

Internal Assessment Test 3 – June 2023

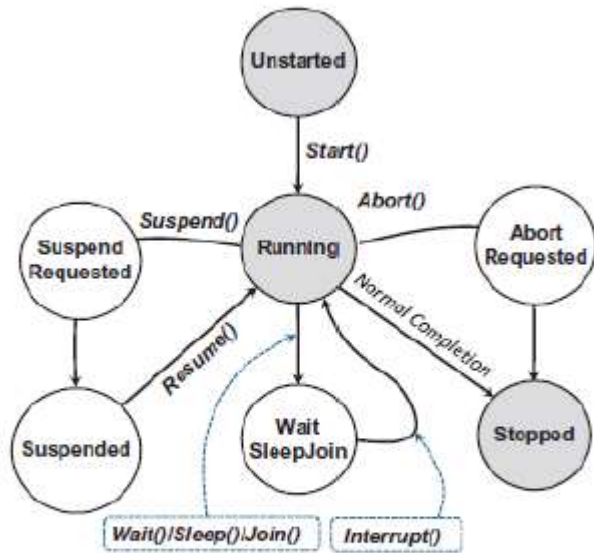
Sub:	Cloud Computing and Applications				Sub Code:	18CS643	Branch :	CSE
Date:	4.06.2023	Duration:	90 mins	Max Marks:	50	Sem / Sec:	6 A,B,C	OBE

