacenters. Roughly 17 of the 48 services listed leverage Google's search engine in some specific way. Some of these search-related sites search through selected content such as Books, Images, Scholar, Trends, and more. Other sites such

as Blog Search, Finance, News, and some others take the search results and format them into an Aggregation page. The below figure shows one of these aggregation pages: Google Finance.

Google's Finance page at http://www.google.com/finance/ is an example of an aggregation page provided by results from Google's search engine.

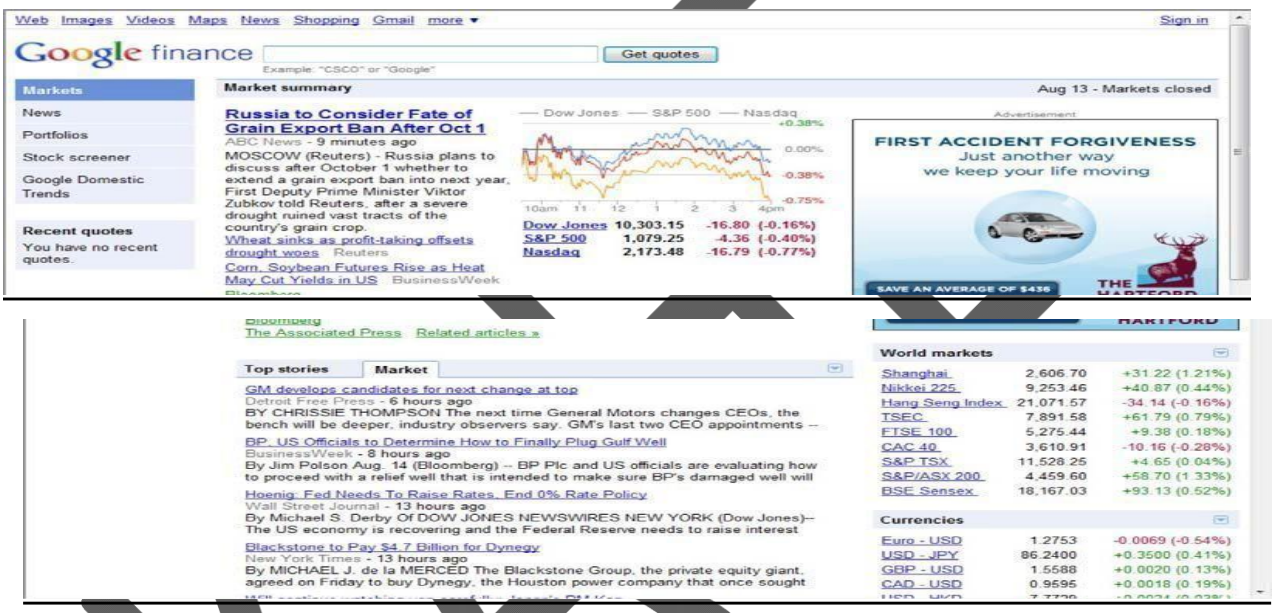

Enterprise offerings

As Google has built out its portfolio, it has released special versions of its products for the enterprise. The following are among Google's products aimed at the enterprise market:

• **Google Commerce Search** (http://www.google.com/commercesearch/): This is a search service for online retailers that markets their products in their site searches with a number of navigation, filtering, promotion, and analytical functions.

• **Google Site Search** (http://www.google.com/sitesearch/): Google sells its search engine customized for enterprises under the Google Site Search service banner. The user enters a search string in the site's search, and Google returns the results from that site.

• **Google Search Appliance** (http://www.google.com/enterprise/gsa): This server can be deployed within an organization to speed up both local (Intranet) and Internet searching. The three versions of the Google Search Appliance can store an index of up to 300,000 (GB-1001), 10 million (GB-5005), or 30 million (GB-8008) documents. Beyond indexing, these appliances have document management features, perform custom searches, cache content, and give local support to Google Analytics and Google Sitemaps.

• **Google Mini** (http://www.google.com/enterprise/mini/): The Mini is the smaller version of the GSA that stores 300,000 indexed documents.

Google Apps for Business is the commercial versions of the company's productivity suites.

Many of Google's productivity applications are quite capable, but none is a state-of-the-art client you might expect to find in a locally installed office suite. When compared one-on-one to Microsoft Office applications, Google's online offerings give users the essential features for a fraction of the Microsoft Office price.

AdWords

AdWords (http://www.google.com/AdWords) is a targeted ad service based on matching

advertisers and their keywords to users and their search profiles. This service transformed Google from a competent search engine into an industry giant and is responsible for the majority of Google's revenue stream. AdWords' two largest competitors are Microsoft adcenter (http://adcenter microsoft.com/) and Yahoo! Search Marketing (http://searchmarketing.yahoo.com/).

Ads are displayed as text, banners, or media and can be tailored based on geographical location, frequency, IP addresses, and other factors. AdWords ads can appear not only on Google.com, but on AOL search, Ask.com, and Netscape, along with other partners. Other partners belonging to the Google Display Network can also display AdSense ads. In all these cases, the AdWords system determines which ads to match to the user searches.

Using OpenNebula

OpenNebula is an open and flexible tool that fits into existing data center's environments to build any type of cloud deployment. OpenNebula can be primarily used as a virtualization tool to manage your virtual infrastructure, which is usually referred to as private cloud. OpenNebula supports a hybrid cloud to combine local infrastructure with public cloud-based infrastructure, enabling highly scalable hosting environments. OpenNebula also supports public clouds by providing cloud's interfaces to expose its functionality for virtual machine, storage, and network management. OpenNebula is one of the technologies being enhanced in the Reservoir Project , European research initiatives in virtualized infrastructures, and cloud computing.

OpenNebula architecture is shown in Figure which illustrates the existence of public and private clouds and also the resources being managed by its virtual manager.

OpenNebula is an open-source alternative to these commercial tools for the dynamic management of VMs on distributed resources. This tool is supporting several research lines in advance reservation of capacity, probabilistic admission control, placement

optimization, resource models for the efficient management of groups of virtual machines, elasticity support, and so on. These research lines address the requirements from both types of clouds namely, private and public.

OpenNebula and Haizea. Haizea is an open-source virtual machine-based lease management architecture developed by Sotomayor et al. ; it can be used as a scheduling backend for OpenNebula. Haizea uses leases as a fundamental resource provisioning abstraction and implements those leases as virtual machines, taking into account the overhead of using virtual machines when scheduling leases.

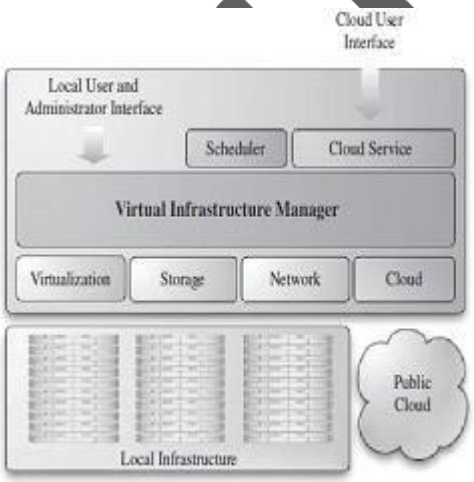

Haizea also provides advanced functionality such as:

- Advance reservation of capacity.
- Best-effort scheduling with backfilling.
- Resource preemption (using VM suspend/resume/migrate).
- Policy engine, allowing developers to write pluggable scheduling policies in Python.

Aneka

Manjrasoft Aneka is a .NET-based platform and framework designed for building and deploying distributed applications on clouds. It provides a set of APIs for transparently exploiting distributed resources and expressing the business logic of applications by using the preferred programming abstractions. Aneka is also a market-oriented cloud

platform since it allows users to build and schedule applications, provision resources, and monitor results using pricing, accounting, and QoS/SLA services in private and/or public cloud environments.

It allows end users to build an enterprise/private cloud setup by exploiting the power of computing resources in the enterprise data centers, public clouds such as Amazon EC2 , and hybrid clouds by combining enterprise private clouds managed by Aneka with resources from Amazon EC2 or other enterprise clouds built and managed using technologies such as XenServer.

Aneka also provides support for deploying and managing clouds. By using its Management Studio and a set ofWeb interfaces, it is possible to set up either public or private clouds, monitor their status, update their configuration, and perform the basic management operations.

Aneka Architecture.

Aneka platform, consists of a collection of physical and virtualized resources connected through a network. Each of these resources hosts an instance of the Aneka container representing the runtime environment where the distributed applications are executed. The container provides the basic management features of the single node and leverages all the other operations on the services that it is hosting. The services are broken up into fabric, foundation, and execution services. Fabric services directly interact with the node through the platform abstraction layer (PAL) and perform hardware profiling and dynamic resource provisioning. Foundation services identify the core system of the Aneka middleware, providing a set of basic features to enable Aneka containers to perform specialized and specific sets of tasks. Execution services directly deal with the scheduling and execution of applications in the cloud.

POSSIBLE QUESTIONS

6 MARKS

1. Discuss about Google Web services with real-time environment.
2. Describe the role of Open Nebula in the cloud environment.
3. Discuss the features of Microsoft Azure in the cloud computing.
4. Elaborate the role of Amazon EC2 in the cloud.

KARPAGAM ACADEMY OF HIGHER EDUCTATION DEPARTMENT OF
III B.Sc CS
CLOUD COMPUTING[17CSU601A] - UNIT III

S.No	**Questions**	**opt1**
1	_____ is a cloud operating system built on top of Microsoft datacenters infrastructure	Microsoft Windows
2	_____ are the core components of Microsoft Windows Azure,	Storage services
3	Expand BLOBs	Binary Large Objects
4	_____ is a comprehensive middleware for developing, deploying, and managing applica tions on the cloud	Microsoft Windows
5	is one of the most important and heavily trafficked Web sites in the world	Yahoo.com
6	_____ is a Platform as a Service (PaaS) cloud-based Web hosting service on Google's infrastructure.	SQL Azure
7	is a service that allows developers to quickly access data persisted on Windows Azure storage.	Web Server Gateway
8	_____ supports databases with a maximum size of 1 GB or 5 GB.	Web Edition
9	_____supports databases with a maximum size from 10 GB to 50 GB	Web Edition
10	In Which Year did Amazon.com made its Web service platform available to	2006
11	_____ is the world's largest online retailer with net sales in $24.51 billion, according to their 2009 annual report.	Google.com
12	_____ is the hourly rate with no long-term commitment.	On-Demand Instance
13	There are currently _____different EC2 service zones or regions	Three
14	_____is the prototypical cloud computing Services Company.	Google
15	_____supports some of the largest Web sites and services in the world	Yahoo

16	SEO stands for_____	Search Engine Optimization
17	Online content that isn't indexed by search engines belongs to what has come to be called the	Crawl Web
18	World's number two Web site, is called_____	Facebook
19	_____ is a prominent example of a site that isn't indexed in search engines	Facebook

20	The success of the ad is measured by what is called the_____	Money through Click
21	_____ supports Dynamic Web services based on common standards.	SQL Azure
22	_____ can be used to run and scale PHP Web applications on Azure.	Worker Roles
23	_____ can be used to host Tomcat and serve JSP-based applications	Web Roles
24	A single block blob can reach ___ in dimension	100 GB
25	_____ type of blob is optimized for random access and can be used to host data different from streaming	Page blobs
26	Access to SQL Azure is based on the _____ protocol	User Datagram
27	_____ is foundation layer + set of developer services	Azure Platform
28	_____provides access to e-mail and display names within your app	Amazon web services
29	_____ eliminates the need for an application to develop its own authentication system	Amazon web services
30	_____ has a distributed datastore system that supports queries and transactions.	Google App Engine
31	The Datastore in Google App Engine is _____	Non-relational
32	The _____ uses an optimistic concurrency control and maintains strong consistency.	Entity Group
33	_____ manage entities as a single group, and entity groups are stored together in the system so operations can be performed faster.	Queries
34	Applications can use the_____ to determine whether a user belongs to a specific group	Admin API
35	Applications running in GAE are isolated from the underlying operating system, which Google describes as running in a	Sandbox
36	The pricing scheme of Google in the Outgoing bandwidth measured in GB is _____ per GB.	$0.20
37	The pricing scheme of Google in the Incoming bandwidth measured in GB is _____ per GB	$0.20
38	The pricing scheme of Google in the Stored data measured in GB per month is _____ per GB/month	$0.20
39	The pricing scheme of Google in the CPU time measured in CPU hours is _____ per hour.	$0.10

40	WSGI stands for_____	Window Server Gateway Interface
41	_____ supports the feature Task queues and task scheduling	Google App Engine
42	_____ can appear not only on Google.com, but on AOL search, Ask.com, and Netscape, along with other partners	AdSense ads
43	The ____system determines which ads to match to the user searches	AdWords
44	Advertisers bid on _____that are used to match a user to their product or service	Exact words
45	_____can be deployed within an organization to speed up both local (Intranet) and Internet searching	Google Site Search
46	_____is the smaller version of the Google Search Appliance .	Google Mini
47	_____ includes Pages without links.	Crawl Web
48	_____ includes Private or limited access Web pages and sites	Crawl Web
49	_____ includes Information contained in sources available through executable code such as JavaScript	Deep Web
50	_____ can be useful in allowing content that isn't browsable to be crawled	Deep Web
51	Web crawlers are also called as _____	Spiders or Robots
52	Content on pages is scanned up to a certain number of words and placed into an _____	Header
53	_____is an example of an aggregation page provided by results from Google's search engine	Google's app page
54	_____which lowers the amount of development effort required to get an application up and running.	Azure Cloud Framework
55	Google's cloud computing services falls under_____umbrellas	2
56	The bulk of Google's income comes from the sales of_____	Social Networking
57	Google is always tweaking the algorithm to prevent _____strategies from gaming the system	Search Engine Results Page
58	_____ represent the units of deployment of Web applications within the Azure infrastructure	Web Roles

