MCA 502 - Cloud Computing

Theory 80 + 20 marks
Lab: 80 + 20 marks

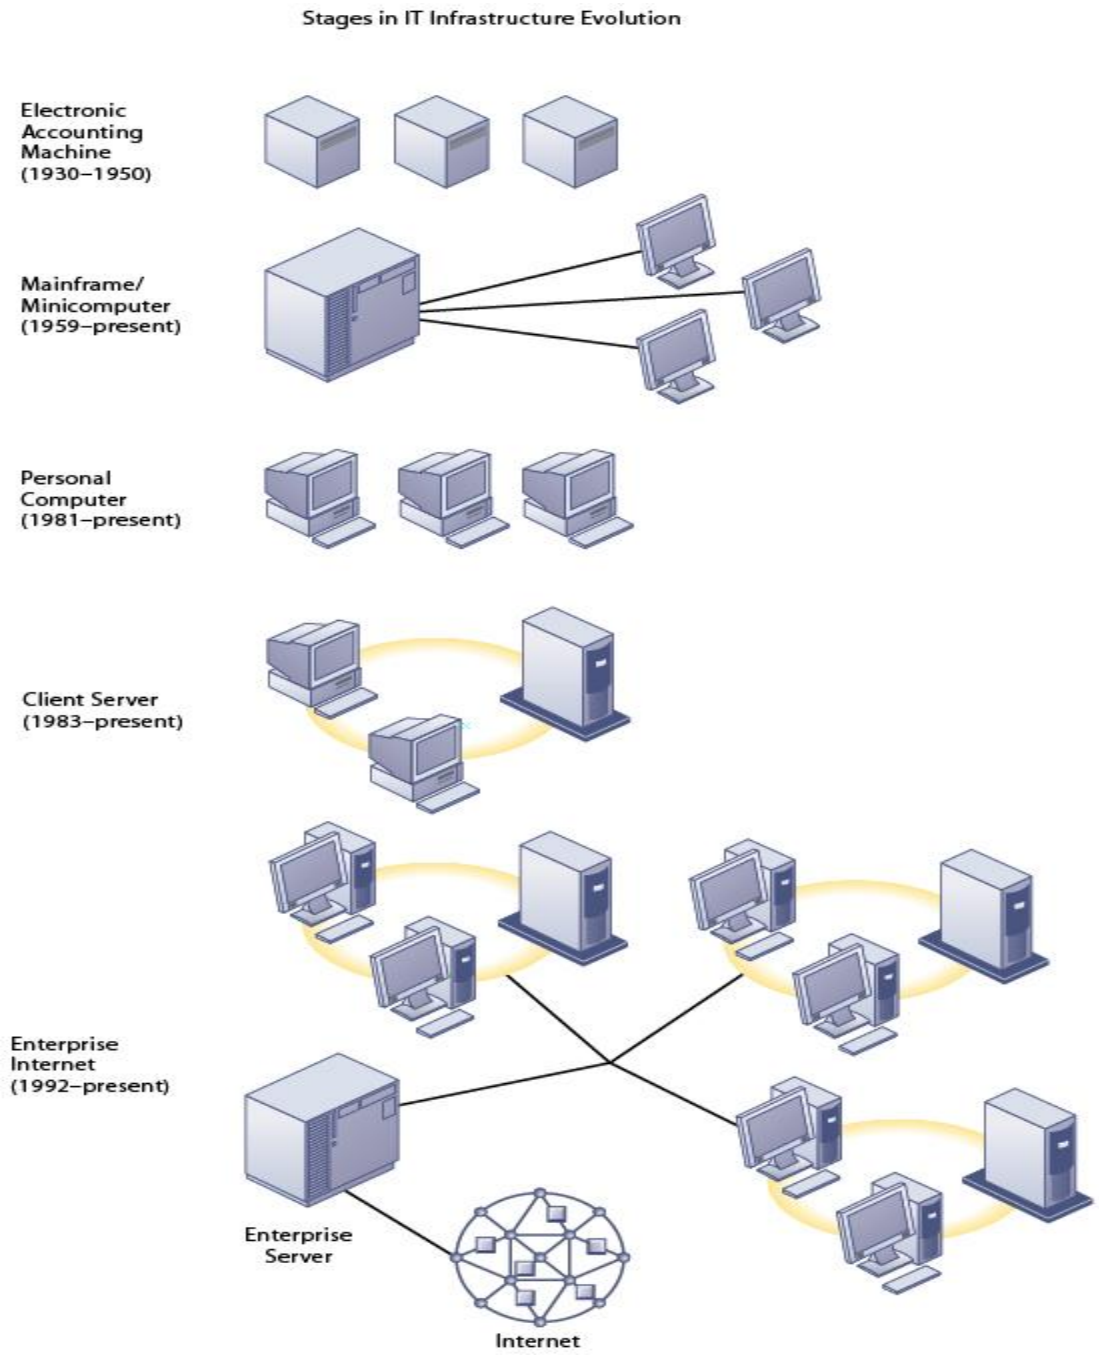
Stages in IT Infrastructure Evolution
Electronic Accounting Machine (1930–1950)
Mainframe/ Minicomputer (1959–present)
Personal Computer (1981–present)
Client Server (1983–present)
Enterprise Internet (1992–present)
Enterprise Server
Internet