Answer any FIVE FULL Questions

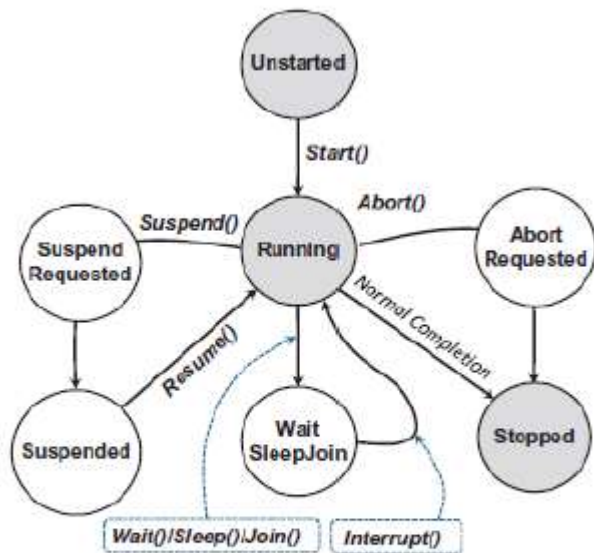
MARKS CO RB T

1 (a)	<p>How does Aneka threads differ from local thread? Explain with diagram.</p> <ul style="list-style-type: none"> <li>▶ Local threads belong to the <b>hosting process</b>.</li> <li>▶ To create a thread, only necessary to provide a <b>pointer</b> to a method</li> <li>▶ Aneka threads live in the <b>context</b> of a <b>distributed application</b></li> <li>▶ <b>multiple distributed applications</b> can be managed within a single process</li> <li>▶ <b>thread creation</b> also requires the specification of the reference to the application to which the thread belongs</li> </ul> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="background-color: #d3d3d3;">.Net Threading API</th> <th style="background-color: #d3d3d3;">Aneka Threading API</th> </tr> </thead> <tbody> <tr><td><i>System.Threading</i></td><td><i>Aneka.Threading</i></td></tr> <tr><td><i>Thread</i></td><td><i>AnekaThread</i></td></tr> <tr><td><i>Thread.ManagedThreadId (int)</i></td><td><i>AnekaThread.Id (string)</i></td></tr> <tr><td><i>Thread.Name</i></td><td><i>AnekaThread.Name</i></td></tr> <tr><td><i>Thread.ThreadState (ThreadState)</i></td><td><i>AnekaThread.State</i></td></tr> <tr><td><i>Thread.IsAlive</i></td><td><i>AnekaThread.IsAlive</i></td></tr> <tr><td><i>Thread.IsRunning</i></td><td><i>AnekaThread.IsRunning</i></td></tr> <tr><td><i>Thread.IsBackground</i></td><td><i>AnekaThread.IsBackground[false]</i></td></tr> <tr><td><i>Thread.Priority</i></td><td><i>AnekaThread.Priority[ThreadPriority.Normal]</i></td></tr> <tr><td><i>Thread.IsThreadPoolThread</i></td><td><i>AnekaThread.IsThreadPoolThread [false]</i></td></tr> <tr><td><i>Thread.Start</i></td><td><i>AnekaThread.Start</i></td></tr> <tr><td><i>Thread.Abort</i></td><td><i>AnekaThread.Abort</i></td></tr> <tr><td><i>Thread.Sleep</i></td><td>[Not provided]</td></tr> <tr><td><i>Thread.Interrupt</i></td><td>[Not provided]</td></tr> <tr><td><i>Thread.Suspend</i></td><td>[Not provided]</td></tr> <tr><td><i>Thread.Resume</i></td><td>[Not provided]</td></tr> <tr><td><i>Thread.Join</i></td><td><i>AnekaThread.Join</i></td></tr> </tbody> </table> <p>Aneka threads live and execute in a <b>distributed environment</b>, so their lifecycle is different from life cycle of threads.</p> <ul style="list-style-type: none"> <li>▶ Not possible to <b>map state value</b> of local threads to Aneka threads.</li> <li>▶ <b>Local threads</b> : most of the state transitions are triggered by the developer</li> </ul>	.Net Threading API	Aneka Threading API	<i>System.Threading</i>	<i>Aneka.Threading</i>	<i>Thread</i>	<i>AnekaThread</i>	<i>Thread.ManagedThreadId (int)</i>	<i>AnekaThread.Id (string)</i>	<i>Thread.Name</i>	<i>AnekaThread.Name</i>	<i>Thread.ThreadState (ThreadState)</i>	<i>AnekaThread.State</i>	<i>Thread.IsAlive</i>	<i>AnekaThread.IsAlive</i>	<i>Thread.IsRunning</i>	<i>AnekaThread.IsRunning</i>	<i>Thread.IsBackground</i>	<i>AnekaThread.IsBackground[false]</i>	<i>Thread.Priority</i>	<i>AnekaThread.Priority[ThreadPriority.Normal]</i>	<i>Thread.IsThreadPoolThread</i>	<i>AnekaThread.IsThreadPoolThread [false]</i>	<i>Thread.Start</i>	<i>AnekaThread.Start</i>	<i>Thread.Abort</i>	<i>AnekaThread.Abort</i>	<i>Thread.Sleep</i>	[Not provided]	<i>Thread.Interrupt</i>	[Not provided]	<i>Thread.Suspend</i>	[Not provided]	<i>Thread.Resume</i>	[Not provided]	<i>Thread.Join</i>	<i>AnekaThread.Join</i>	9 M	CO2	L2
.Net Threading API	Aneka Threading API																																							
<i>System.Threading</i>	<i>Aneka.Threading</i>																																							
<i>Thread</i>	<i>AnekaThread</i>																																							
<i>Thread.ManagedThreadId (int)</i>	<i>AnekaThread.Id (string)</i>																																							
<i>Thread.Name</i>	<i>AnekaThread.Name</i>																																							
<i>Thread.ThreadState (ThreadState)</i>	<i>AnekaThread.State</i>																																							
<i>Thread.IsAlive</i>	<i>AnekaThread.IsAlive</i>																																							
<i>Thread.IsRunning</i>	<i>AnekaThread.IsRunning</i>																																							
<i>Thread.IsBackground</i>	<i>AnekaThread.IsBackground[false]</i>																																							
<i>Thread.Priority</i>	<i>AnekaThread.Priority[ThreadPriority.Normal]</i>																																							
<i>Thread.IsThreadPoolThread</i>	<i>AnekaThread.IsThreadPoolThread [false]</i>																																							
<i>Thread.Start</i>	<i>AnekaThread.Start</i>																																							
<i>Thread.Abort</i>	<i>AnekaThread.Abort</i>																																							
<i>Thread.Sleep</i>	[Not provided]																																							
<i>Thread.Interrupt</i>	[Not provided]																																							
<i>Thread.Suspend</i>	[Not provided]																																							
<i>Thread.Resume</i>	[Not provided]																																							
<i>Thread.Join</i>	<i>AnekaThread.Join</i>																																							

