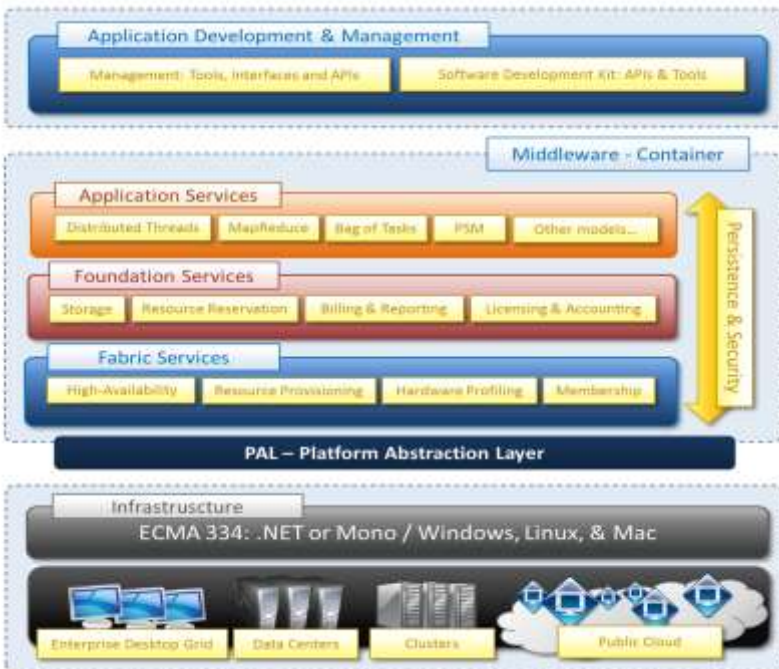


Internal Assessment Test 2 – May 2023

Sub:	Cloud Computing and Applications	Sub Code:	18CS643	Branch:	CSE
Date:	24.05.2023	Duration:	90 mins	Max Marks:	50
		Sem / Sec:	6 A,B,C		OBE

Answer any FIVE FULL Questions

		MARKS	CO	RBT
1 (a)	<p>Explain the Aneka Framework with a neat diagram. Explain anatomy of the Aneka framework and PAL and 3 types of services.</p>  <p>The diagram illustrates the Aneka Framework architecture, which is layered as follows from top to bottom:</p> <ul style="list-style-type: none"> Application Development & Management: Includes Management (Tools, Interfaces and APIs) and Software Development Kit (APIs & Tools). Middleware - Container: This layer is divided into three service categories: <ul style="list-style-type: none"> Application Services: Distributed Threads, MapReduce, Bag of Tasks, PSM, Other models... Foundation Services: Storage, Resource Reservation, Billing & Reporting, Licensing & Accounting. Fabric Services: High-Availability, Resource Provisioning, Hardware Profiling, Membership. A vertical arrow on the right side of the Middleware - Container layer is labeled "Persistence & Security". PAL – Platform Abstraction Layer: A dark blue bar separating the services from the infrastructure. Infrastructure: ECMA 334: .NET or Mono / Windows, Linux, & Mac. This layer includes: <ul style="list-style-type: none"> Enterprise Desktop Grid Data Centers Clusters Public Cloud 	7M	CO2	L2
	<ul style="list-style-type: none"> • Services operate at container level • Except for platform abstraction layer <ul style="list-style-type: none"> ◦ they provide developers, users, and administrators with all features offered by the framework • Services also constitute the extension and customization point of Aneka Clouds • Infrastructure allows for integration • The framework includes the basic services for infrastructure and node management, application execution,, accounting and system monitoring • is a lightweight software layer designed to host services and interact with the underlying operating system and hardware • Role: deploy services and some basic capabilities such as communication channels through which it interacts with other nodes in the Aneka Cloud. • Almost all operations performed within Aneka are carried out by the services managed by the container. • Services stack resides on top of the Platform Abstraction Layer(PAL) • It provides a uniform view of the soft- ware and hardware environment in which the container is running 			