59	_____service is optimal to store large text or binary files.	Blocks
60	A connection is the Service Bus element that is priced by Azure on a _____basis	Yearly

COMPUTER SCIENCE

opt2	opt3	opt4	Answer
Oracle Azure	VB Azure	Java Azure	Microsoft Windows Azure
Compute services	Product services	Infrastructure services	Compute services
Blocking Large	Between Large	Blocking Last Objects	Binary Large Objects
AppFabric	SQL Azure	Oracle Azure	AppFabric
Google.com	Amazon.com	MSN.com	Amazon.com
Web Server Gateway	Amazon Web	Google App Engine	Google App Engine
Azure Cache	Amazon Web	SQL Azure	Azure Cache
Business Edition	Standard Edition	Special Edition	Web Edition
Standard Edition	Business Edition	Special Edition	Business Edition
2005	2007	2008	2006
Amazon.com	Force.com	Rackspace.com	Amazon.com
Reserved Instances	Spot Instance	Timing Instance	On-Demand Instance
Four	Five	Six	Four
Yahoo	Amazon	IBM	Google
Amazon	Google	IBM	Google

Small Engine Optimization	Secret Engine	Sophisticated Engine	Search Engine Optimization
Deep Web	Dark Web	Open Web	Deep Web
Google	Yahoo	MSN	Facebook
Google	Yahoo	MSN	Facebook

Click-through rate	Minimum Click	Maximum through Click	Click-through rate
Azure Cloud	Google App Engine	Amazon EC2	Google App Engine
Web Roles	Virtual machine	Storage Roles	Web Roles
Virtual machine	Worker roles	Storage Roles	Worker roles
200 GB	150 GB	160 GB	200 GB
Block blobs	Storage blobs	Cluster blobs	Page blobs
Transmission Control	Tabular Data Stream	Stream Oriented	Tabular Data Stream
Google Platform	Amazon Platform	Azure Infrastructure	Azure Platform
App Engine	SQL Azure	Azure Cloud	App Engine
SQL Azure	App Engine	Azure Cloud	App Engine
Azure Cloud	SQL Azure	Amazon EC2	Google App Engine
Relational	Schema based	Standard	Non-relational
User API	Datastore	Attribute Group	Datastore
Datastore	Entity Group	Transactions	Transactions
User API	Third party API	API	User API
Web pages	Well known ports	Protocols	Sandbox
$0.22	$0.10	$0.12	$0.12
$0.22	$0.10	$0.12	$0.10
$0.10	$0.12	$0.15	$0.15
$0.22	$0.20	$0.12	$0.10

Web Server Gateway Interface	Web Server Gateway Information	Window Server Gateway Information	Web Server Gateway Interface
Azure Cloud	SQL Azure	Amazon EC2	Google App Engine
AdDeep web ads	AdWords ads	AdServer ads	AdWords ads
AdSense	AdDeep web	AdServer	AdWords
Keywords	Retrieved words	Keys	Keywords
Google Commerce Search	Google Search Appliance	Google Mini	Google Search Appliance
Google Commerce	Google Site	Google App	Google Mini
Deep Web	Dark Web	Open Web	Deep Web
Dark Web	Deep Web	Open Web	Deep Web
Crawl Web	Dark Web	Open Web	Deep Web
Crawl Web	Dark Web	Sitemaps	Sitemaps
Dark Robots	Dark Spiders	Shadowers	Spiders or Robots
Index	Table	Database	Index
Google's Store page	Google's Finance	Google's Box page	Google's Finance page
Google App Engine Framework	SQL Azure Framework	Amazon EC2 Framework	Google App Engine Framework
3	4	5	2
Applications	Target advertising	Blogger	Target advertising
Search Engine Data	Search Engine Algorithm	Search Engine Optimization	Search Engine Optimization
Virtual machine	Worker roles	Storage Roles	Web Roles

Blobs	Storage	Cluster	Blobs
Weekly	Pay-as-you-go	Monthly	Pay-as-you-go

KARPAGAM ACADEMY	_____ is a .NET-based platform and framework designed for building and deploying distributed applications on clouds	Haizea
62	____clouds are easier with Azure	Private
63	_____ is a system for implementing on-premise private and hybrid clouds.	Eucalyptus

Manjrasoft Aneka	AdWords	google Analytics	Manjrasoft Aneka
Public	Community	Hybrid	Hybrid
Amazon	Microsoft Azure	Google App Engine	Eucalyptus

Cloud Information Security Objectives

- Developing secure software is based on applying the secure software design principles that form the fundamental basis for software assurance.
- Software assurance has been given many definitions, and it is important to understand the concept. The Software Security Assurance Report 2 defines software assurance as "the basis for gaining justifiable confidence that software will consistently exhibit all properties required to ensure that the software, in operation, will continue to operate dependably despite the presence of sponsored (intentional) faults.
- The Data and Analysis Center for Software (DACS) 4 requires that software must exhibit the following three properties to be considered secure:

  - Dependability - Software that executes predictably and operates correctly under a variety of conditions, including when under attack or running on a malicious host.
  - Trustworthiness - Software that contains a minimum number of vulnerabilities or no vulnerabilities or weaknesses that could sabotage the software's dependability. It must also be resistant to malicious logic.
  - Survivability (Resilience)

    Software that is resistant to or tolerant of attacks and has the ability to recover as quickly as possible with as little harm as possible Seven complementary principles that support information assurance are confidentiality, integrity, availability, authentication, authorization, auditing, and accountability. These concepts are summarized in the following sections.

Confidentiality, Integrity, and Availability

Confidentiality, integrity, and availability are sometimes known as the CIA triad of information system security, and are important pillars of cloud software assurance.

Confidentiality

Confidentiality refers to the prevention of intentional or unintentional unauthorized disclosure of information.

Integrity

The concept of cloud information integrity requires that the following three principles are met:

- Modifications are not made to data by unauthorized personnel or processes.
- Unauthorized modifications are not made to data by authorized personnel or processes.
- The data is internally and externally consistent - in other words, the internal information is consistent both among all sub-entities and with the real-world, external situation.

Availability

Availability ensures the reliable and timely access to cloud data or cloud computing resources by the appropriate personnel. Availability guarantees that the systems are functioning properly when needed. In addition, this concept guarantees that the security services of the cloud system are in working order.

Cloud Security Services

Authentication

Authentication is the testing or reconciliation of evidence of a user's identity. It establishes the user's identity and ensures that users are who they claim to be. For example, a user presents an identity (user ID) to a computer login screen and then has to provide a password. The computer system authenticates the user by verifying that the password corresponds to the individual presenting the ID.

Authorization

Authorization refers to rights and privileges granted to an individual or process that enable access to computer resources and information assets. Once a user's identity and authentication are established, authorization levels determine the extent of system rights a

user can hold.

Auditing

To maintain operational assurance, organizations use two basic methods: system audits and monitoring. These methods can be employed by the cloud customer, the cloud provider, or both, depending on asset architecture and deployment.

Relevant Cloud Security Design Principles

Historically, computer software was not written with security in mind; but because of the increasing frequency and sophistication of malicious attacks against information systems, modern software design methodologies include security as a primary objective. With cloud computing systems seeking to meet multiple objectives, such as cost, performance, reliability, maintainability, and security, trade-offs have to be made.

A completely secure system will exhibit poor performance characteristics or might not function at all. Technically competent hackers can usually find a way to break into a computer system, given enough time and resources.

The goal is to have a system that is secure enough for everyday use while exhibiting reasonable performance and reliability characteristics.

The CIA Triad

The three fundamental tenets of information security confidentiality, integrity, and availability (CIA) define an organization's security posture. All of the information security controls and safeguards, and all of the threats, vulnerabilities, and security processes are subject to the CIA yardstick.

Confidentiality

Confidentiality is the prevention of the intentional or unintentional unauthorized disclosure of contents. Loss of confidentiality can occur in many ways.

For example, loss of confidentiality can occur through the intentional release of private company information or through a misapplication of network rights.

Integrity

Integrity is the guarantee that the message sent is the message received and that the message is not intentionally or unintentionally altered. Loss of integrity can occur through an intentional attack to change information (for example, a website defacement) or, more commonly, unintentionally (data is accidentally altered by an operator). Integrity also contains the concept of non repudiation of a message source, which we will describe later.

Some of the elements used to ensure integrity include the following:

□ Firewall services

□ Communications security management

□ Intrusion detection services

Availability

This concept refers to the elements that create reliability and stability in networks and systems. It ensures that connectivity is accessible when needed, allowing authorized users to access the network or systems.

Also included in that assurance is the guarantee that security services for the security practitioner are usable when they are needed. The concept of availability also tends to include areas in an information system (IS) that are traditionally not thought of as pure security (such as guarantee of service, performance, and up time), yet are obviously affected by breaches such as a denial-of-service (DoS) attack.

Privacy and Compliance Risks

One area that is greatly affected by cloud computing is privacy. It's important to remember that although the control of cloud computing privacy has many threats and vulnerabilities in common with non cloud processes and infrastructure, it also has unique security issues.

For example, a successful identity theft exploit can result in a privacy loss that has a huge impact on an enterprise. The organization can suffer short-term losses due to remediation, investigation, and restitution costs. It can also incur longer term problems for the organization due to loss of credibility, confidence, and negative publicity.

Another mistake organizations often make is in assigning responsibility for privacy controls to the IT dept, rather than a business unit that owns the data. Information systems security frameworks have defined, standardized processes that apply to cloud computing and its potential privacy breaches.

An individual's right to privacy is embodied in the fundamental principles of privacy:

□ Notice — Regarding the collection, use, and disclosure of personally identifiable information (PII)

□ Choice — To opt out or opt in regarding disclosure of PII to third parties

□ Access — By consumers to their PII to permit review and correction of information

□ Security — To protect PII from unauthorized disclosure

□ Enforcement — Of applicable privacy policies and obligations

Privacy Legislation

The following list summarizes some important legislation and recommended guidelines for privacy:

□ The Cable Communications Policy Act provides for discretionary use of PII by cable operators internally but imposes restrictions on disclosures to third parties.

□ The Children's Online Privacy Protection Act (COPPA) is aimed at providing protection to

children under the age of 13.

□ Customer Proprietary Network Information Rules apply to telephone companies and restrict their use of customer information both internally and to third parties.

□ The Financial Services Modernization Act (Gramm-Leach-Bliley) requires financial institutions to provide customers with clear descriptions of the institution's policies and procedures for protecting the PII of customers.

□ The Telephone Consumer Protection Act restricts communications between companies and consumers, such as telemarketing.

□ The 1973 U.S. Code of Fair Information Practices states that:

1. There must not be personal data record-keeping systems whose very existence is secret.

2. There must be a way for a person to find out what information about them is in a record and how it is used.

3. There must be a way for a person to prevent information about them, which was obtained for one purpose, from being used or made available for another purpose without their consent.

4. Any organization creating, maintaining, using, or disseminating records of identifi able personal data must ensure the reliability of the data for their intended use and must take precautions to prevent misuses of that data.

Threats to Infrastructure, Data, and Access Control

To properly understand the threats that cloud computing presents to the computing infrastructure, it's important to understand communications security techniques to prevent, detect, and correct errors so that integrity, availability, and the confidentiality of transactions

over networks may be maintained.

This includes the following:

□ Communications and network security as it relates to voice, data, multimedia, and facsimile transmissions in terms of local area, wide area, and remote access networks

□ Internet/intranet/extranet in terms of firewalls, routers, gateways, and various protocols

Common Threats and Vulnerabilities

A threat is simply any event that, if realized, can cause damage to a system and create a loss of confidentiality, availability, or integrity. Threats can be malicious, such as the intentional modification of sensitive information, or they can be accidental such as an error in a transaction calculation or the accidental deletion of a file.

Vulnerability is a weakness in a system that can be exploited by a threat. Reducing the vulnerable aspects of a system can reduce the risk and impact of threats on the system. For example, a password-generation tool, which helps users choose robust passwords, reduces the chance that users will select poor passwords (the vulnerability) and makes the password more difficult to crack.

Common threats to both cloud and traditional infrastructure include the following:

□ Eavesdropping — Data scavenging, traffic or trend analysis, social engineering, economic or political espionage, sniffing, dumpster diving, keystroke monitoring, and shoulder surfi ng are all types of eavesdropping to gain information or to create a foundation for a later attack. Eavesdropping is a primary cause of the failure of confidentiality.

□ Fraud — Examples of fraud include collusion, falsified transactions, data manipulation, and other altering of data integrity for gain.