- ▶ In Aneka : triggered by the **Aneka Middleware**
- ▶ **Aneka threads** : has more states
- ▶ they support **file staging**, scheduled by the middleware, which can queue them for a period of time
- ▶ They support **reservation of nodes** : state indicating execution failure due to missing reservation credentials.



System.Threading.Threadlife  
Life cycle



System.Threading.Threadlife  
Life cycle

## *Unstarted*

- ▶ On invoking Start() method moves to *Started* state
- ▶ Moves to *StagingIn* state (if there are files to upload for execution) or *Queued* state
- ▶ If there is an error, goes to *Failed* state.
- ▶ *Rejected* State is reached for invalid reservation token
- ▶ From the *Queued* state, if there is a free node, it moves to the *Running* state
- ▶ If thread generates an exception, moves to *Failed* state
- ▶ On successful completion, moves to the *Completed* state
- ▶ If output files are yet to be retrieved, it moves to *StagingOut* state
- ▶ If developer directly calls Abort() method, goes into the *Abort* state.

Unstarted-[Started]-[Queued]-Running-Completed/Aborted/Failed

## Thread Synchronization

.NET base class: provides **monitors, semaphores, reader-writer locks, and basic synchronization constructs** at language level

- ▶ **Aneka** provides **minimal support for thread synchronization**: limited to **join** operation.
- ▶ .NET provides more **stringent access** to shared data.
- ▶ This is not as strict in distributed environments, where there is **no shared memory** among threads.
- ▶ Providing locking facility in distributed environment leads to **distributed deadlocks** that are hard to detect.

## Thread Priorities

System.Threading.Thread

- ▶ ThreadPriority enumeration: **Highest, AboveNormal, Normal, BelowNormal, or Lowest.**
- ▶ Aneka **does not support thread priorities**
- ▶ But for interface compatibility purposes, it has
- ▶ Aneka.Threading.Thread
- ▶ Priority property - **ThreadPriority** is always set to **normal**

## Type Serialization

	<p>Aneka threads execute in a distributed environment in which the <b>object code in the form of libraries</b> and live instances information are moved over the network</p> <p>► <b>Local threads</b> execute all within the same address space and share memory – they do not need objects to be copied or transferred into a different address space.</p> <p>► In Aneka threads : <b>state of enclosing instance</b> needs to be transferred and reconstructed in the remote machine : at class level this is called <i>type serialization</i></p>			
<p>(b)</p>	<p>In the following code snippet, what is the use of [Serializable]?</p> <pre>[Serializable] public class Cosine(double x) {     double x;     public double Result;      public Cosine(double x) {         this.x = x;     }     public void CosineCompute() {         result = System.math.cos(x * System.Math.PI/180);     } }</pre> <p>means it is possible to convert an instance of the type into a binary array.</p>	<p>1M</p>	<p>CO2</p>	<p>L2</p>
<p>2 (a)</p>	<p>Explain Google File System with a neat diagram</p> <p>GoogleFileSystem(GFS). GFS [54] is the storage infrastructure that supports the execution of distributed applications in Google’s computing cloud. The system has been designed to be a fault- tolerant, highly available, distributed file system built on commodity hardware and standard Linux operating systems. Rather than a generic implementation of a distributed file system, GFS specifically addresses Google’s needs in terms of distributed storage for applications, and it has been designed with the following assumptions:</p> <ul style="list-style-type: none"> <li>• The system is built on top of commodity hardware that often fails.</li> <li>• The system stores a modest number of large files; multi-GB files are common and should be treated efficiently, and small files must be supported, but there is no need to optimize for that.</li> <li>• The workloads primarily consist of two kinds of reads: large streaming reads and small random reads.</li> <li>• The workloads also have many large, sequential writes that append data to files.</li> <li>• High-sustained bandwidth is more important than low latency.</li> </ul> <p>The architecture of the file system is organized into a single master, which contains the metadata of the entire file system, and a collection of chunk servers, which provide storage space.</p> <p>From a logical</p>	<p>5M</p>	<p>CO2</p>	<p>L2</p>