- The PAL is responsible for **detecting the supported hosting environment and providing the corresponding implementation to interact with it to support the activity of the container**
- The core infrastructure is based on .NET technology and allows the Aneka container to be portable over different platforms and operating system.
- **ECMA-334 / ECMA-335** : compatible environment can host and run an instance of the Aneka container.
- **Common Language Infrastructure(CLI)** (Introduced in ECMA-334 standard) : defines a common runtime environment and application model for executing programs but does not provide any interface to access the hardware or to collect performance data from the hosting operating system
- each operating system has a different file system organization and stores that information differently
- PAL addresses this heterogeneity
- provides the container with a uniform interface for accessing the relevant hardware and
- operating system information
- Allows the rest of the container to run unmodified on any supported platform.

Fabric Services

- lowest level of the software stack representing the Aneka Container
- They provide access to the resource-provisioning subsystem and to the
- monitoring facilities implemented in Aneka.
- **Resource-provisioning services** :dynamically provide new nodes on demand by relying on virtualization technologies.
- **monitoring services** allow for *hardware profiling*
- implement a basic monitoring infrastructure that can be used by **all the**
- **services installed in the container**

Foundation Services

- related to the **logical management of the**
- **distributed system** built on top of the infrastructure and provide **supporting services** for the execution of distributed applications.
- Foundation Services provide a **uniform approach to managing distributed applications**
- allow developers to concentrate only on the **logic** that distinguishes a specific
- programming model from the others.
- The Fabric Services and Foundation
- Services constitute the **core of the Aneka**
- **middleware**

Application Services

- manage the **execution of applications**
- A layer that differentiates according to the specific **programming model**

- used for developing distributed applications on top of Aneka.
- Two major types of activities that are common across all the supported models.
- Scheduling
- Service
- Application Services
- Execution

(b)i) What is the use of Aneka SDK?
 provides APIs for developing applications on top of existing programming models

ii) What are the 3 models for tasks generated by the user in Aneka SDK?
 Task Model, Thread Model, Parameter Sweep model

iii) What is used when the workUnit is generated by runtime rather than user?
 MapReduce

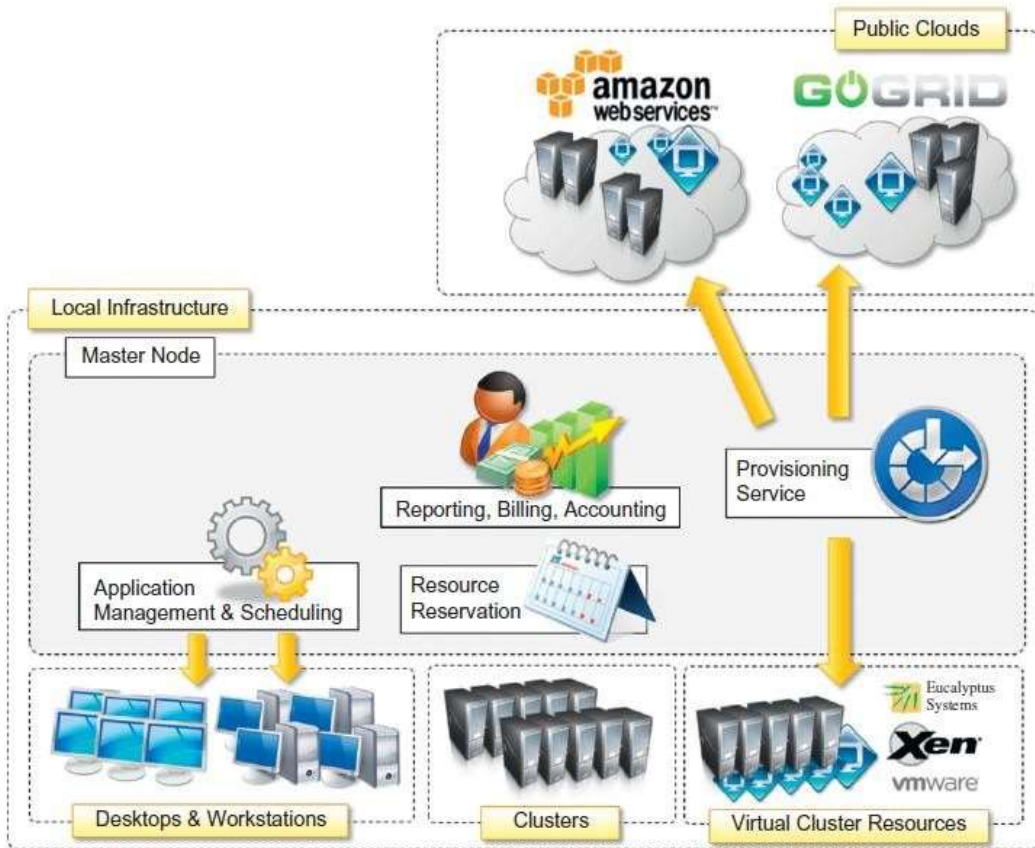
iv) What are W and M in AnekaApplication<W,M>
 Worker, Application Manager