□ Theft — Examples of theft include the theft of information or trade secrets for profit or unauthorized disclosure, and physical theft of hardware or software.

□ Sabotage — Sabotage includes denial-of-service (DoS) attacks, production delays, and data integrity sabotage.

□ External attack — Examples of external attacks include malicious cracking, scanning, and probing to gain infrastructure information, demon dialing to locate an unsecured modem line, and the insertion of a malicious code or virus. Many network attacks and abuses share a commonality with traditional infrastructure and cloud infrastructure. Attacks against computers, networks, and cryptographic systems have a variety of motivations. Some attacks are aimed at disrupting service, others focus on illegally acquiring sensitive information, and others attempt to deceive or defraud. In general, such attacks target the CIA components of information security.

Logon Abuse

Logon abuse can refer to legitimate users accessing services of a higher security level that would normally be restricted to them. Unlike network intrusion, this type of abuse focuses primarily on those users who might be legitimate users of a different system or users who have a lower security classification.

Masquerading is the term used when one user pretends to be another user, such as an attacker socially engineering passwords from an Internet Service Provider (ISP).

Inappropriate System Use

This style of network abuse refers to the non business or personal use of a network by otherwise authorized users, such as Internet surfing to inappropriate content sites (travel, pornography, sports, and so forth). As per the International Information Systems Security Certification Consortium (ISC) Code of Ethics and the Internet Advisory Board (IAB) recommendations, the use of networked services for other than business purposes can be considered abuse of the system.

While most employers do not enforce extremely strict Web surfing rules, occasional harassment litigation may result from employees accessing pornography sites and employees operating private Web businesses using the company's infrastructure.

Eavesdropping

This type of network attack consists of the unauthorized interception of network traffic. Certain network transmission methods, such as satellite, wireless, mobile, PDA, and so on, are vulnerable to eavesdropping attacks. Tapping refers to the physical interception of a transmission medium (like the splicing of a cable or the creation of an induction loop to pick up electromagnetic emanations from copper).

Accountability

Accountability is another facet of access control. Individuals on a system are responsible for their actions. This accountability property enables system activities to be traced to the proper individuals. Accountability is supported by audit trails that record events on both the system and the network. Audit trails can be used for intrusion detection and for the reconstruction of past events.

The following measures compensate for both internal and external access violations:

□ Backups

□ RAID (Redundant Array of Independent Disks)technology

□ Fault tolerance

□ Business continuity planning

□ Insurance

Architectural Considerations

A variety of factors affect the implementation and performance of cloud security architecture. There are general issues involving regulatory requirements, adherence to standards, security management, information classification, and security awareness. Then there are more specific architecturally related areas, including trusted hardware and software, providing for a secure execution environment, establishing secure communications, and hardware augmentation through micro architectures. These important concepts are addressed in this section.

General Issues

A variety of topics influence and directly affect the cloud security architecture. They include such factors as compliance, security management, administrative issues, controls, and security awareness.

Compliance

In a public cloud environment, the provider does not normally inform the clients of the storage location of their data. In fact, the distribution of processing and data storage is one of the cloud's fundamental characteristics. However, the cloud provider should cooperate to consider the client's data location requirements.

Security Management

Security architecture involves effective security management to realize the benefits of cloud computation. Proper cloud security management and administration should identify management issues in critical areas such as access control, vulnerability analysis, change control, incident response, fault tolerance, and disaster recovery and business continuity planning.

Controls

The objective of cloud security controls is to reduce vulnerabilities to a tolerable level and minimize the effects of an attack. To achieve this, an organization must determine what impact an attack might have, and the likelihood of loss.

Complementary Actions

Additional activities involved in cloud security management include the following:

□ Management and monitoring of service levels and service-level agreements

□ Acquisition of adequate data to identify and analyze problem situations through instrumentation and dashboards.

□ Reduction of the loss of critical information caused by lack of controls.

□ Proper management of data on an organization's distributed computing resources. Data centralized on the cloud reduces the potential for data loss in organizations with large numbers of laptop computers and other personal computing devices.

□ Monitoring of centrally stored cloud information, as opposed to having to examine data distributed throughout an organization on a variety of computing and storage devices.

□ Provisioning for rapid recovery from problem situations.

The purpose of computer security awareness, training, and education is to enhance security by doing the following:

□ Improving awareness of the need to protect system resources

□ Developing skills and knowledge so computer users can perform their jobs more securely

□ Building in-depth knowledge, as needed, to design, implement, or operate security programs for organizations and systems

An effective computer security awareness and training program requires proper planning, implementation, maintenance, and periodic evaluation. In general, a computer security awareness and training program should encompass the following seven steps:

1. Identify program scope, goals, and objectives.

2. Identify training staff.

3. Identify target audiences.

4. Motivate management and employees.

5. Administer the program.

6. Maintain the program.

7. Evaluate the program.

Making cloud system users and providers aware of their security responsibilities and teaching them correct practices helps change their behavior. It also supports individual accountability because without knowledge of the necessary security measures and to how to use them, personnel cannot be truly accountable for their actions.

Security Awareness

As opposed to training, the security awareness of an organization refers to the degree to which its personnel are collectively aware of the importance of security and security controls. In addition to the benefits and objectives previously mentioned, security awareness programs also have the following benefits:

□ They can reduce the unauthorized actions attempted by personnel.

□ They can significantly increase the effectiveness of the protection controls.

□ They help to prevent the fraud, waste, and abuse of computing resources.

Personnel are considered "security aware" when they clearly understand the need for security, how security affects viability and the bottom line, and the daily risks to cloud computing resources.

It is important to have periodic awareness sessions to orient new employees and refresh senior employees. The material should always be direct, simple, and clear. It should be fairly motivational and should not contain a lot of techno-jargon, and you should convey it in a style that the audience easily understands. These sessions are most effective when they demonstrate how the security interests of the organization parallel the interests of the audience.

Training and Education

Training is different from awareness in that it provides security information in a more formalized manner, such as classes, workshops, or individualized instruction. The following types of training are related to cloud security:

□ Security-related job training for operators and specific users
□ Awareness training for specific departments or personnel groups with security-sensitive positions
□ Technical security training for IT support personnel and system administrators
□ Advanced training for security practitioners and information systems auditors
□ Security training for senior managers, functional managers, and business unit managers

Trusted Cloud Computing

Trusted cloud computing can be viewed as a computer security architecture that is designed to protect cloud systems from malicious intrusions and attacks, and ensure that computing resources will act in a specific, predictable manner as intended. A trusted cloud computing system will protect data in use by hypervisors and applications, protect against unauthorized access to information, provide for strong authentication, apply encryption to protect sensitive data that resides on stolen or lost devices, and support compliance through hardware and software mechanisms.

Trusted Computing Characteristics

In a cloud computational system, multiple processes might be running concurrently. Each process has the capability to access certain memory locations and to execute a subset of the computer's instruction set. The execution and memory space assigned to each process is called a protection domain. This domain can be extended to virtual memory, which increases the apparent size of real memory by using disk storage. The purpose of establishing a protection domain is to protect programs from all unauthorized modification or executional interference.

A trusted computing base (TCB) is the total combination of protection mechanisms within a computer system, which includes the hardware, software, and firmware, that are trusted to enforce a security policy.

Another element associated with trusted computing is the trusted platform module (TPM). The TPM stores cryptographic keys that can be used to attest to the operating state of a computing platform and to verify that the hardware and software configuration has not been modified

Virtual Private Networks

Another important method to secure cloud communications is through a virtual private network (VPN). A VPN is created by building a secure communications link between two nodes by emulating the properties of a point-to-point private link. A VPN can be used to facilitate secure remote access into the cloud, securely connect two networks together, or create a secure data tunnel within a network.

The portion of the link in which the private data is encapsulated is known as the tunnel. It may be referred to as a secure, encrypted tunnel, although it's more accurately defined as an encapsulated tunnel, as encryption may or may not be used. To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information. Most often the data is encrypted for confidentiality. This encrypted part of the link is considered the actual virtual private network connection. Figure shows a common VPN configuration with example IP addresses for remote access into an organization's intranet through the Internet. Address 192.168.123.2 designates the organization's router.

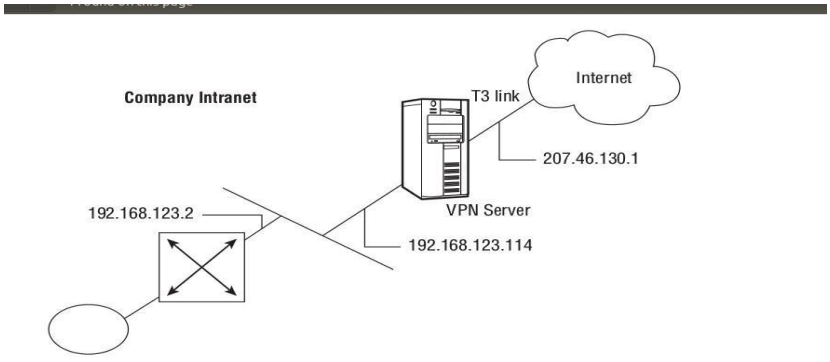

Figure 6-1: VPN configuration

The two general types of VPNs relevant to cloud computing are remote access and network-to-network. These VPN types are described in the following sections.

Remote Access VPNs

A VPN can be configured to provide remote access to corporate resources over the public Internet to maintain confidentiality and integrity. This configuration enables the

remote user to utilize whatever local ISP is available to access the Internet without forcing the user to make a long-distance or 800 call to a third-party access provider. Using the connection to the local ISP, the VPN software creates a virtual private network between the dial-up user and the corporate VPN server across the Internet.

Network-to-Network VPNs

A VPN is commonly used to connect two networks, perhaps the main corporate LAN and a remote branch office LAN, through the Internet. This connection can use either dedicated lines to the Internet or dial-up connections to the Internet. However, the corporate hub router that acts as a VPN server must be connected to a local ISP with a dedicated line if the VPN server needs to be available 24/7. The VPN software uses the connection to the local ISP to create a VPN tunnel between the branch office router and the corporate hub router across the Internet.

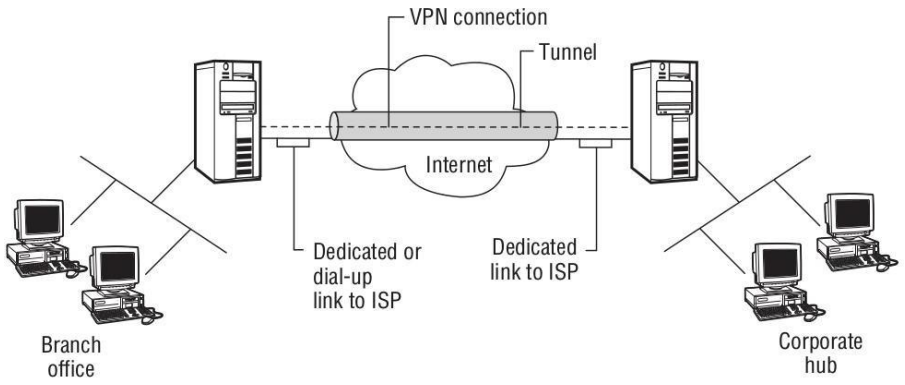

Figure 6-3: A network-to-network VPN configuration

Identity Management and Access Control

Identity management and access control are fundamental functions required for secure

cloud computing. The simplest form of identity management is logging on to a computer system with a user ID and password. However, true identity management, such as is required for cloud computing, requires more robust authentication, authorization, and access control. It should determine what resources are authorized to be accessed by a user or process by using technology such as biometrics or smart cards, and determine when a resource has been accessed by unauthorized entities.

Identity Management

Identification and authentication are the keystones of most access control systems. Identification is the act of a user professing an identity to a system, usually in the form of a username or user logon ID to the system. Identification establishes user accountability for the actions on the system. User IDs should be unique and not shared among different individuals. In many large organizations, user IDs follow set standards, such as first initial followed by last name, and so on. In order to enhance security and reduce the amount of information available to an attacker, an ID should not reflect the user's job title or function. Authentication is verification that the user's claimed identity is valid, and it is usually implemented through a user password at logon. Authentication is based on the following three factor types:

□ Type 1 — Something you know, such as a personal identification number (PIN) or password

□ Type 2 — Something you have, such as an ATM card or smart card

□ Type 3 — Something you are (physically), such as a fingerprint or retina scan

Passwords

Because passwords can be compromised, they must be protected. In the ideal case, a password should be used only once. This "one-time password," or OTP, provides maximum

security because a new password is required for each new logon. A password that is the same for each logon is called a static password.

A password that changes with each logon is termed a dynamic password. The changing of passwords can also fall between these two extremes. Passwords can be required to change monthly, quarterly, or at other intervals, depending on the criticality of the information needing protection and the password's frequency ofuse.

Memory Cards

Memory cards provide nonvolatile storage of information, but they do not have any processing capability. A memory card stores encrypted passwords and other related identifying information. A telephone calling card and an ATM card are examples of memory cards.

Smart Cards

Smart cards provide even more capability than memory cards by incorporating additional processing power on the cards. These credit-card-size devices comprise microprocessor and memory and are used to store digital signatures, private keys, passwords, and other personal information.