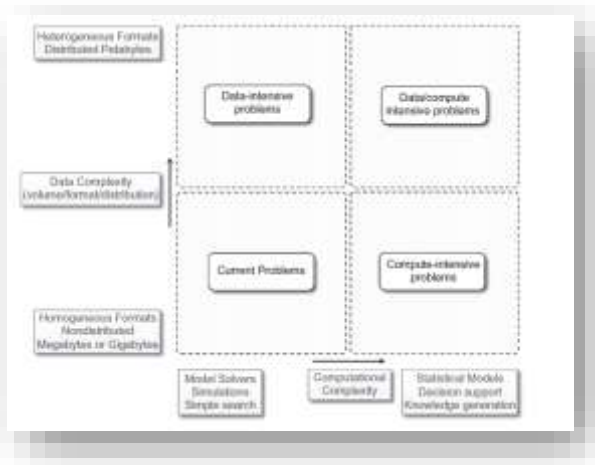
	<p>Point of view the system is composed of a collection of software daemons, which implement either the master server or the chunk server.</p> <p>A file is a collection of chunks for which the size can be configured at file system level. Chunks are replicated on multiple nodes in order to tolerate failures.</p> <p>Clients lookup the master server and identify the specific chunk of a file they want to access.</p> <p>Once the chunk is identified, the interaction happens between the client and the chunk server. Applications interact through the filesystem with a specific interface supporting the usual operations for file creation, deletion, read, and write. The interface also supports snapshots and record append operations that are frequently performed by applications. GFS has been conceived by considering that failures in a large distributed infrastructure are common rather than a rarity; therefore, specific attention has been given to implementing a highly available, lightweight, and fault-tolerant infrastructure. The potential single point of failure of the single-master architecture has been addressed by giving the possibility of replicating the master node on any other node belonging to the infrastructure. Moreover, a stateless daemon and extensive logging capabilities facilitate the system's recovery from failures.</p>			
(b)	<p>Explain Amazon S3 [OR] Apache Cassandra.</p> <ul style="list-style-type: none"> <li>▶ is a distributed object store for managing large amounts of <b>structured data</b> spread across many <b>commodity servers</b></li> <li>▶ The system is designed to <b>avoid a single point of failure</b></li> <li>▶ offer a highly <b>reliable service</b></li> <li>▶ initially developed by <b>Facebook</b></li> <li>▶ It is now part of the Apache incubator initiative</li> <li>▶ Used by <b>Facebook, Digg, and Twitter</b></li> <li>▶ 2<sup>nd</sup> gen distributed database that builds on concept of <ul style="list-style-type: none"> <li>▶ <b>AmazonDynamo</b> : fully distributed</li> <li>▶ <b>Google BigTable</b> : inherits the column family concept</li> </ul> </li> <li>▶ The <b>data model</b> exposed by Cassandra is based on the concept of a table that is implemented as a distributed <b>multidimensional map</b> indexed by a <b>key</b></li> <li>▶ value corresponding to a key : <b>highly structured object</b> and constitutes the <b>row of a table</b></li> <li>▶ Cassandra organizes the <b>row of a table into columns</b></li> <li>▶ sets of columns can be grouped <b>into column families</b></li> <li>▶ APIs : insertion, retrieval, and deletion <ul style="list-style-type: none"> <li>▶ <b>Insertion : at row level</b></li> <li>▶ <b>Retrieval and delete</b> : at column level.</li> </ul> </li> </ul> <p>Amazon S3</p> <ul style="list-style-type: none"> <li>▶ Amazon S3 is the online storage service provided by <b>Amazon</b></li> <li>▶ the system is claimed to support high availability, reliability, scalability, infinite storage, and low latency at commodity cost</li> </ul>	5M	CO2	L2