v) List any 5 properties in the WorkUnit class.
 Id, completionTime, Exception, InputFiles, MaximumExecutionTime, Name, NodeId, OutputFiles, Preemptable, Priority, QueuedTime, ReservationId, ResubmitMode, ScheduleTime, State, SubmissionTime, UserCredential

3M CO2 L2

2 (a) Explain the Hybrid cloud deployment of Aneka in detail including the configuration of various nodes with a neat diagram

5M CO2 L2



constitutes the static deployment of Aneka that can be elastically scaled on demand when additional resources are required.

There is already an existing infrastructure that can be used that may consist of desktops and workstations, cluster or virtual cluster resources.

The following scenarios are useful for hybrid cloud:

Dynamic Resource Provisioning

Resource Reservation

Workload Partitioning


Accounting, Monitoring, Reporting

In a hybrid scenario, heterogeneous resources can be used for different purposes. In the case of a private cloud deployment, desktop machines can be reserved for low priority work-load outside the common working hours. Any additional computing capability demand can be primarily addressed by the local virtualization facilities, and if more computing power is required, it is possible to leverage external IaaS providers.

Since part of the infrastructure is local, a cost in data transfer to the external IaaS infrastructure cannot be avoided.


It is then important to select the most suitable option to address application needs.


- **master node** : has all the services.
- Maintained in one single copy
- Provides the intelligence of the Aneka Cloud
- Cloud
- **Index Service (or Membership Catalogue)**
–
- **is mandatory**
- Other services except for those that are mandatory may be present in the master node or other nodes.



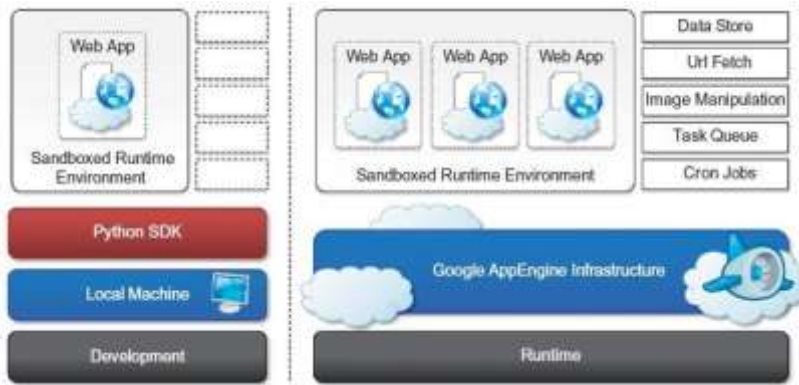
Worker Node

- The **workforce** of the Aneka Cloud
- generally configured for the **execution of applications**.
- **mandatory services**
- **specific execution services** of each of the supported programming models in the Cloud.
- Common Configuration
- Index Service
- Heartbeat Service
- Logging Service
- Allocation Service
- Monitoring Service
- Execution Services for the supported programming models