Biometrics

An alternative to using passwords for authentication in logical or technical access control is biometrics. Biometrics is based on the Type 3 authentication mechanism something you are. Biometrics is defined as an automated means of identifying or authenticating the identity of a living person based on physiological or behavioral characteristics.

In biometrics, identification is a one-to-many search of an individual's characteristics from a database of stored images. Authentication is a one-to-one search to verify a claim to an identity made by a person. Biometrics is used for identification in physical controls and for authentication in logical controls.

There are three main performance measures in biometrics:

□ False rejection rate (FRR) or Type I Error — The percentage of valid subjects that are falsely rejected.

□ False acceptance rate (FAR) or Type II Error — The percentage of invalid subjects that are falsely accepted.

□ Crossover error rate (CER) — The percentage at which the FRR equals the FAR. The smaller the CER, the better the device is performing.

The following are typical biometric characteristics that are used to uniquely authenticate an individual's identity:

Fingerprints — Fingerprint characteristics are captured and stored. Typical CERs are 4–5%.

□ Retina scans — The eye is placed approximately two inches from a camera and an invisible light source scans the retina for blood vessel patterns. CERs are approximately 1.4%.

□ Iris scans — A video camera remotely captures iris patterns and characteristics. CER values are around 0.5%.

□ Hand geometry — Cameras capture three-dimensional hand characteristics. CERs are approximately 2%.

□ Voice — Sensors capture voice characteristics, including throat vibrations and air pressure, when the subject speaks a phrase. CERs are in the range of 10%.

□ Handwritten signature dynamics — The signing characteristics of an individual making a signature are captured and recorded. Typical characteristics including writing pressure and pen direction. CERs are not published at this time.

Other types of biometric characteristics include facial and palm scans.

Implementing Identity Management

Realizing effective identity management requires a high-level corporate commitment and dedication of suffi cient resources to accomplish the task. Typical undertakings in putting identity management in place include the following:

□ Establishing a database of identities and credentials

□ Managing users' access rights

□ Enforcing security policy

□ Developing the capability to create and modify accounts

□ Setting up monitoring of resource accesses

□ Installing a procedure for removing access rights

□ Providing training in proper procedures

An identity management effort can be supported by software that automates many of the required tasks.

The Open Group and the World Wide Web Consortium (W3C) are working toward a standard for a global identity management system that would be interoperable, provide for privacy, implement accountability, and be portable.

Identity management is also addressed by the XML-based eXtensible Name Service (XNS) open protocol for universal addressing. XNS provides the following capabilities:

□ A permanent identification address for a container of an individual's personal data and contact information

□ Means to verify whether an individual's contact information is valid

□ A platform for negotiating the exchange of information among different entities

Access Control

Access control is intrinsically tied to identity management and is necessary to preserve the confidentiality, integrity, and availability of cloud data.

These and other related objectives flow from the organizational security policy. This policy is a high-level statement of management intent regarding the control of access to information and the personnel who are authorized to receive that information.

Three things that must be considered for the planning and implementation of access control mechanisms are threats to the system, the system's vulnerability to these threats, and the risk that the threats might materialize. These concepts are defined as follows:

□ Threat — An event or activity that has the potential to cause harm to the information systems or networks

□ Vulnerability — A weakness or lack of a safeguard that can be exploited by a threat, causing harm to the information systems or networks

□ Risk — The potential for harm or loss to an information system or network; the probability that a threat will materialize

Controls

Controls are implemented to mitigate risk and reduce the potential for loss. Two important control concepts are separation of duties and the principle of least privilege.

Separation of duties requires an activity or process to be performed by two or more entities for successful completion.

Thus, the only way that a security policy can be violated is if there is collusion among the entities. For example, in a financial environment, the person requesting that a check be issued for payment should not also be the person who has authority to sign the check.

Least privilege means that the entity that has a task to perform should be provided with the minimum resources and privileges required to complete the task for the minimum necessary period of time.

Control measures can be administrative, logical (also called technical), and physical in their implementation.

□ Administrative controls include policies and procedures, security awareness training, background checks, work habit checks, a review of vacation history, and increased supervision.

□ Logical or technical controls involve the restriction of access to systems and the protection of information. Examples of these types of controls are encryption, smart cards, access control lists, and transmission protocols.

□ Physical controls incorporate guards and building security in general, such as the locking of doors, the securing of server rooms or laptops, the protection of cables, the separation of duties, and the backing up of files.

Controls provide accountability for individuals who are accessing sensitive information in a cloud environment. This accountability is accomplished through access control mechanisms that require identification and authentication, and through the audit function.

These controls must be in accordance with and accurately represent the organization's security policy. Assurance procedures ensure that the control mechanisms correctly implement the security policy for the entire life cycle of a cloud information system.

In general, a group of processes that share access to the same resources is called a protection domain, and the memory space of these processes is isolated from other running processes.

Models for Controlling Access

Controlling access by a subject (an active entity such as an individual or process) to an object (a passive entity such as a file) involves setting up access rules. These rules can be classified into three categories or models.

Mandatory Access Control

The authorization of a subject's access to an object depends upon labels, which indicate the subject's clearance, and the classification or sensitivity of the object.

For example, the military classifies documents as unclassified, confidential, secret, and top secret. Similarly, an individual can receive a clearance of confidential, secret, or top secret and can have access to documents classified at or below his or her specified clearance level. Thus, an individual with a clearance of "secret" can have access to secret and confidential documents with a restriction.

This restriction is that the individual must have a need to know relative to the classified documents involved. Therefore, the documents must be necessary for that individual to complete an assigned task.

Even if the individual is cleared for a classification level of information, the individual should not access the information unless there is a need to know. Rule-based access control is

a type of mandatory access control because rules determine this access (such as the correspondence of clearance labels to classification labels), rather than the identity of the subjects and objects alone.

Discretionary Access Control

With discretionary access control, the subject has authority, within certain limitations, to specify what objects are accessible.

For example, access control lists (ACLs) can be used. An access control list is a list denoting which users have what privileges to a particular resource. For example, a tabular listing would show the subjects or users who have access to the object, e.g., file X, and what privileges they have with respect to that file.

An access control triple consists of the user, program, and file, with the corresponding access privileges noted for each user.

This type of access control is used in local, dynamic situations in which the subjects must have the discretion to specify what resources certain users are permitted to access. When a user within certain limitations has the right to alter the access control to certain objects, this is termed a user-directed discretionary access control.

An identity-based access control is a type of discretionary access control based on an individual's identity. In some instances, a hybrid approach is used, which combines the features of user-based and identity-based discretionary access control.

Nondiscretionary Access Control

A central authority determines which subjects can have access to certain objects based on the organizational security policy.

The access controls might be based on the individual's role in the organization (role-based) or the subject's responsibilities and duties (task-based). In an organization with frequent personnel changes, nondiscretionary access control is useful because the access controls are based on the individual's role or title within the organization.

Therefore, these access controls don't need to be changed whenever a new person assumes that role. Access control can also be characterized as context-dependent or content-dependent.

Context-dependent access control is a function of factors such as location, time of day, and previous access history. It is concerned with the environment or context of the data. In content-dependent access control, access is determined by the information contained in the item being accessed.

Single Sign-On (SSO)

Single sign-on (SSO) addresses the cumbersome situation of logging on multiple times to access different resources. When users must remember numerous passwords and IDs, they might take shortcuts in creating them that could leave them open to exploitation.

In SSO, a user provides one ID and password per work session and is automatically logged on to all the required applications. For SSO security, the passwords should not be stored or transmitted in the clear. SSO applications can run either on a user's workstation or on authentication servers.

The advantages of SSO include having the ability to use stronger passwords, easier administration of changing or deleting the passwords, and less time to access resources. The major disadvantage of many SSO implementations is that once users obtain access to the system through the initial logon, they can freely roam the network resources without any restrictions.

Authentication mechanisms include items such as smart cards and magnetic badges. Strict controls must be enforced to prevent a user from changing configurations that another authority sets.

SSO can be implemented by using scripts that replay the users' multiple logins or by using authentication servers to verify a user's identity, and encrypted authentication tickets to permit access to system services.

Enterprise access management (EAM) provides access control management services to Web-based enterprise systems that include SSO. SSO can be provided in a number of ways.

For example, SSO can be implemented on Web applications residing on different servers in the same domain by using non persistent, encrypted cookies on the client interface.

This task is accomplished by providing a cookie to each application that the user wishes to access. Another solution is to build a secure credential for each user on a reverse proxy that is situated in front of the Web server. The credential is then presented each time a user attempts to access protected Web applications.

KARPAGAM ACADEMY OF HIGHER EDUCTATION DEPARTME
III B.Sc CS
CLOUD COMPUTING[17CSU601A] - UN

S.NO	QUESTIONS	OPTION 1	OPTION 2
1	DACS stands for_____	Data and Analysis Center for	Data and Analysis Center for
2	CIA stands for_____	Confidential, Independent and Analysis	Confidentiality, Integrity, and Availability
3	_____ ensures the reliable and timely access to cloud data or cloud computing resources by the appropriate personnel.	Integrity	Availability
4	_____refers to the prevention of intentional or unintentional unauthorized disclosure of	Integrity	Security
5	There are ____ types of cloud security services.	3	2
6	_____ is the testing or reconciliation of evidence of a user's identity.	Authentication	Authorization
7	_____refers to rights and privileges granted to an individual.	Authentication	Auditing

KARPAGAM
ACADEMY OF HIGHER EDUCATION
(Deemed to be University)
(Established Under Section 3 of UGC Act, 1956)

8	DoS stands for_____	Denial-of-System	Denial-of-Security

9	_____can be regarded as the collection, use, and disclosure of personally identifiable information.	Notice	Access
10	The Telephone Consumer Protection Act restricts communications between companies and consumers, such as_____	Telesales	Telegraph
11	A____is simply any event that, if realized, can cause damage to a system and create a loss of confidentiality, availability, or integrity	Threat	Theft
12	____ is a weakness in a system that can be exploited by a threat.	Eavesdropping	Virus
13	_____includes denial-of-service attacks, production delays, and data integrity.	Sabotage	External attack
14	_____ is a primary cause of the failure of confidentiality.	Vulnerability	Threat
15	_____can refer to legitimate users accessing services of a higher security level that would normally be restricted to them.	Sabotage	Eavesdropping
16	_____refers to the physical interception of a transmission medium.	Hub	Cable
17	_____can be used for intrusion detection and for the reconstruction of past events.	Audit trails	Accountability
18	RAID stands for_____	Redundant Array of Important Disks	Redundant Array of Independent Disks
19	The objective of cloud security controls is to reduce _____to a tolerable level and minimize the effects of an attack.	Vulnerabilities	Threat
20	_____can reduce the unauthorized actions attempted by personnel.	Trusted Cloud Computing	Security Awareness
21	A_____ is the total combination of protection mechanisms within a computer system.	Trusted platform module	Trusted computing base

22	The_____stores cryptographic keys that can be used to attest to the operating state of a computing platform.	Trusted computing network	Trusted computing base
23	VPN stands for_____	Virtual Public Network	Virtual Private Network
24	The simplest form of _____ is logging on to a computer system with a user ID and password.	Access Management	Identity Control
25	_____ is verification that the user's claimed identity is valid	Authentication	Authorization
26	_____provides maximum security because a new password is required for each new logon.	OTP	OPT
27	A password that is the same for each logon is called a _____	One Time Password	Static Password
28	A password that changes with each logon is termed a _____	One Time Password	Static Password

:NT OF COMPUTER SCIENCE

NIT IV

OPTION 3	OPTION 4	ANSWER
Data and Analysis Center	Data and Analysis Center	Data and Analysis Center for Software
Cloud Integrity Assurance	Cloud Independent Analysis	Confidentiality, Integrity, and Availability
Confidentiality	Accessibility	Availability
Confidentiality	Privacy	Confidentiality
4	5	3
Auditing	Accessibility	Authentication
Authorization	Accessibility	Authorization

Denial-of- Software	Denial-of- Service	Denial-of-Service

Security	Enforcement	Notice
Telemarketing	Teleconsulting	Telemarketing
Vulnerability	Fraud	Threat
Fraud	Vulnerability	Vulnerability
Eavesdropping	Vulnerability	Sabotage
Sabotage	Eavesdropping	Eavesdropping
Inappropriate System Use	Logon abuse	Logon abuse
Fiber optics	Tapping	Tapping
Fault Tolerance	Backups	Audit trails
Random Access of Independent Disks	Random Access of Important Disks	Redundant Array of Independent Disks
Fraud	Theft	Vulnerabilities
Security Provision	Security Access	Security Awareness
Trusted computing network	Trusted software module	Trusted computing base

Trusted platform module	Trusted software module	Trusted platform module
Virtual Pin Network	Virtual Protected	Virtual Private Network
Access Control	Identity Management	Identity Management
Auditing	Accessibility	Authentication
POT	TOP	OTP
Special Password	Dynamic Password	Static Password
Special Password	Dynamic Password	Dynamic Password