- ▶ offers a flat storage space organized into buckets, which are attached to an **Amazon Web Services (AWS)** account
- ▶ Each **bucket** can store multiple objects, each identified by a unique key
- ▶ Objects are identified by **unique URLs and exposed through HTTP**
- ▶ Allows very simple **get-put** semantics
- ▶ Because of use of **HTTP**, there is no need for any specific library for accessing the storage system
- ▶ a **POSIX-like client library** has been developed to **mount S3 buckets** as part of the local file system.
- ▶ Since the buckets are linked to AWS accounts, the owner of a bucket can decide to make it visible to other accounts or the public

3 (a) What is data intensive computing? List the challenges in data-intensive computing and explain each one in detail.

9M CO2 L2

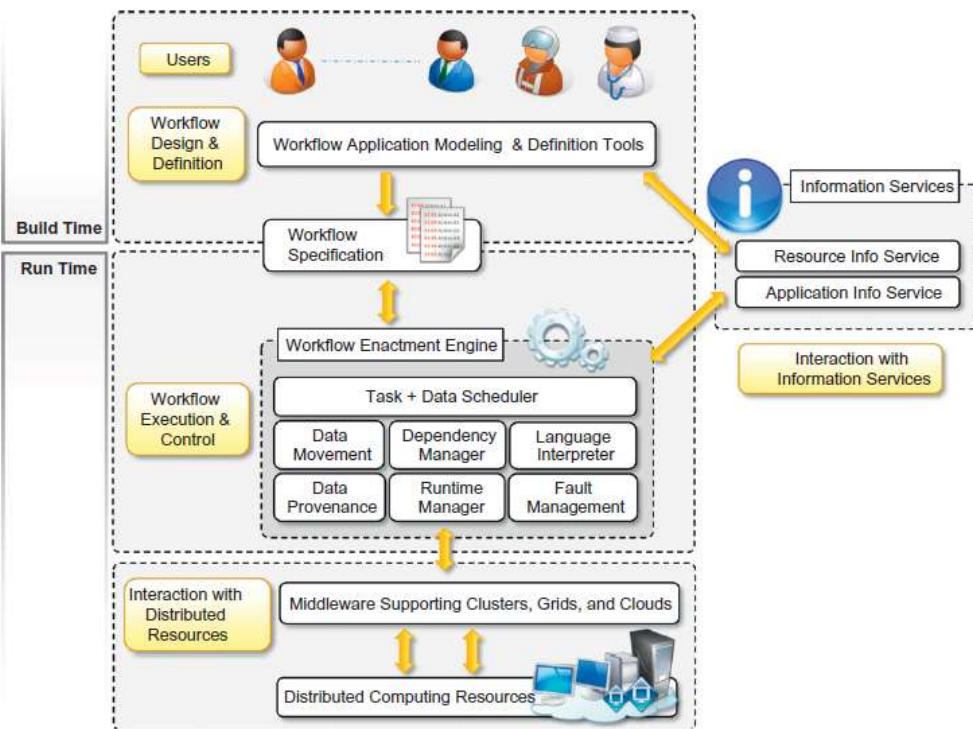
**Concerned with the production, manipulation and analysis of large-scale data in the range of hundreds of megabytes to petabytes of data.**