<p>b)</p>	 <p>Describe the following aspects pertaining to the infrastructure of a public cloud provider.</p> <ol style="list-style-type: none"> i) Regions a set of Datacenters that are connected through a dedicated low-latency network. ii) Availability Zones Availability zones (AZs) are isolated data centers located within specific regions in which public cloud services originate and operate iii) Edge Locations Edge locations are AWS data centers designed to deliver services with the lowest latency possible. iv) Data centers A data center stores and shares applications and data. It comprises components that include switches, storage systems, servers, routers, and security v) Are all services available in all regions? No, some services are not available in all regions. vi) Why is there a difference in pricing for the same services across different regions? The cost of an Azure service can vary between locations based on demand and local infrastructure costs vii) How is a service made highly available with a guaranteed SLA? Locally redundant, zone redundant, geo-redundant. -run shadow services in different zones. 	5M	CO3	L3
<p>3 (a)</p>	<ul style="list-style-type: none"> • Write short notes on AWS • is a platform that allows the development of flexible applications by providing solutions for elastic infrastructure scalability, messaging, and data storage. • accessible through SOAP or RESTful • Web service interfaces. • provides a Web-based console where users can handle administration and monitoring 	6M	CO3	L2

	<p>Services</p> <ul style="list-style-type: none"> • An implementation of a content delivery network on top of the Amazon distributed storage infrastructure. • leverages a collection of edge servers strategically located around the globe. • better serve requests for static and streaming Web content • Communication Services : • Virtual Networking : comprises a collection of services that allow AWS users to control the connectivity to and between compute and storage services. • Amazon Virtual Private Cloud (VPC) and Amazon Direct Connect: provides connectivity solutions in terms of infrastructure 			
(b)	<p>i) Why is S3 Glacier Deep Archive the cheapest storage? It is used for archival purposes and used only when access to data is needed rarely, not more than 90 days or 180 days. The technology may be a mix of tape, disk and flash storage.</p> <p>ii) What concept reduces AWS cloud pricing? Economies of Scale</p> <p>iii) Which service sends and receives messages between applications? Amazon SQS(Simple Queue Service)</p>  <p>iv) What is Amazon Lex? chatbot</p>	2M	CO3	L2
c)	<p>Differentiate on-spot, dedicated and reserved instances.</p> <p>On-spot – bid on a resource based on market demand. Highly affordable</p> <p>Dedicated- give dedicated hardware for a particular organization which is single tenant</p> <p>Reserved Instances – resources that are committed to for 1year or 3 years and purchased in advance.</p>	2M	CO3	L3
4 (a)	<p>Describe core components of Google App Engine</p> <ul style="list-style-type: none"> ▶ a PaaS implementation that provides services for developing and hosting scalable Web applications ▶ AppEngine is essentially a distributed and scalable runtime environment ▶ leverages Google’s distributed infrastructure to scale out applications facing a large number of requests ▶ It allocates more computing resources to them and balancing the load among them ▶ Developers can develop applications in Java, Python, and Go(developed by Google) 	6.5M	CO3	L2



► Infrastructure

AppEngine hosts Web applications

primary function is to serve users **requests efficiently**.

For each **HTTP request**

AppEngine locates the servers hosting the application that processes the request

evaluates their load

if necessary, allocates additional resources

Also monitors **application performance**

Collects **statistics** on which the billing is calculated.

► runtime environment

represents the execution context of applications hosted on AppEngine

Sandboxing : provide the application environment with an **isolated and protected** context

It can execute without **causing threat to the server** or be **influenced by other applications**.

Supported Runtimes : Java, Python, and Go.