KARPAGAM	_____provide nonvolatile storage of information, but they do not have any processing	Memory cards	Biometrics
30	An alternative to using passwords for authentication in logical or technical access	Biometrics	Passwords
31	_____is a one-to-one search to verify a claim to an identity made by a person.	Accessibility	Authentication
32	There are ____ main performance measures in biometrics.	3	2
33	False rejection rate is also known as_____	Type II Error	Type I Error
34	False acceptance rate is also known as_____	Type II Error	Type I Error
35	The percentage at which the FRR equals the FAR is called _____	Distance error rate	Total error rate
36	Typical Crossover Error Rates of Fingerprints are _____	3-4%	4-5%
37	Typical Crossover Error Rates of Retina scans are _____	1.4%(app.)	1.5%(app.)
38	Typical Crossover Error Rates of Iris scans are _____	around 0.5%.	around 0.4%.
39	Typical Crossover Error Rates of Hand geometry are _____	approximately 2%.	approximately 1%.
40	Typical Crossover Error Rates of Voice are _____	12%	10%
41	_____addresses the cumbersome situation of logging on multiple times to access different	Super sign-on	Situation sign-on
42	_____provides access control management services to Web-based enterprise systems that include SSO.	Framework access management	Environment access management
43	An _____ is a list denoting which users have what privileges to a particular resource.	File control list	Access control list

44	_____ controls involve the restriction of access to systems and the protection of information.	Administrative	Logical
45	_____controls incorporate guards and building security in general.	Risk	Physical
46	_____are implemented to mitigate risk and reduce the potential for loss.	Security	Privacy
47	A _____ can be violated is if there is collusion among the entities.	Security policy	Privacy policy
48	The potential for harm or loss to an information system or network is called_____	Threat	Risk
49	XNS stands for_____	eXtensible Name Service	eXtensible Number Service
50	_____means to verify whether an individual's contact information is valid.	eXtensible Number Service	eXtensible Name Security

Smart cards	Passwords	Memory cards
Memory cards	Biometrics	Biometrics
Authorization	Auditing	Authentication
4	5	3
Type III Error	Type IV Error	Type I Error
Type III Error	Type IV Error	Type II Error
Crossover error rate	Performance error rate	Crossover error rate
5-6%	2-3%	4-5%
1.6%(app.)	1.7%(app.)	1.4%(app.)
around 0.6%.	around 0.7%.	around 0.5%.
approximately 3%.	approximately 4%.	approximately 2%.
11%	15%	10%
Sign single-on	Single sign-on	Single sign-on
Enterprise access management	Enterprise control access management	Enterprise access management
Program control list	Data control list	Access control list

Physical	Risk	Logical
Administrative	Logical	Physical
Controls	Priviledge	Controls
Network policy	Access policy	Security policy
Vulnerability	Fraud	Risk
eXtensible Name Security	eXtensible Number Security	eXtensible Name Service
eXtensible Name Service	eXtensible Number Security	eXtensible Name Service

KARPAGAM ACADEMY OF	_____ for senior managers, functional managers, and business unit managers.	Technical security training	Awareness training
52	_____for IT support personnel and system administrators.	Technical security training	Awareness training
53	_____for specific departments or personnel groups with security-sensitive positions.	Technical security training	Awareness training
54	Confidentiality, integrity, and availability are important pillars of _____	Cloud security services	Cloud Information Security
55	_____concept guarantees that the security services of the cloud system are in working order.	Availability	Confidentiality
56	_____is the guarantee that the message sent is the message received and that the message is not intentionally or unintentionally	Integrity	Security
57	_____ include management and monitoring of service levels and service-level agreements.	Cloud access service	Cloud Information Security
58	In general, a computer security awareness and training program should encompass steps.	6	7
59	There are___ types of training are related to cloud security.	7	4
60	_____ can be viewed as a computer security architecture.	Cloud Information Security	Trusted cloud computing

Advanced training	Security training	Security training
Advanced training	Security training	Technical security training
Advanced training	Security training	Awareness training
Cloud software assurance	Cloud Independent Analysis	Cloud software assurance
Accessibility	Integrity	Availability
Privacy	Confidentiality	Integrity
Cloud security management	Cloud access security	Cloud security management
5	9	7
6	5	5
Cloud security management	Cloud access security	Trusted cloud computing

Microsoft Azure

Microsoft Windows Azure is a cloud operating system built on top of Microsoft datacenters infrastructure and provides developers with a collection of services for building applications with cloud technology. Services range from compute, storage, and networking to application connectivity, access control, and business intelligence. Any application that is built on the Microsoft technology can be scaled using the Azure platform, which integrates the scalability features into the common Microsoft technologies such as Microsoft Windows Server 2008, SQL Server, and ASP.NET.

Figure following provides an overview of services provided by Azure. These services can be managed and controlled through the Windows Azure Management Portal, which acts as an administrative console for all the services offered by the Azure platform. In this section, we present the core features of the major services available with Azure.

Azure core concepts

The Windows Azure platform is made up of a foundation layer and a set of developer services that can be used to build scalable applications. These services cover compute, storage, networking, and identity management, which are tied together by middleware called AppFabric.

This scalable computing environment is hosted within Microsoft datacenters and accessible through the Windows Azure Management Portal. Alternatively, developers can recreate a Windows Azure environment (with limited capabilities) on their own machines for development and testing purposes. In this section, we provide an overview of the Azure middleware and its services.

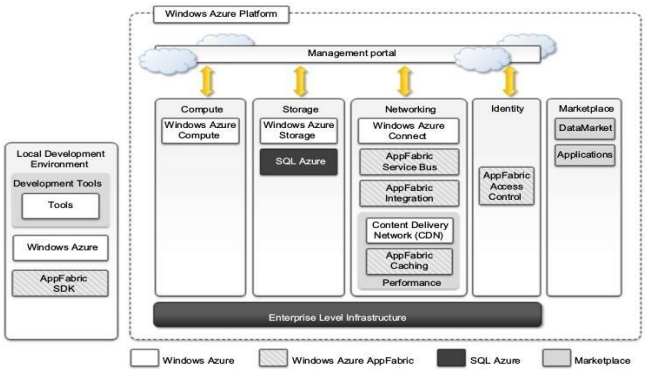

FIGURE 9.3
Microsoft Windows Azure Platform Architecture.

Compute services

Compute services are the core components of Microsoft Windows Azure, and they are delivered by means of the abstraction of roles. A role is a runtime environment that is customized for a specific compute task. Roles are managed by the Azure operating system and instantiated on demand in order to address surges in application demand. Currently, there are three different roles: Web role, Worker role, and Virtual Machine (VM) role.

Web role

The Web role is designed to implement scalable Web applications. Web roles represent the units of deployment of Web applications within the Azure infrastructure. They are hosted on the IIS 7 Web Server, which is a component of the infrastructure that supports Azure. When Azure detects peak loads in the request made to a given application, it instantiates multiple Web roles for that application and distributes the load among them by means of a load balancer.

Since version 3.5, the .NET technology natively supports Web roles; developers can directly develop their applications in Visual Studio, test them locally, and upload to Azure. It is possible to develop ASP.NET (ASP.NET Web Role and ASP.NET MVC 2 Web Role) and WCF (WCF Service Web Role) applications. Since IIS 7 also supports the PHP runtime environment by means of the FastCGI module, Web roles can be used to run and scale PHP Web applications on Azure (CGI Web Role).

Other Web technologies that are not integrated with IIS can still be hosted on Azure (i.e., Java Server Pages on Apache Tomcat), but there is no advantage to using a Web role over a Worker role.

Worker role

Worker roles are designed to host general compute services on Azure. They can be used to quickly provide compute power or to host services that do not communicate with the external world through HTTP. A common practice for Worker roles is to use them to provide background processing for Web applications developed with Web roles.

Developing a worker role is like a developing a service. Compared to a Web role whose computation is triggered by the interaction with an HTTP client (i.e., a browser), a Worker role runs continuously from the creation of its instance until it is shut down. The Azure SDK provides developers

with convenient APIs and libraries that allow connecting the role with the service provided by the runtime and easily controlling its startup as well as being notified of changes in the hosting environment.

As with Web roles, the .NET technology provides complete support for Worker roles, but any technology that runs on a Windows Server stack can be used to implement its core logic. For example, Worker roles can be used to host Tomcat and serve JSP-based applications.

Virtual machine role

The Virtual Machine role allows developers to fully control the computing stack of their compute service by defining a custom image of the Windows Server 2008 R2 operating system and all the service stack required by their applications. The Virtual Machine role is based on the Windows Hyper-V virtualization technology, which is natively integrated in the Windows server technology at the base of Azure. Developers can image a Windows server installation complete with all the required applications and components, save it into a Virtual Hard Disk (VHD) file, and upload it to Windows Azure to create compute instances on demand.

Storage services

Compute resources are equipped with local storage in the form of a directory on the local file system that can be used to temporarily store information that is useful for the current execution cycle of a role. If the role is restarted and activated on a different physical machine, this information is lost.

Windows Azure provides different types of storage solutions that complement compute services with a more durable and redundant option compared to local storage. Compared to local storage, these services can be accessed by multiple clients at the same time and from everywhere, thus becoming a general solution for storage.

Blobs

Azure allows storing large amount of data in the form of binary large objects (BLOBs) by means of the blobs service. This service is optimal to store large text or binary files. Two types of blobs are available:

a) Block blobs. Block blobs are composed of blocks and are optimized for sequential access therefore they are appropriate for media streaming. Currently, blocks are of 4 MB, and a single block blob can reach 200 GB in dimension.
b) Page blobs. Page blobs are made of pages that are identified by an offset from the beginning of the blob. A page blob can be split into multiple pages or constituted of a single page. This type of blob is optimized for random access and can be used to host data different from streaming.

Core infrastructure: AppFabric

AppFabric is a comprehensive middleware for developing, deploying, and managing applica tions on the cloud or for integrating existing applications with cloud services. AppFabric implements an optimized infrastructure supporting scaling out and high availability; sandboxing and multitenancy; state management; and dynamic address resolution and routing. On top of this infra structure, the middleware offers a collection of services that simplify many of the common tasks in a distributed application, such as communication, authentication and authorization, and data access.

These services are available through language-agnostic interfaces, thus allowing developers to build heterogeneous applications.

Access control

AppFabric provides the capability of encoding access control to resources in Web applications and services into a set of rules that are expressed outside the application code base. These rules give a great degree of flexibility in terms of the ability to secure components of the application and define access control policies for users and groups.

Access control services also integrate several authentication providers into a single coherent identity management framework. Applications can leverage Active Directory, Windows Live, Google, Facebook, and other services to authenticate users. This feature also allows easy building of hybrid systems, with some parts existing in the private premises and others deployed in the public cloud.

Service bus

Service Bus constitutes the messaging and connectivity infrastructure provided with AppFabric

for building distributed and disconnected applications in the Azure Cloud and between the private premises and the Azure Cloud. Service Bus allows applications to interact with different protocols and patterns over a reliable communication channel that guarantees delivery.

The service is designed to allow transparent network traversal and to simplify the development of loosely coupled applications, without renouncing security and reliability and letting developers focus on the logic of the interaction rather than the details of its implementation. Service Bus allows services to be available by simple URLs, which are untied from their deployment location.

It is possible to support publish-subscribe models, full-duplex communications point to point as well as in a peer-to-peer environment, unicast and multicast message delivery in one-way communications, and asynchronous messaging to decouple application components.

In order to leverage these features, applications need to be connected to the bus, which provides these services. A connection is the Service Bus element that is priced by Azure on a pay-as-you-go basis. Users are billed on a connections-per-month basis, and they can buy advance "connection packs," which have a discounted price, if they can estimate their needs in advance.

Azure cache

Windows Azure provides a set of durable storage solutions that allow applications to persist their data. These solutions are based on disk storage, which might constitute a bottleneck for the applications that need to gracefully scale along the clients' requests and dataset size dimensions.

Azure Cache is a service that allows developers to quickly access data persisted on Windows Azure storage or in SQL Azure. The service implements a distributed in-memory cache of which the size can be dynamically adjusted by applications according to their needs.

It is possible to store any .NET managed object as well as many common data formats (table rows, XML, and binary data) and control its access by applications. Azure Cache is delivered as a service, and it can be Microsoft Azure easily integrated with applications. This is a particularly true for ASP.NET applications, which already integrate providers for session state and page output caching based on Azure Cache.

The service is priced according the size of cache allocated by applications per month, despite their effective use of the cache. Currently, several cache sizes are available, ranging from 128 MB ($45/month) to 4 GB ($325/month).

SQL Azure

SQL Azure is a relational database service hosted on Windows Azure and built on the SQL Server technologies. The service extends the capabilities of SQL Server to the cloud and provides developers with a scalable, highly available, and fault-tolerant relational database. SQL Azure is accessible from either the Windows Azure Cloud or any other location that has access to the Azure Cloud.

It is fully compatible with the interface exposed by SQL Server, so applications built for SQL Server can transparently migrate to SQL Azure. Moreover, the service is fully manageable using REST APIs, allowing developers to control databases deployed in the Azure Cloud as well as the firewall rules set up for their accessibility.