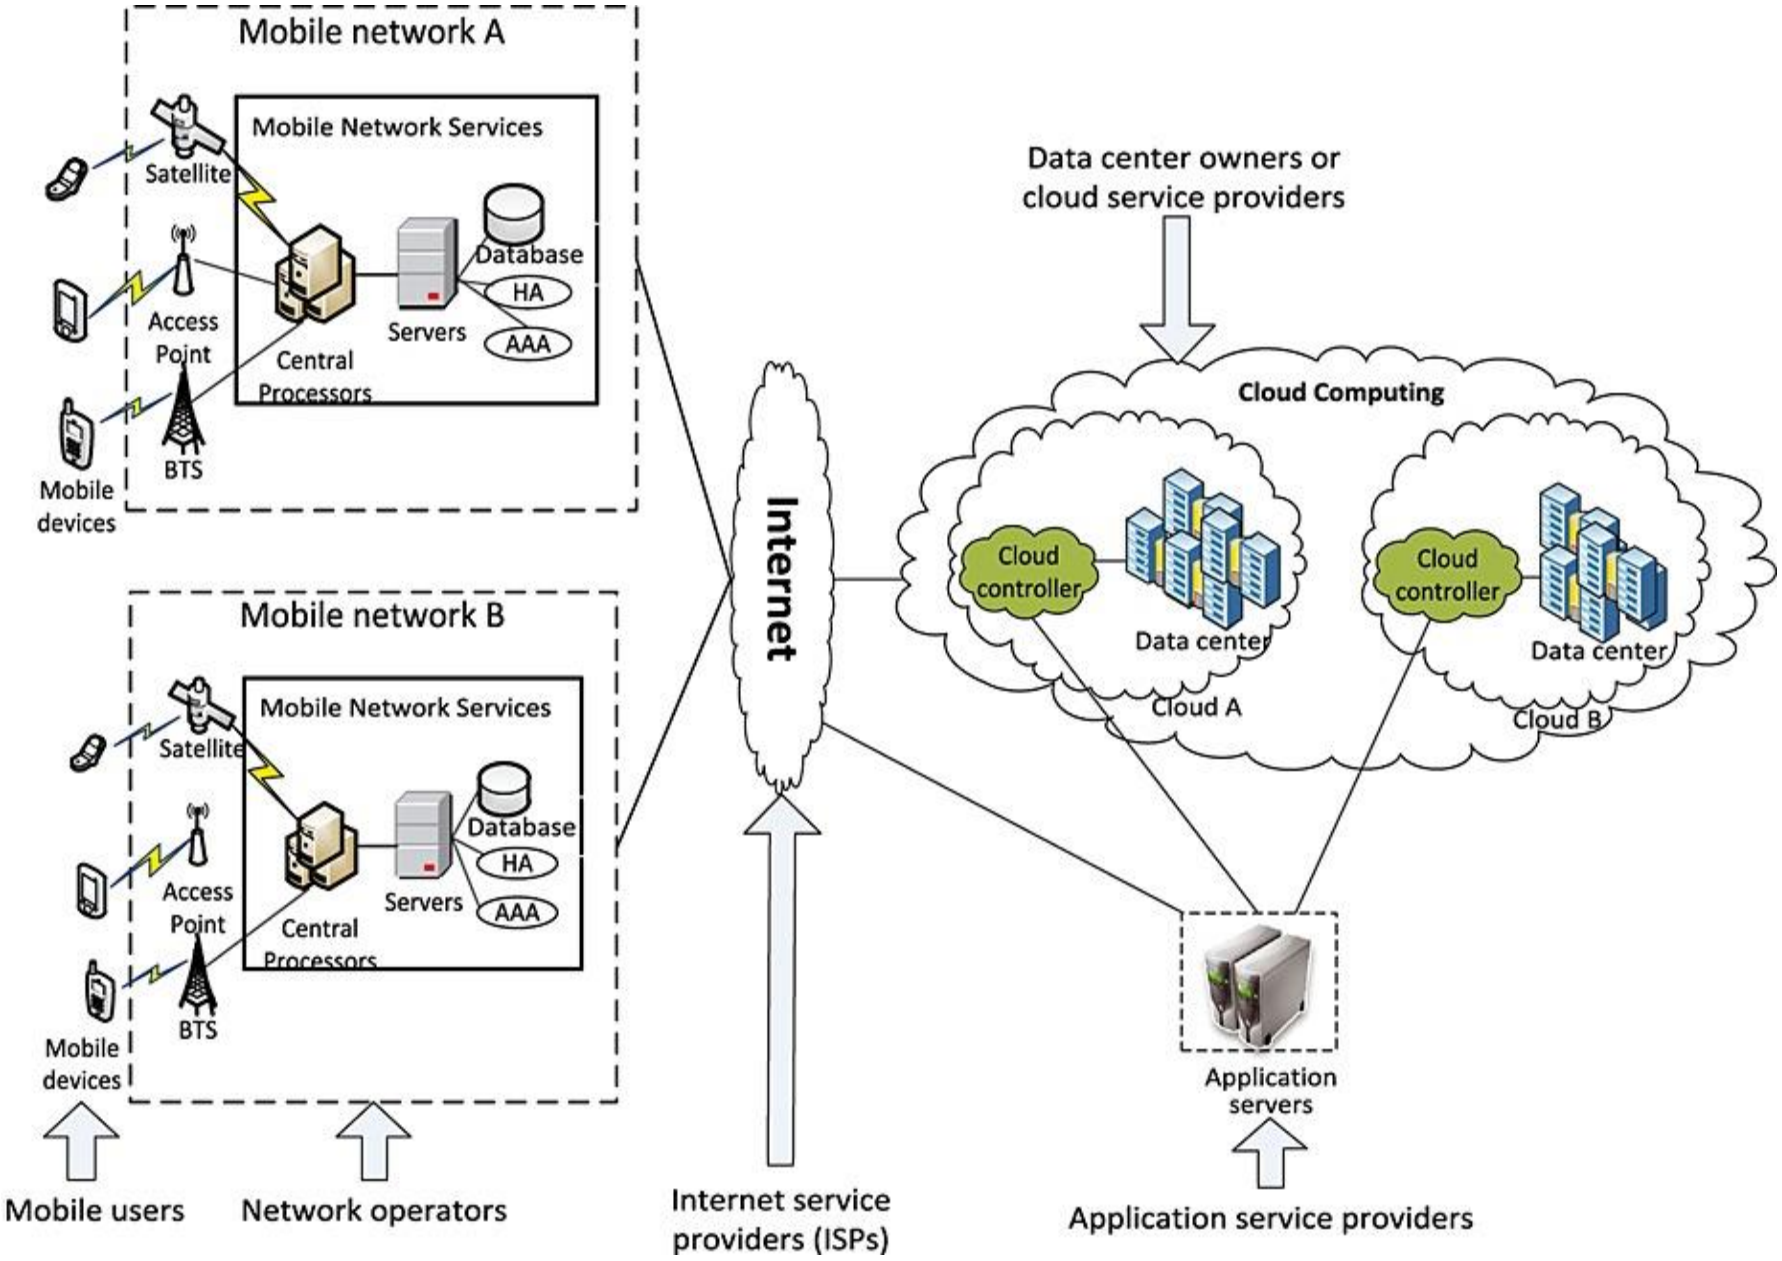
Mobile network A
Mobile Network Services
Satellite
Access Point
Central Processors
Servers
Database
HA
AAA
BTS
Mobile devices
Mobile network B
Mobile Network Services
Satellite
Access Point
Central Processors
Servers
Database
HA
AAA
BTS
Mobile devices
Internet
Data center owners or cloud service providers
Cloud Computing
Cloud controller
Data center
Cloud A
Cloud controller
Data center
Cloud B
Application servers
Mobile users
Network operators
Internet service providers (ISPs)
Application service providers