► Underlying storage

Storage for semi-structured data

Long term storage for static file servers

► **Static File Servers** : components that define the graphical layout of the application (CSS files, plain HTML files, JavaScript files, images, icons, and sound files) or data files.

► Hosted in static servers that are not modified often

► Servers are optimized for storing static content.

► **Data Store:** allows developers to store semi-structured data


► Designed **to scale and optimized to quickly access** data.

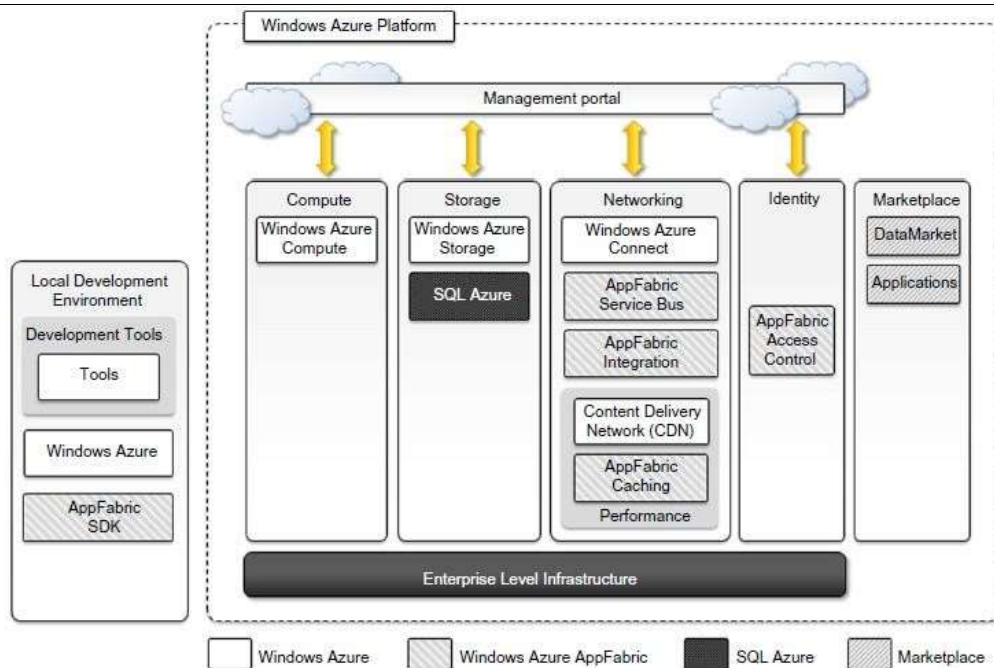
► Object described in terms of **entity and properties**.

► A large object database where the object can be retrieved **using a key**.

► type of the **key and the structure** of the object can vary.

► Provides facilities to **create indexes** on data.

	<ul style="list-style-type: none"> ▶ Uses Optimistic concurrency control: If one user tries to update an entity that is already being updated, the control returns and the operation fails ▶ Set of scalable services : <ul style="list-style-type: none"> • UrlFetch: Applications can make synchronous and asynchronous Web requests • MemCache : <ul style="list-style-type: none"> ○ This is a distributed in-memory cache that is optimized for fast access and provides developers ○ with a volatile store for the objects that are frequently accessed • Mail and instant Messaging <ul style="list-style-type: none"> ○ sends email on behalf of the application to specific user accounts • Account management : using Google account, this allows Web applications to offload the implementation of authentication capabilities to Google's authentication system. • Image Manipulation : Often simple operations, such as adding watermarks or applying simple <ul style="list-style-type: none"> ○ filters, are required • Task Queues: allow applications to submit a task for a later execution • Cron jobs : to perform an operation at a specific time of the day. 			
(b)	<p>i) List any 2 security features implemented by Google to protect the infrastructure/data. multiple security layers for physical security used such as biometric identification, laser-based intrusion detection, etc Ingress and egress filtering at various points in our network helps prevent IP Spoofing use custom hardware chip (Titan) on servers, devices, peripherals to authenticate google devices</p> <p>ii) What is sandboxing? use custom hardware chip (Titan) on servers, devices, peripherals to authenticate google devices</p> <p>iii) No resources in GCP can be used without being associated with a project</p> <p>iv) What are the use of quotas? Prevents uncontrolled consumption of resources</p> <p>v) Which container does Google Kubernetes Engine support? Docker</p> <div style="text-align: center;">  <p>Google Kubernetes Engine</p> </div>	3.5M	CO3	L2
5 (a)	Explain the Windows Azure platform architecture.	[07]	CO 3	L2