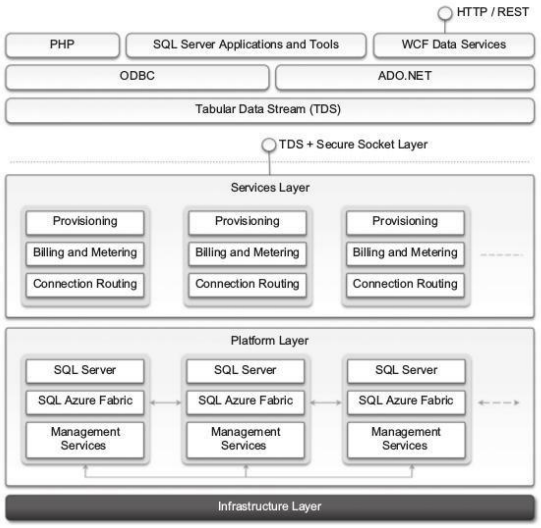

FIGURE 9.4
SQL Azure architecture.

On the SQL Azure side, access to data is mediated by the service layer, which provides provisioning, billing, and connection-routing services. These services are logically part of server instances, which are managed by SQL Azure Fabric. This is the distributed database middleware that constitutes the infrastructure of SQL Azure and that is deployed on Microsoft datacenters.

Developers have to sign up for a Windows Azure account in order to use SQL Azure. Once the account is activated, they can either use the Windows Azure Management Portal or the REST APIs to

create servers and logins and to configure access to servers. SQL Azure servers are abstractions that closely resemble physical SQL Servers: They have a fully qualified domain name under the database.windows.net (i.e., server-name.database.windows.net) domain name. This simplifies the management tasks and the interaction with SQL Azure from client applications.

SQL Azure ensures that multiple copies of each server are maintained within the Azure Cloud and that these copies are kept synchronized when client applications insert, update, and delete data on them.

Currently, the SQL Azure service is billed according to space usage and the type of edition. Currently, two different editions are available: Web Edition and Business Edition. The former is suited for small Web applications and supports databases with a maximum size of 1 GB or 5 GB.

The latter is suited for independent software vendors, line-of-business applications, and enterprise applications and supports databases with a maximum size from 10 GB to 50 GB, in increments of 10 GB. Moreover, a bandwidth fee applies for any data transfer trespassing the Windows Azure Cloud or the region where the database is located. A monthly fee per user/database is also charged and is based on the peak size the database reaches during the month.

Amazon EC2

Amazon.com is one of the most important and heavily trafficked Web sites in the world. It provides a vast selection of products using an infrastructure based on Web services. As Amazon.com has grown, it has dramatically grown its infrastructure to accommodate peak traffic times. Over time the company has made its network resources available to partners and affiliates, which also has improved its range of products.

Starting in 2006, Amazon.com made its Web service platform available to developers on a usage-basis model.

Through hardware virtualization on Xen hypervisors, Amazon.com has made it possible to create private virtual servers that you can run worldwide. These servers can be provisioned with almost any kind of application software you might envisage, and they tap into a range of support services that not only make distributed cloud computing applications possible, but make them robust. Some very large Web sites are running on Amazon.com's infrastructure without their client audience being any the

wiser.

Amazon Web Services is based on SOA standards, including HTTP, REST, and SOAP transfer protocols, open source and commercial operating systems, application servers, and browser-based access. Virtual private servers can provision virtual private clouds connected through virtual private networks providing for reasonable security and control by the system administrator.

AWS has a great value proposition: You pay for what you use. While you may not save a great deal of money over time using AWS for enterprise class Web applications, you encounter very little barrier to entry in terms of getting your site or application up and running quickly and robustly. AWS has much to teach us about the future of cloud computing and how virtual infrastructure can be best leveraged as a business asset. Understanding Amazon Web Services The Amazon is the world's largest river.

Amazon.com is the world's largest online retailer with net sales in $24.51 billion, according to their 2009 annual report. The company is a long way past selling books and records. While Amazon.com is not the earth's biggest retailer (that spot is reserved for Wal-Mart), Amazon.com offers the largest number of retail product SKUs through a large ecosystem of partnerships. By any measure, Amazon.com is a huge business.

To support this business, Amazon.com has built an enormous network of IT systems to support not only average, but peak customer demands. Amazon Web Services (AWS) takes what is essentially unused infrastructure capacity on Amazon.com's network and turns it into a very profitable business.

AWS is having enormous impact in cloud computing. Indeed, Amazon.com's services represent the largest pure Infrastructure as a Service (IAAS) play in the marketplace today. The structure of Amazon.com's Amazon Web Services (AWS) is therefore highly educational in understanding just how disruptive cloud computing can be to traditional fixed asset IT deployments, how virtualization enables a flexible approach to system rightsizing, and how dispersed systems can impart reliability to mission critical systems.

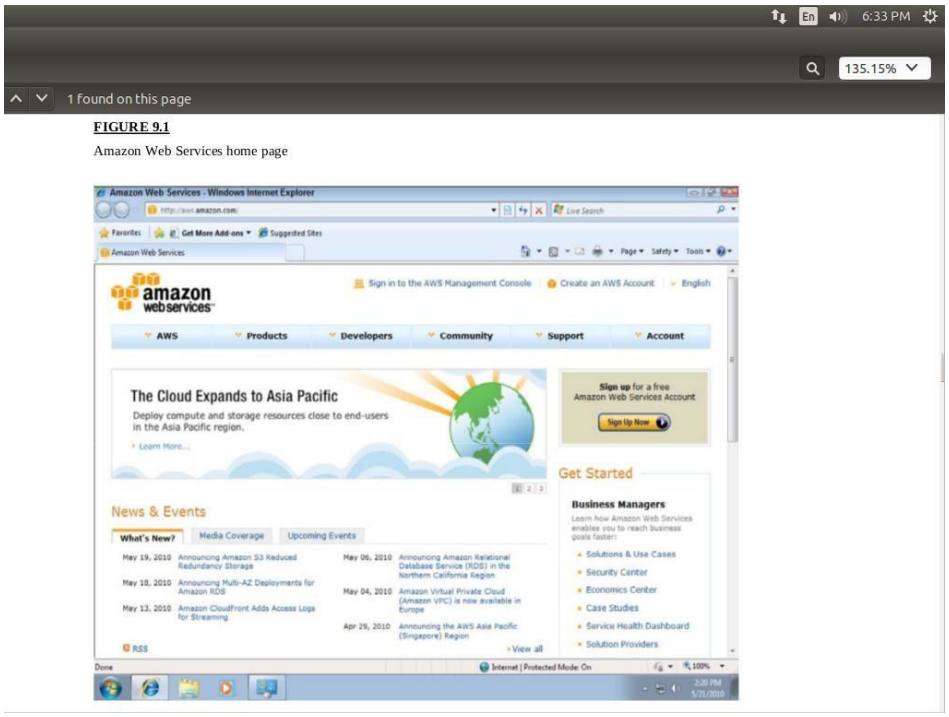

FIGURE 9.1
Amazon Web Services home page

Amazon Web Services represents only a small fraction of Amazon's overall business sales at the moment, but it is a rapidly growing component. Amazon doesn't break down its sales by individual areas in its annual report, but according to Randy Bias who blogs on the site Cloudscaling.com(http://cloudscaling.com/blog/cloud-computing/amazons-ec2-generating20manually) the largest component of Amazon's offerings is Amazon's Elastic Compute Cloud (EC2), which generates in excess of $220 million annually as of October 2009. EC2 is estimated to run on over 40,000+ servers worldwide divided into six availability zones. (You learn about EC2 later in this chapter.) EC2 is an Infrastructure as a Service (IaaS) play, a market that was pegged to be around $400-$600 M/year and growing 10%-20%/year even in the face of a dramatic market slowdown. Rackspace Cloud (http://www rackspacecloud.com/), EC2's nearest competitor, is pegged to be around 10% the size of EC2 by Bias.

Amazon Elastic Compute Cloud (EC2; http://aws.amazon.com/ec2/), is the central application in the AWS portfolio. It enables the creation, use, and management of virtual private servers running the Linux or Windows operating system over a Xen hypervisor. Amazon Machine Instances are sized at

various levels and rented on a computing/hour basis. Spread over data centers worldwide, EC2 applications may be created that are highly scalable, redundant, and fault tolerant.

Working with the Elastic Compute Cloud (EC2)

Amazon Elastic Compute Cloud (EC2) is a virtual server platform that allows users to create and run virtual machines on Amazon's server farm. With EC2, you can launch and run server instances called Amazon Machine Images (AMIs) running different operating systems such as Red Hat Linux and Windows on servers that have different performance profiles. You can add or subtract virtual servers elastically as needed; cluster, replicate, and load balance servers; and locate your different servers in different data centers or "zones" throughout the world to provide fault tolerance. The term elastic refers to the ability to size your capacity quickly as needed.

The difference between an instance and a machine image is that an instance is the emulation of a hardware platform such as X86, IA64, and so on running on the Xen hypervisor. A machine image is the software and operating system running on top of the instance. A machine image may be thought of as the contents of a boot drive, something that you could package up with a program such as Ghost, Acronis, or TrueImage to create a single file containing the exact contents of a volume. A machine image should be composed of a hardened operating system with as few features and capabilities as possible and locked down as much as possible.

Consider a situation where you want to create an Internet platform that provides the following:• A high transaction level for a Web application

• A system that optimizes performance between servers in your system

• Data driver information services

• Network security

• The ability to grow your service on demand

Implementing that type of service might require a rack of components that included the following:

• An application server with access to a large RAM allocation

• A load balancer, usually in the form of a hardware appliance such as F5's BIG-IP

• A database server

• Firewalls and network switches

• Additional rack capacity at the ISP

A physical implementation of these components might cost you something in the neighborhood of $25,000 depending upon the scale of your application. With AWS, you might be able to have an equivalent service for as little as $1,000 and have a high level of availability and reliability to boot. This difference may surprise you, but it is understandable when you consider that AWS can run its services with a much greater efficiency than your company would alone and therefore amortize its investment in hardware over several customers. That is the promise and the potential of cloud computing realized and why large Web sites such as Recovery.gov have moved to AWS.

Amazon Machine Images

AMIs are operating systems running on the Xen virtualization hypervisor. Each virtual private server is accorded a size rating called its EC2 Compute Unit, which is pegged to the equivalent of a 1.0–1.2 GHz 2007 Opteron or 2007 Xeon processor. Table 9.1 shows the current set of Instance types, which broadly fall into the following three classes:

1. Standard Instances: The standard instances are deemed to be suitable for standard server applications.

2. High Memory Instances: High memory instances are useful for large data throughput applications such as SQL Server databases and data caching and retrieval.

3. High CPU Instances: The high CPU instance category is best used for applications that are processor- or compute-intensive. Applications of this type include rendering, encoding, data analysis, and others.

Pricing models

The pricing of these different AMI types depends on the operating system used, which data center the AMI is located in (you can select its location), and the amount of time that the AMI runs. Rates are quoted based on an hourly rate. Additional charges are applied for:

• the amount of data transferred

• whether Elastic IP Addresses are assigned

• your virtual private server's use of Amazon Elastic Block Storage (EBS)

• whether you use Elastic Load Balancing for two or more servers

• other features

AMIs that have been saved and shut down incur a small one-time fee, but do not incur additional hourly fees.

The three different pricing models for EC2 AMIs are as follows:

• On-Demand Instance: This is the hourly rate with no long-term commitment.

• Reserved Instances: This is a purchase of a contract for each instance you use with a significantly lower hourly usage charge after you have paid for the reservation.

• Spot Instance: This is a method for bidding on unused EC2 capacity based on the current spot price. This feature offers a significantly lower price, but it varies over time or may not be available when there is no excess capacity.

Pricing varies by zone, instance, and pricing model. A chart of the different current prices may be found at http://aws.amazon.com/ec2/. This page also includes current Amazon Elastic Block Store volume and snapshot charges to Amazon S3, as well as data transfer rates. Figure below shows the AWS Simple Monthly Calculator that you can find at http://calculator.s3.amazonaws.com/calc5 html to help you estimate your monthly charges.

System images and software

You can choose to use a template AMI system image with the operating system of your choice or create your own system image that contains your custom applications, code libraries, settings, and data. Security can be set through passwords, Kerberos tickets, or certificates.

These operating systems are offered:

• Red Hat Enterprise Linux

• OpenSuse Linux

• Ubuntu Linux

• Sun OpenSolaris

• Fedora

• Gentoo Linux

• Oracle Enterprise Linux

• Windows Server 2003/2008 32-bit and 64-bit up to Data Center Edition

• Debian

Most of the system image templates that Amazon AWS offers are based on Red Hat Linux, Windows Server, Oracle Enterprise Linux, and OpenSolaris from the list above. Table 9.2 lists some of the more common enterprise applications that are available from AWS either as part of its canned templates or for use in building your own AMI system image. Hundreds of free and paid AMIs can be found on AWS.

System images and software

You can choose to use a template AMI system image with the operating system of your choice or create your own system image that contains your custom applications, code libraries, settings, and data.

Security can be set through passwords, Kerberos tickets, or certificates.

These operating systems are offered:

• Red Hat Enterprise Linux

• OpenSuse Linux

• Ubuntu Linux

• Sun OpenSolaris

• Fedora

• Gentoo Linux

• Oracle Enterprise Linux

• Windows Server 2003/2008 32-bit and 64-bit up to Data Center Edition

• Debian

Most of the system image templates that Amazon AWS offers are based on Red Hat Linux, Windows Server, Oracle Enterprise Linux, and OpenSolaris from the list above.