- ▶ Data-intensive applications not only deal with **huge volumes** of data but, very often, also exhibit **compute-intensive** properties
- ▶ Datasets are commonly persisted in **several formats** and **distributed** across different locations.
- ▶ Processing requirements **scale linearly** with the data size and they can be easily processed parallel.
- ▶ **Scalable algorithms** that can **search and process** massive datasets
- ▶ New **metadata management** technologies that can scale to handle **complex, heterogeneous, and distributed** data sources
- ▶ Advances in high-performance computing platforms aimed at providing a better support for **accessing in-memory multiterabyte data structures**
- ▶ **High-performance, highly reliable, petascale distributed file systems**
- ▶ **Data signature-generation** techniques for data reduction and rapid processing
- ▶ New approaches to **software mobility** for delivering algorithms that are able to move the computation to where the data are located

	<ul style="list-style-type: none"> <li>▶ Specialized <b>hybrid interconnection architectures</b> that provide better support for <b>filtering multigigabyte datastreams</b> coming from high-speed networks and scientific instruments</li> <li>▶ Flexible and high-performance software integration techniques.</li> </ul>			
(b)	<p>According to the Backblaze stats for Q1 2023, “The annualized failure rate (AFR) for hard drives has increased over the last three years. For example, Backblaze recorded an AFR of 0.93% in 2020 and 1.10% in 2021. The AFR for 2022 was up to 1.37%. The main reason for the increase is due to the age of the drives.”</p> <p>How is this challenge of disk failure addressed in data-intensive computing? Data replication, checksum</p>	1M	CO2	L3

4 (a)	<p>What is a workflow?</p> <ul style="list-style-type: none"> <li>▶ <i>An automation of a business process, in whole or in part, during which, documents, information, or tasks are passed from one participant ( a resource, human or machine) to another for action, according to a set of procedural rules.</i></li> <li>▶ DAG(Directed Acyclic Graph)</li> </ul> <p>▶</p> <p>▶</p>	5M	CO2	L2
(b)	<p>Explain workflow technologies for designing and extracting workflow based technologies.</p>	5M	CO2	L2



- ▶ design tools allow users to visually compose a workflow application
- ▶ normally stored in the form of an XML document
- ▶ controls the execution of the workflow by leveraging a distributed infrastructure
- ▶ Some frameworks can natively support the execution of workflow applications
  
- ▶ Kepler
- ▶ DAGMan
- ▶ Cloudbus WfMS
- ▶ Offspring

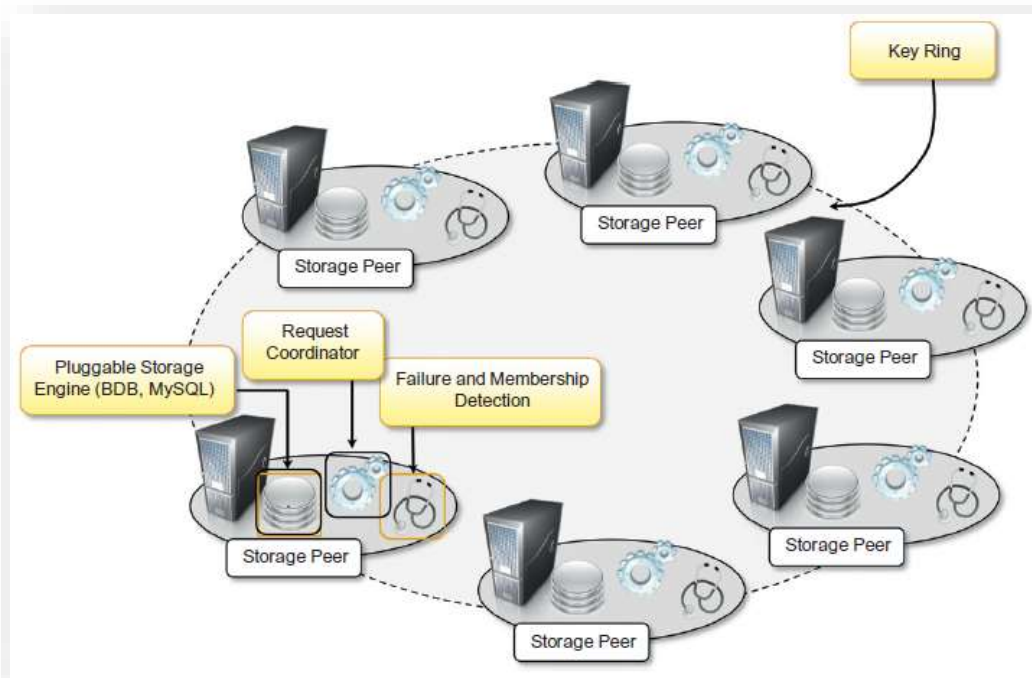
5 (a) Explain Amazon Dynamo architecture with a neat diagram.