AppFabric

- Middleware for developing, deploying, and managing applications on cloud or for integrating existing applications with cloud services.
 - Scaling out and high availability; sandboxing and multitenancy; state management; and dynamic address resolution and routing.
 - Simplify many common tasks in distributed application, such as communication, authentication and authorization, and data access.

Access control

- AppFabric provides capability of encoding access control to resources in Web applications.
- Services into set of rules that are expressed outside application code base.
- Applications can leverage Active Directory, Windows Live, Google, Facebook, and other services to authenticate users.

Service bus

- Messaging and connectivity infrastructure.
- Designed to allow transparent network traversal and to simplify development of loosely coupled applications letting developers focus on logic of interaction.
- Applications need to be connected to bus, which provides these services.


Azure cache

- provides a set of durable storage solutions that allow applications to persist their data.
- Azure Cache is a service that allows developers to quickly access data persisted on Windows Azure storage or in SQL Azure.
- implements a distributed in-memory cache of which, size can be dynamically adjusted by
 - applications.

- **Web Role** – hosted in IIS 7 web server.
 - designed to implement scalable Web applications.
 - Web roles represent the units of deployment of Web applications within the Azure infrastructure
 - NET technology natively supports Web roles
 - 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.
- **Worker role** - designed to host compute services in Azure.
- **Virtual Machine Role** – gives a finer grained control over the virtual resources.
- The Virtual Machine role is based on the Windows Hyper-V virtualization technology
- Developers can image a Windows server installation complete with all the required applications and components, save it into a Virtual Hard Disk (VHD)
- upload it to Windows Azure to create compute instances on demand.

Storage

- Windows Azure provides different types of storage solutions that
- complement compute services with a more durable and redundant option compared to local storage.
- **Blobs**
- Store large amount of data in the form of binary large objects (BLOBs).
- **Azure drive**
 - Entire file system in the form of single Virtual Hard Drive (VHD) file.
 - NTFS file system, providing persistent and durable storage.
- **Tables**
 - semistructured storage solution. Tables are more similar to spreadsheets.
 - Handle large amounts of data and queries returning huge result sets.
 - Currently, table can contain up to 100 TB of data, and rows can have up to 255 properties, with a maximum of 1 MB for each row.
- **Queues**
 - Queue storage allows applications to communicate by exchanging messages through durable queues, thus avoiding lost or unprocessed messages.
 - Applications enter messages into a queue, and other applications can read them in a first-in, first-out
- (FIFO) style.
- Simplify the development and integration of applications.
- **Windows Azure virtual network**
 - includes **Windows Azure Connect** and **Windows Azure Traffic Manager**.