When you create a virtual private server, you can use the Elastic IP Address feature to create what amounts to a static IP v4 address to your server. This address can be mapped to any of your AMIs and is associated with your AWS account. You retain this IP address until you specifically release it from your AWS account. Should a machine instance fail, you can map your Elastic IP Address to fail over to a different AMI. You don't need to wait until a DNS server updates the IP record assignment, and you can use a form to configure the reverse DNS record of the Elastic IP address change.

There are currently four different EC2 service zones or regions:

• US East (Northern Virginia)

• US West (Northern California)

• EU (Ireland)

• Asia Pacific (Singapore)

Google Web Services:

Google is the prototypical cloud computing Services Company, and it supports some of the largest Web sites and services in the world. In this chapter, you learn about Google's applications and

services for users and the various developer tools that Google makes available. At the center of Google's core business is the company's search technology. Google uses automated technology to index the Web. It makes its search service available to users as a standard search engine and to developers as a collection of special search tools limited to various areas of content. The application of Google's searches to content aggregation has led to enormous societal changes and to a growing trend of disintermediation.

The most important commercial part of Google's activities is its targeting advertising business: AdWords and AdSense. Google has developed a range of services including Google Analytics that supports its targeted advertising business.

Google applications are cloud-based applications. The range of application types offered by Google spans a variety of types: productivity applications, mobile applications, media delivery, social interactions, and many more. The different applications are listed in this chapter. Google has begun to commercialize some of these applications as cloud-based enterprise application suites that are being widely adopted.

Google has a very large program for developers that span its entire range of applications and services. Among the services highlighted are Google's AJAX APIs, the Google Web Toolkit, and in particular Google's relatively new Google Apps Engine hosting service. Using Google App Engine, you can create Web applications in Java and Python that can be deployed on Google's infrastructure and scaled to a large size.

Exploring Google Applications

Few companies have had as much impact on their industries as Google has had on the computer industry and on the Internet in particular. Some companies may have more Internet users (Microsoft comes to mind) or have a stock valuation higher than Google (Apple currently fits that description), but Google remains both a technology and thought leader for all things Internet.

For a company whose motto is “Don't be evil,” the impact of consumer tracking and targeted advertising, free sourcing applications, and the relentless assault on one knowledge domain after another has had a profound impact on the lives of many people. I call it the Google Effect.

The bulk of Google's income comes from the sales of target advertising based on information that Google gathers from your activities associated with your Google account or through cookies

placed on your system using its AdWords system. In 2009, Google's revenue was $23.6 billion, and it controlled roughly 65 percent of the search market through its various sites and services.

The company is highly profitable, and that has allowed Google to create a huge infrastructure as well as launch many free cloud-based applications and services that this chapter details. These applications are offered mostly on a free usage model that represents Google's Software as a Service portfolio. A business model that offers cloud-based services for free that are "good enough" is very compelling. While Google is slowly growing a subscription business selling these applications to enterprises, its revenue represents only a small but growing part of Google's current income.

Google's cloud computing services falls under two umbrellas. The first and best-known offerings are an extensive set of very popular applications that Google offers to the general public. These applications include Google Docs, Google Health, Picasa, Google Mail, Google Earth, and many more. You can access a jump table of Google's cloud-based user applications by following the "More" and "Even More" links on Google's home page to the More Google Products page at http://www.google.com/intl/en/options/ shown in Figure.

The second of Google's cloud offerings is its Platform as a Service developer tools. In April 2008, Google introduced a development platform for hosted Web applications using Google's infrastructure called the Google App Engine (GAE). The goal of GAE is to allow developers to create and deploy Web applications without worrying about managing the infrastructure necessary to have their applications run. GAE applications may be written using many high-level programming languages (most prominently Java and Python) and the Google App Engine Framework, which lowers the amount of development effort required to get an application up and running. Goggle also allows a certain free level of service so that the application must exceed a certain level of processor load, storage usage, and network bandwidth (Input/Output) before charges are assessed.

More Google Products equals fewer commercial products.

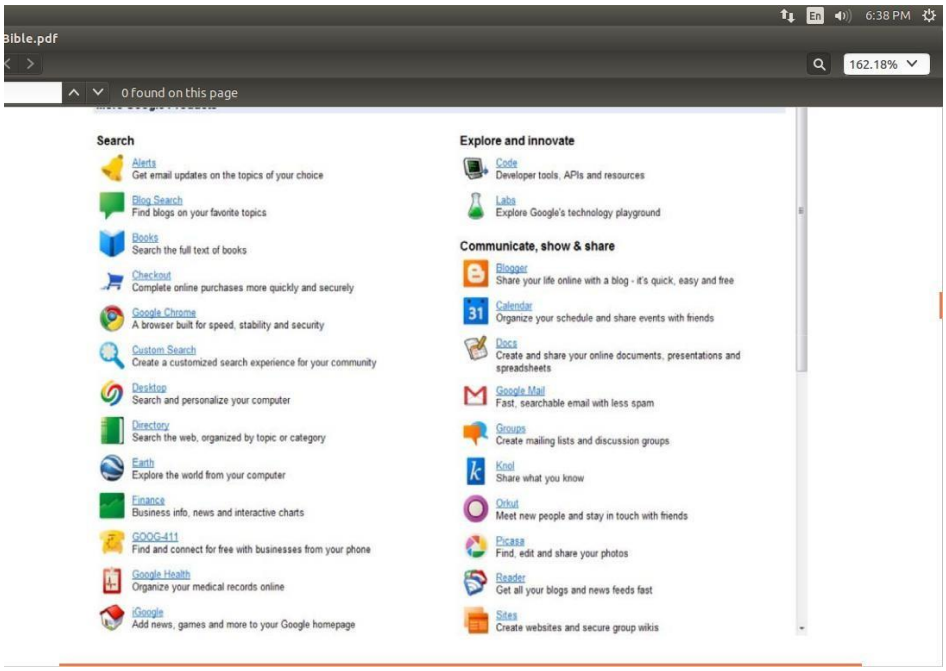

Google App Engine applications must be written to comply with Google's infrastructure. This narrows the range of application types that can be run on GAE; it also makes it very hard to port applications to GAE. After an application is deployed on GAE, it is also difficult to port that application to another platform.

Even with all these limitations, the Google App Engine provides developers a low-cost option on which to create an application that can run on a world-class cloud infrastructure with all the attendant benefits that this type of deployment can bestow.

Google Application Portfolio

It is fair to say that nearly all the products in Google's application and service portfolio are cloud computing services in that they all rely on systems staged worldwide on Google's one million plus servers in nearly 30 datacenters. Roughly 17 of the 48 services listed leverage Google's search engine in some specific way. Some of these search-related sites search through selected content such as Books, Images, Scholar, Trends, and more. Other sites such as Blog Search, Finance, News, and some others take the search results and format them into an Aggregation page. Figure shows one of these aggregation pages: Google Finance.

Google's Finance page at http://www.google.com/finance/ is an example of an aggregation page provided by results from Google's search engine.

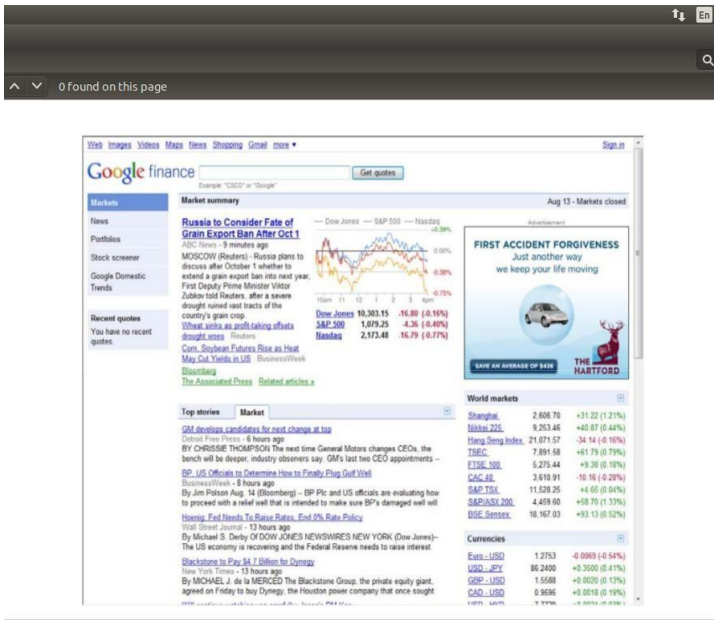

Indexed search

Google's search technology is based on automated page indexing and information retrieval by Web crawlers, also called spiders or robots. Content on pages is scanned up to a certain number of words and placed into an index. Google also caches copies of certain Web pages and stores copies of documents it finds such as DOC or PDF files in its cache.

Google uses a patented algorithm to determine the importance of a particular page based on the number of quality links to that page from other sites, along with other factors such as the use of keywords, how long the site has been available, and traffic to the site or page. That factor is called the PageRank, and the algorithm used to determine PageRank is a trade secret. Google is always tweaking the algorithm to prevent Search Engine Optimization (SEO) strategies from gaming the system. Based on this algorithm, Google returns what is called a Search Engine Results Page (SERP) for a query that is parsed for its keywords.

It is really important to understand what Google (and other search engines) offers and what it doesn't offer. Google does not search all sites. If a site doesn't register with the search engine or isn't the target of a prominent link at another site, that site may remain undiscovered. Any site can place directions in their ROBOTS.TXT file indicating whether the site can be searched or not, and if so what pages can be searched. Google developed something called the Sitemaps protocol, which lets a Web site list in an XML file information about how the Google robot can work with the site. Sitemaps can be useful in allowing content that isn't browsable to be crawled; they also can be useful as guides to finding media information that isn't normally considered, such as AJAX, Flash, or Silverlight media. The Sitemaps protocol has been widely adopted in the industry.

Online content that isn't indexed by search engines belongs to what has come to be called the "Deep Web"—that is, content on the World Wide Web that is hidden. Any site that suppresses Web crawlers rom indexing it is part of the Deep Web. You need go no further than the world's number two Web site, Facebook, for a prominent example of a site that isn't indexed in search engines.

Entire networks exist that aren't searchable, particularly peer-to-peer networks. Ian Clarke's Freenet, which is a P2P network, supports both "darknet" and "opennet" connections. Freenet (http://freenetproject.org/) has been downloaded by millions of people.

The Deep Web includes:

• Database generated Web pages or dynamic content

• Pages without links

• Private or limited access Web pages and sites

• Information contained in sources available through executable code such as JavaScript

• Documents and files that aren't in a form that can be searched, which includes not only media files, but information in non-standard file formats Although efforts are underway to enable information on the Deep Web to be searchable, the amount of information stored that is not accessible is many times larger than the amount of information that can currently be accessed. Some estimates at the size of the Dark Web suggest that it could be an order of magnitude larger than the content contained in the world's search engines.

Enterprise offerings

As Google has built out its portfolio, it has released special versions of its products for the enterprise.

The following are among Google's products aimed at the enterprise market:

• Google Commerce Search (http://www.google.com/commercesearch/): This is a search service for online retailers that markets their products in their site searches with a number of navigation, filtering, promotion, and analytical functions.

• Google Site Search (http://www.google.com/sitesearch/): Google sells its search engine customized for enterprises under the Google Site Search service banner. The user enters a search string in the site's search, and Google returns the results from that site.

• Google Search Appliance (http://www.google.com/enterprise/gsa): This server can be deployed within an organization to speed up both local (Intranet) and Internet searching. The three versions of the Google Search Appliance can store an index of up to 300,000 (GB-1001), 10 million (GB-5005), or 30 million (GB-8008) documents. Beyond indexing, these appliances have document management features, perform custom searches, cache content, and give local support to Google Analytics and Google Sitemaps.

• Google Mini (http://www.google.com/enterprise/mini/): The Mini is the smaller version of the GSA that stores 300,000 indexed documents.

AdWords

AdWords (http://www.google.com/AdWords) is a targeted ad service based on matching advertisers and their keywords to users and their search profiles. This service transformed Google from a competent search engine into an industry giant and is responsible for the majority of Google's revenue stream. AdWords' two largest competitors are Microsoft adcenter (http://adcenter microsoft.com/) and Yahoo! Search Marketing (http://searchmarketing.yahoo.com/).

Ads are displayed as text, banners, or media and can be tailored based on geographical location, frequency, IP addresses, and other factors. AdWords ads can appear not only on Google.com, but on AOL search, Ask.com, and Netscape, along with other partners. Other partners belonging to the Google Display Network can also display AdSense ads. In all these cases, the AdWords system determines which ads to match to the user searches.

Here's how the system works: Advertisers bid on keywords that are used to match a user to their product or service. If a user searches for a term such as “develop abdominal muscles,” Google returns products based on those terms. You might see an ad with Chuck Norris selling a modern-day version of a torture rack that, if it doesn't give you a six-pack, at least makes your wallet lighter. Up to 12 ads per search can be returned.

Google gets paid for the ad whenever a user clicks it. The system is referred to as pay-per-click advertising, and the success of the ad is measured by what is called the click-through rate (CTR). Google calculates a quality score for ads based on the CTR, the strength of the connection between the ad and the keywords, and the advertiser's history with Google. This quality score is a Google trade secret and is used to price the minimum bid of a keyword.