[09]

CO2

L2





- ▶ is composed of a collection of storage peers organized in a ring that share same key space.
- ▶ **Key space** : partitioned among the storage peers
- ▶ Each peer is configured with access to a local storage facility where original objects and replicas are stored
- ▶ each node provides facilities for distributing the updates among the rings and to detect failures and unreachable nodes
- ▶ Dynamo implements the capability of being an *always-writable* store, where consistency of data is resolved in the background
- ▶ which requires applications to build their own data models on top of simple building blocks provided by the store.
- ▶ there are no **referential integrity constraints**, relationships are not embedded in the storage model, and therefore **join operations** are not supported
- ▶ is the distributed key-value store that supports the management of information of several of the business services offered by Amazon Inc.
- ▶ **Goal** : provide an incrementally scalable and highly available storage system
  - ▶ To reach reliability at a massive scale.
  - ▶ Serve 10 million requests per day.
- ▶ **Interface** : get/put semantics
- ▶ **ACID properties** are sacrificed for a more reliable and efficient infrastructure.
- ▶ Creates an **eventually consistent model** : in the long term, all users will see the same data.

(b) In ACID properties, which property is relaxed to guarantee high availability in Amazon Dynamo and generally in all distributed databases ?

[1]

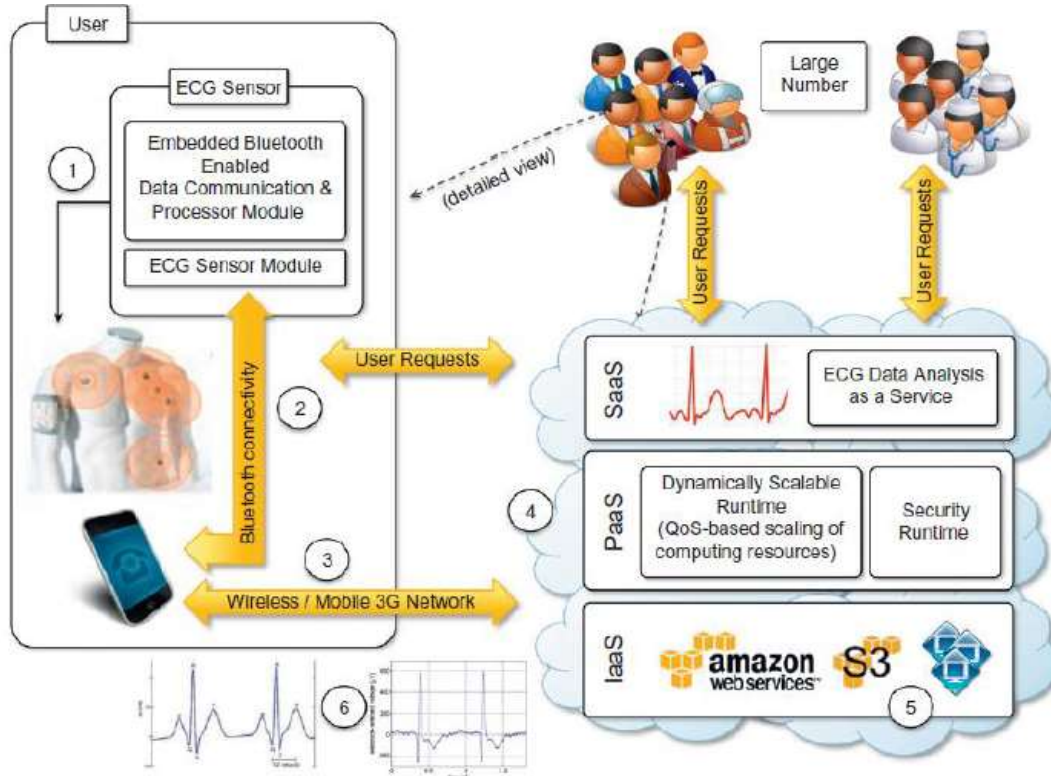
CO2

L3

6 Explain how cloud computing can support ECG monitoring with a neat diagram.

10M

CO3 L2



**Wearable computing devices** equipped with ECG sensors constantly monitor the patient's heartbeat.

- ▶ transmitted to the patient's **mobile device**, eventually forwarded to **cloud-hosted Web service** for analysis
- ▶ Web service forms the front-end that is entirely hosted in the cloud
- ▶ Leverages SaaS, PaaS and IaaS.
- ▶ Web service constitute the SaaS application : **store ECG data in the Amazon S3 service**
- ▶ issue a **processing request** to the scalable cloud platform
- ▶ runtime platform is composed of a dynamically sizable number of instances running the workflow engine and Aneka.