	<ul style="list-style-type: none"> • Windows Azure Connect allows easy setup of IP-based network connectivity among machines hosted on private premises and roles deployed on Azure Cloud. • Windows Azure Traffic Manager provides load-balancing features for services listening to HTTP or HTTPS ports and hosted on multiple roles. • Windows Azure content delivery network <ul style="list-style-type: none"> ○ Content delivery network solution that improves content deliver capabilities of Windows Azure Storage and Microsoft Windows Update, ○ Bing maps. 			
(b)	<p>i) Serverless computing is offered in Azure via Azure Functions</p> <p>ii) Use of Azure CDN? a globally distributed network of servers that can efficiently deliver web content to consumers.</p> <p>iii) How to connect to a Windows VM or Linux VM in Azure? SSH into Linux VM, RDP into Windows VM</p> <p>iv) What is PowerBI? Power BI is a collection of software services, apps, and connectors that work together to turn your unrelated sources of data into coherent, visually immersive, and interactive insights.</p>  <p>v) The typical web role in Windows Server VM in Azure is provided by which server? IIS Server</p>	[03]	CO3	L2
6	<p>Explain storage services in various cloud platforms.</p> <p>Azure Storage</p> <p>Windows Azure provides different types of storage solutions that</p> <ul style="list-style-type: none"> • complement compute services with a more durable and redundant option compared to local storage. • Blobs • Store large amount of data in the form of binary large objects (BLOBs). • Two types of blobs are available: <ul style="list-style-type: none"> ○ Block blobs. composed of blocks optimized for sequential access; ○ Blocks are of 4 MB, and a single block blob can reach 200 GB. ○ Page blobs. pages that are identified by an offset from the beginning of blob. ○ Split into multiple pages or constituted of single page. ○ Optimized for random access. ○ Maximum dimension of a page blob can be 1 TB. • Azure drive 	7M	CO2	L2

- Entire file system in the form of single Virtual Hard Drive (VHD) file.
- NTFS file system, providing persistent and durable storage.
- **Tables**
 - semistructured storage solution. Tables are more similar to spreadsheets.
 - Handle large amounts of data and queries returning huge result sets.
 - Currently, table can contain up to 100 TB of data, and rows can have up to 255 properties, with a maximum of 1 MB for each row.

GCP Storage

- **Static File Servers** : components that define the graphical layout of the application (CSS files, plain HTML files, JavaScript files, images, icons, and sound files) or data files.
 - Hosted in static servers that are not modified often
 - Servers are optimized for storing static content.
- **Data Store**: allows developers to store semi-structured data
 - Designed **to scale and optimized to quickly access** data.
 - Object described in terms of **entity and properties**.
 - A large object database where the object can be retrieved **using a key**.
 - type of the **key and the structure** of the object can vary.
 - Provides facilities to **create indexes** on data.
 - Uses **Optimistic concurrency control**: If one user tries to update an entity that is already being updated, the control returns and the operation fails

Amazon Storage

- S3 Key concepts: accessible through Representational State Transfer (REST) interface
- The storage is organized in a two-level hierarchy
- Stored objects cannot be manipulated like standard files
- Content is not immediately available to users.
- Requests will occasionally fail.
- A bucket is a container of objects.
- virtual drive hosted on the S3 distributed storage
- provides users with a flat store to which they can add objects.
- A bucket is located in a specific geographic location eventually replicated fault tolerance and better content distribution.
- Users create a bucket by sending a PUT request to <http://s3.amazonaws.com/> with the name of the bucket and,
- They may want to specify the availability zone, additional information about the preferred location.
- Content of a bucket can be listed by sending a GET request
- The deletion of a bucket is performed by a DELETE request
- **Amazon RDS**: A relational database service that relies on the EC2 infrastructure and is managed by Amazon

	<ul style="list-style-type: none"> ○ 2 features : multi-AZ deployment and read replicas ○ Optimal for Oracle and MySQL migrated to AWS ● Amazon SimpleDB: a lightweight, highly scalable, and flexible data storage solution for applications that do not require a fully relational model for their data 			
	<p>i) What does BLOB stand for? Binary Large Object</p> <p>ii) Which access tier for storage is optimized for highest storage costs, but the lowest access costs (Hot/Cool/Cold/Archive) Hot</p> <p>iii) What are some of the items that BLOB storage is used for? Images, video, audio, logs</p> <p>iv) What is locally redundant storage and geo redundant storage? replicates your storage account three times within a single data center in the primary region replicates your data to another physical location in the secondary region to protect against regional outages</p>	3M	CO2	L2

CI Signature

CCI Signature

HoD Signature