Working with the Google App Engine

Google App Engine (GAE) is a Platform as a Service (PaaS) cloud-based Web hosting service on Google's infrastructure. This service allows developers to build and deploy Web applications and have Google manage all the infrastructure needs, such as monitoring, failover, clustering, machine instance management, and so forth. For an application to run on GAE, it must comply with Google's platform standards, which narrows the range of applications that can be run and severely limits those

applications' portability.

GAE supports the following major features:

• Dynamic Web services based on common standards

• Automatic scaling and load balancing• Authentication using Google's Accounts API

• Persistent storage, with query access sorting and transaction management features

• Task queues and task scheduling

• A client-side development environment for simulating GAE on your local system

• One of either two runtime environments: Java or Python

When you deploy an application on GAE, the application can be accessed using your own domain name or using the Google Apps for Business URL.

Google App Engine currently supports applications written in Java and in Python, although there are plans to extend support to more languages in the future. The service is meant to be language-agnostic. A number of Java Virtual Machine languages are compliant with GAE, as are several Python Web frameworks that support the Web Server Gateway Interface (WSGI) and CGI. Google has its own Webapp framework designed for use with GAE.

The AppScale (http://appscale.cs.ucsb.edu/) open-source framework also may be used for running applications on GAE.To encourage developers to write applications using GAE, Google allows for free application development and deployment up to a certain level of resource consumption. Resource limits are described on Google's quota page at http://code.google.com/appengine/docs/quotas.html, and the quota changes from time to time.

Google uses the following pricing scheme:

• CPU time measured in CPU hours is $0.10 per hour.

• Stored data measured in GB per month is $0.15 per GB/month.

• Incoming bandwidth measured in GB is $0.10 per GB.

• Outgoing bandwidth measured in GB is $0.12 per GB.

• Recipients e-mailed is $0.0001 per recipient.

Applications running in GAE are isolated from the underlying operating system, which Google describes as running in a sandbox. This allows GAE to optimize the system so Web requests can be matched to the current traffic load. It also allows applications to be more secure because applications

can connect only to computers using the specified URLs for the e-mail and fetch services using HTTP or HTTPS over the standard well-known ports. URL fetch uses the same infrastructure that retrieves Web pages on Google. The mail service also supports Gmail's messaging system.

Applications also are limited in that they can only read files; they cannot write to the file system directly. To access data, an application must use data stored in the memcache (memory cache), the datastore, or some other persistent service. Memcache is a fast in-memory key-value cache that can be used between application instances. For persistent data storage of transactional data, the datastore is used. Additionally, an application responds only to a specific HTTP request—in real-time, part of a queue, or scheduled—and any request is terminated if the response requires more than 30 seconds to complete.

GAE has a distributed datastore system that supports queries and transactions. This datastore is non-relational or "schema-less," but it does store data objects or entities that are assigned properties. In your queries, you can use entities filtered by kind or type and also sorted by properties.

The datastore uses an optimistic concurrency control and maintains strong consistency. An application can execute transactions with multiple operations, and they either all succeed or fail as a unit. To support the distributed nature of the datastore, the concept of an entity group is employed. Transactions manage entities as a single group, and entity groups are stored together in the system so operations can be performed faster.

The App Engine relies on the Google Accounts API for user authentication, the same system used when you log into a Google account. This provides access to e-mail and display names within your app, and it eliminates the need for an application to develop its own authentication system. Applications can use the User API to determine whether a user belongs to a specific group and even whether that person is an administrator for your application.

Many applications have been built and are running on Google App Engine. To get some idea of the range of applications that have been developed, you may want to visit the Google App Engine Gallery.

This gallery is found at http://appgallery.appspot.com/ and is shown in Figure. It is searchable by keyword and category.


Google
Code
Search
Sign in
Google App Engine
Home
Docs
FAQ
Articles
Group
Terms
Download
Download the SDK
Visit the App Gallery
Introduction
What Is Google App Engine?
Getting Started
APIs
The Python Runtime
Datastore API
Images API
Mail API
Memcache API
URL Fetch API
Users API
Tools and Configuration
The webapp Framework
Configuring an App
Configuring Indexes
The Dev Web Server
Uploading an App
The Bulk Data Uploader
The Admin Console
Featured Application
SnapABug
TimZon - Sep 23, 2009
About SnapABug
SnapABug is a widget that can simply be installed on any website. It enables users to easily request support while providing the support team with a screenshot of users web page along with additional debug information; Helping website operator to help their customers faster and more accurately. The SnapABug Help widget is free.
Tags: customer support help widget snapabug
Editor's Picks
Search for apps by tag
Search
Add your own Application
The Google App Engine Application Gallery is a great place to get feedback on your application, find out what others are doing and enable others to see your hard work.
Sharing your application is as easy as filling out a simple form.
Add your app now »
Find apps by category
Tools
Communication
Games
News
Finance
Sports
Lifestyle
Technology
Enterprise

S.NO	QUESTIONS	OPTION 1	OPTION 2	OPTION 3
1	_____ is a cloud operating system built on top of Microsoft datacenters infrastructure	Microsoft Windows Azure	Oracle Azure	VB Azure
2	_____ are the core components of Microsoft	Storage services	Compute services	Product services
3	Expand BLOBs	Binary Large Objects	Blocking Large Objects	Between Large Objects
4	_____ is a comprehensive middleware for developing, deploying, and managing applica	Microsoft Windows Azure	AppFabric	SQL Azure
5	_____is one of the most important and heavily trafficked Web sites in the world	Yahoo.com	Google.com	Amazon.com
6	_____ is a Platform as a Service (PaaS) cloud-based Web hosting service on Google's infrastructure.	SQL Azure	Web Server Gateway Interface	Amazon Web Services
7	_____is a service that allows developers to quickly access data persisted on Windows Azure storage.	Web Server Gateway Interface	Azure Cache	Amazon Web Services
8	_____ supports databases with a maximum size of 1 GB or 5 GB.	Web Edition	Business Edition	Standard Edition
9	_____supports databases with a maximum size from 10 GB	Web Edition	Standard Edition	Business Edition
10	In Which Year did Amazon.com made its Web service platform	2006	2005	2007
11	_____ is the world's largest online retailer with net sales in $24.51 billion, according to their	Google.com	Amazon.com	Force.com
12	_____ is the hourly rate with no long-term commitment.	On-Demand Instance	Reserved Instances	Spot Instance
13	There are currently _____different EC2 service	Three	Four	Five

14	_____is the prototypical cloud computing Services Company.	Google	Yahoo	Amazon
15	_____supports some of the largest Web sites and services in	Yahoo	Amazon	Google
16	SEO stands for_____	Search Engine Optimization	Small Engine Optimization	Secret Engine Optimization
17	Online content that isn't indexed by search engines belongs to what has come to be called the_____	Crawl Web	Deep Web	Dark Web
18	World's number two Web site, is called_____	Facebook	Google	Yahoo
19	_____ is a prominent example of a site that isn't indexed in search engines	Facebook	Google	Yahoo
20	The success of the ad is measured by what is called the_____	Money through Click	Click-through rate	Minimum Click through rate
21	_____ supports Dynamic Web services based on common	SQL Azure	Azure Cloud	Google App Engine
22	_____ can be used to run and scale PHP Web applications on	Worker Roles	Web Roles	Virtual machine Roles
23	_____ can be used to host Tomcat and serve JSP-based	Web Roles	Virtual machine Roles	Worker roles
24	A single block blob can reach ___ in dimension	100 GB	200 GB	150 GB
25	_____type of blob is optimized for random access and can be used to host data different from	Page blobs	Block blobs	Storage blobs
26	Access to SQL Azure is based on the _____ protocol	User Datagram	Transmission Control	Tabular Data Stream
27	_____ is foundation layer + set of developer services	Azure Platform	Google Platform	Amazon Platform
28	_____provides access to e-mail and display names within your	Amazon web services	App Engine	SQL Azure
29	_____eliminates the need for an application to develop its own authentication system	Amazon web services	SQL Azure	App Engine
30	_____has a distributed datastore system that supports queries and transactions.	Google App Engine	Azure Cloud	SQL Azure
31	The Datastore in Google App Engine is _____	Non-relational	Relational	Schema based

32	The _____ uses an optimistic concurrency control and maintains strong consistency.	Entity Group	User API	Datastore
33	_____manage entities as a single group, and entity groups are stored together in the system so operations can be performed	Queries	Datastore	Entity Group
34	Applications can use the_____ to determine whether a user	Admin API	User API	Third party API
35	Applications running in GAE are isolated from the underlying	Sandbox	Web pages	Well known ports
36	The pricing scheme of Google in the Outgoing bandwidth measured in GB is _____ per GB.	$0.20	$0.22	$0.10
37	The pricing scheme of Google in the Incoming bandwidth measured in GB is ____ per GB	$0.20	$0.22	$0.10
38	The pricing scheme of Google in the Stored data measured in GB per month is ____ per GB/month	$0.20	$0.10	$0.12
39	The pricing scheme of Google in the CPU time measured in CPU hours is _____ per hour.	$0.10	$0.22	$0.20
40	WSGI stands for_____	Window Server Gateway Interface	Web Server Gateway Interface	Web Server Gateway Information
41	_____ supports the feature Task queues and task scheduling	Google App Engine	Azure Cloud	SQL Azure
42	_____ can appear not only on Google.com, but on AOL search,	AdSense ads	AdDeep web ads	AdWords ads
43	The ____system determines which ads to match to the user	AdWords	AdSense	AdDeep web
44	Advertisers bid on _____that are used to match a user to their product or service	Exact words	Keywords	Retrieved words
45	_____can be deployed within an organization to speed up both local (Intranet) and Internet	Google Site Search	Google Commerce Search	Google Search Appliance
46	_____is the smaller version of the Google Search Appliance .	Google Mini	Google Commerce	Google Site
47	_____ includes Pages without links.	Crawl Web	Deep Web	Dark Web

48	_____ includes Private or limited access Web pages and	Crawl Web	Dark Web	Deep Web
49	_____ includes Information contained in sources available	Deep Web	Crawl Web	Dark Web
50	_____ can be useful in allowing content that isn't browsable to	Deep Web	Crawl Web	Dark Web
51	Web crawlers are also called as	Spiders or Robots	Dark Robots	Dark Spiders
52	Content on pages is scanned up to a certain number of words and placed into an _____	Header	Index	Table
53	_____is an example of an aggregation page provided by results from Google's search engine	Google's app page	Google's Store page	Google's Finance page
54	_____which lowers the amount of development effort required to get an application up and running.	Azure Cloud Framework	Google App Engine Framework	SQL Azure Framework
55	Google's cloud computing services falls under_____umbrellas	2	3	4
56	The bulk of Google's income comes from the sales of_____	Social Networking sites	Applications	Target advertising
57	Google is always tweaking the algorithm to prevent _____strategies from gaming the system	Search Engine Results Page	Search Engine Data	Search Engine Algorithm
58	_____represent the units of deployment of Web applications	Web Roles	Virtual machine Roles	Worker roles
59	_____service is optimal to store large text or binary files.	Blocks	Blobs	Storage
60	A connection is the Service Bus element that is priced by Azure on a _____basis	Yearly	Weekly	Pay-as-you-go

OPTION 4	ANSWER
Java Azure	Microsoft Windows Azure
Infrastructure services	Compute services
Blocking Last Objects	Binary Large Objects
Oracle Azure	AppFabric
MSN.com	Amazon.com
Google App Engine	Google App Engine
SQL Azure	Azure Cache
Special Edition	Web Edition
Special Edition	Business Edition
2008	2006
Rackspace.com	Amazon.com
Timing Instance	On-Demand Instance
Six	Four

IBM	Google
IBM	Google
Sophisticated Engine Optimization	Search Engine Optimization
Open Web	Deep Web
MSN	Facebook
MSN	Facebook
Maximum through Click	Click-through rate
Amazon EC2	Google App Engine
Storage Roles	Web Roles
Storage Roles	Worker roles
160 GB	200 GB
Cluster blobs	Page blobs
Stream Oriented	Tabular Data Stream
Azure Infrastructure	Azure Platform
Azure Cloud	App Engine
Azure Cloud	App Engine
Amazon EC2	Google App Engine
Standard	Non-relational

Attribute Group	Datastore
Transactions	Transactions
API	User API
Protocols	Sandbox
$0.12	$0.12
$0.12	$0.10
$0.15	$0.15
$0.12	$0.10
Window Server Gateway Information	Web Server Gateway Interface
Amazon EC2	Google App Engine
AdServer ads	AdWords ads
AdServer	AdWords
Keys	Keywords
Google Mini	Google Search Appliance
Google App	Google Mini
Open Web	Deep Web

Open Web	Deep Web
Open Web	Deep Web
Sitemaps	Sitemaps
Shadowers	Spiders or Robots
Database	Index
Google's Box page	Google's Finance page
Amazon EC2 Framework	Google App Engine Framework
5	2
Blogger	Target advertising
Search Engine Optimization	Search Engine Optimization
Storage Roles	Web Roles
Cluster	Blobs
Monthly	Pay-as-you-go