Introduction:

Computing describes the way how computers and computer systems work and how they are constructed and programmed.

1. DISTRIBUTED COMPUTING

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

- It also refers to the processing in which different parts of a program run concurrently on two or more processors that are part of the same computer.
- Both types of processing require that a program be segmented—divided into segments that can run concurrently. Over the decades, distributed computing has been an essential component of scientific computing.
- It comprises of a set of processes that cooperate to achieve a common specific goal. Mostly social network sites are implemented by using the concept of distributed computing systems
- There are two chief distributed computing standards: CORBA and DCOM

2. UTILITY COMPUTING

- Due to its flexibility and economy, utility computing is one of the most popular IT service models.
- Grid computing, cloud computing and managed IT services are based on the concept of utility computing.
- Utility computing is the process of providing computing service through an on-demand, pay-peruse billing method.
- Utility computing is a computing business model in which the provider owns, operates and manages the computing resources, infrastructure and the subscribers accesses it as and when required on a rental or metered basis .
- Utility computing usually envisions some form of virtualization so that the amount of storage or computing power available is considerably larger than that of a single time-sharing computer. For this purpose, multiple servers are used on the back end. These might be a dedicated computer cluster specifically built for the purpose of being rented out.
- **This model is based on that used by conventional utilities such as telephone services, electricity and gas**

3.CLUSTER COMPUTING

A cluster computing is a type of parallel or distributed processing system, which consists of a collection of interconnected stand-alone computers working together as a single integrated computing resource.

The components of a cluster are commonly, but not always, linked to each other through fast local area networks.
A computer node will be one or multiprocessor system (PCs, workstations, or SMPs) with memory, I/O facilities, and an operating system.

A cluster generally refers to two or additional computers (nodes) connected along. The nodes can exist in a single cabinet or be physically separated and connected via a LAN. An inter-connected (LAN-based) cluster of computers will seem as a single system to users and applications.

Multiple High Performance Computers (PCs, Workstations, or SMPs)

4.GRID COMPUTING

- Grid computing is a distributed architecture that combines computer resources from various domains to reach a main objective.
- In grid computing, the computers run independent tasks and are loosely linked by the Internet can work on a task together, thus functioning as a supercomputer.
- Typically, a grid works on various tasks among a network, but it is additionally capable of performing on specialized applications. It is designed to resolve issues that are too big for a supercomputer while maintaining the flexibility to process various smaller problems.
- Computing grids deliver a multiuser infrastructure that accommodates the discontinuous demands of large information processing.
- A grid is connected by parallel nodes that form a computer cluster, which runs on an operating system like Linux or free software
- The grid technology is applied to a wide range of applications, such as mathematical, scientific or educational tasks through several computing resources. It is often used in structural analysis, Web services such as ATM banking, backoffice infrastructures, and scientific or marketing research

HISTORY OF CLOUD COMPUTING

Cloud computing is believed to have been invented by **Joseph Carl Robnett Licklider** in the 1960s with his work on ARPANET to connect people and data from anywhere at any time.

In 1969, J. C. R. Licklider helped develop the ARPANET (Advanced Research Projects Agency Network), a “very” primitive version of the Internet. JCR, or “Lick” was both a psychologist and a computer scientist, and promoted a vision called the “*Intergalactic Computer Network*,” in which everyone on the planet would be interconnected by way of computers, and able to access information from anywhere.

Cloud computing is one of the greatest developments in technology over the past decade, and allows online file access across several applications spread over a vast geographical area.

Cloud computing creates virtual space or a cloud of files and applications that may be used and shared by all members of an organization, no matter where they are.

Cloud Computing is the practice of using large groups of remote servers, hosted on Internet, to store and access applications and computer data, instead of saving them on the local server or personal computer.

NIST definition: "Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

This cloud model is composed of five essential characteristics, three service models, and four deployment models.