**Tasks : set of operations involving the extraction of the waveform from the heartbeat data**

- ▶ **comparison of the waveform with a reference waveform** to detect anomalies.
- ▶ If anomalies are found, **doctors and first-aid personnel** can be notified to act on a specific patient.

CO-PO and CO-PSO Mapping																			
Course Outcomes		Blooms Level	Modules covered	PO 1	PO 2	PO 3	PO 4	PO 5	PO 6	PO 7	PO 8	PO 9	PO 10	PO 11	PO 12	PSO 1	PSO 2	PSO 3	PSO 4
CO 1	Explain cloud computing, virtualization and classification of services of cloud computing	L2	1,2	3	2	2	2	0	3	3	0	0	0	0	0	3	0	2	2
CO 2	Illustrate architecture and programming examples in cloud	L3	2,3,4	3	2	0	2	2	3	3	0	0	0	0	0	3	2	0	2
CO 3	Describe the platforms for development of cloud applications with examples	L2	4,5	3	3	3	3	2	3	3	0	0	0	0	0	3	2	0	2

### CO PO Mapping

COGNITIVE LEVEL	REVISED BLOOMS TAXONOMY KEYWORDS
L1	List, define, tell, describe, identify, show, label, collect, examine, tabulate, quote, name, who, when, where, etc.
L2	summarize, describe, interpret, contrast, predict, associate, distinguish, estimate, differentiate, discuss, extend
L3	Apply, demonstrate, calculate, complete, illustrate, show, solve, examine, modify, relate, change, classify, experiment, discover.
L4	Analyze, separate, order, explain, connect, classify, arrange, divide, compare, select, explain, infer.
L5	Assess, decide, rank, grade, test, measure, recommend, convince, select, judge, explain, discriminate, support, conclude, compare, summarize.

PROGRAM OUTCOMES (PO), PROGRAM SPECIFIC OUTCOMES (PSO)				CORRELATION LEVELS	
PO1	Engineering knowledge	PO7	Environment and sustainability	0	No Correlation
PO2	Problem analysis	PO8	Ethics	1	Slight/Low
PO3	Design/development of solutions	PO9	Individual and team work	2	Moderate/Medium
PO4	Conduct investigations of complex problems	PO10	Communication	3	Substantial/High
PO5	Modern tool usage	PO11	Project management and finance		
PO6	The Engineer and society	PO12	Life-long learning		
PSO1	Develop applications using different stacks of web and programming technologies				

PSO2	Design and develop secure, parallel, distributed, networked, and digital systems
PSO3	Apply software engineering methods to design, develop, test and manage software systems.
PSO4	Develop intelligent applications for business and industry

---