The NIST Cloud Definition Framework

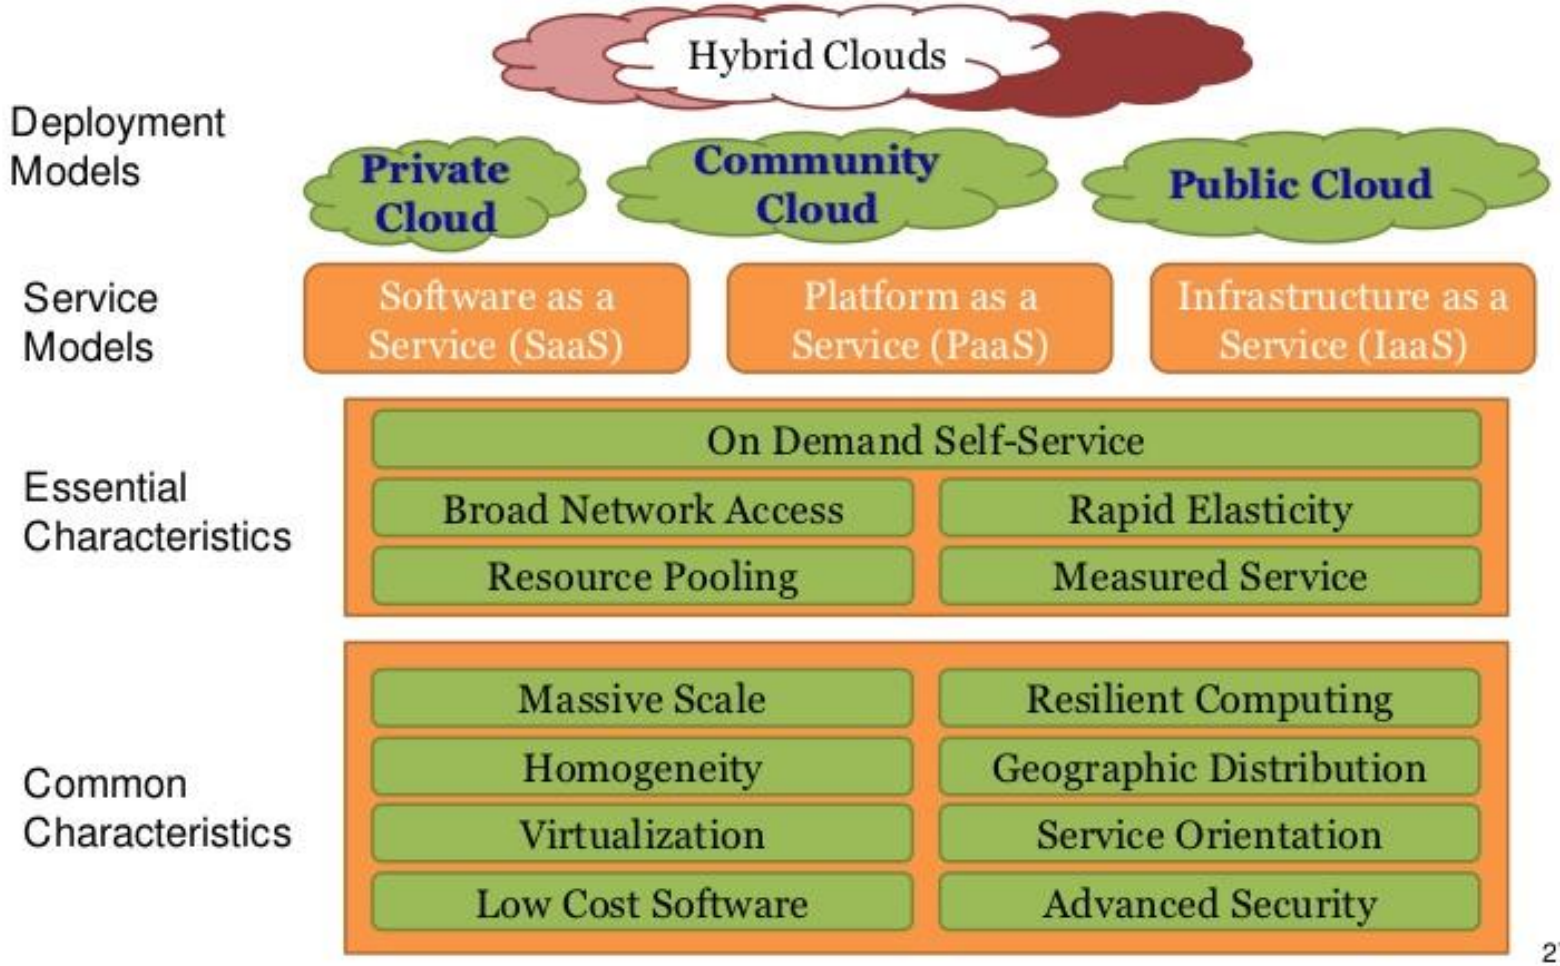

The need for Cloud Computing

Most organizations spend several million dollars each year on procuring the right hardware and software for their employees. These purchases include not just the computers and laptops, but also software and software licenses. For every additional employee a new software license has to be bought. This is financially draining for any organization, big or small.

It drastically reduces the cost of hiring software engineers and other specialized workers for IT support, as the streamlined software available on the cloud eliminates the need for a lot of heterogeneous hardware and software.

Also, an organization that avails of **cloud computing** can use the processing power of the entire network or cloud that it is a part of. So it never falls short on processing power.

Another huge advantage is that since **cloud computing** eliminates the need for an organization to buy hardware, a plenty of physical space is saved. There is no need to reserve space on the company premises for giant servers, for instance.

Essential Characteristics:

On-demand self-service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

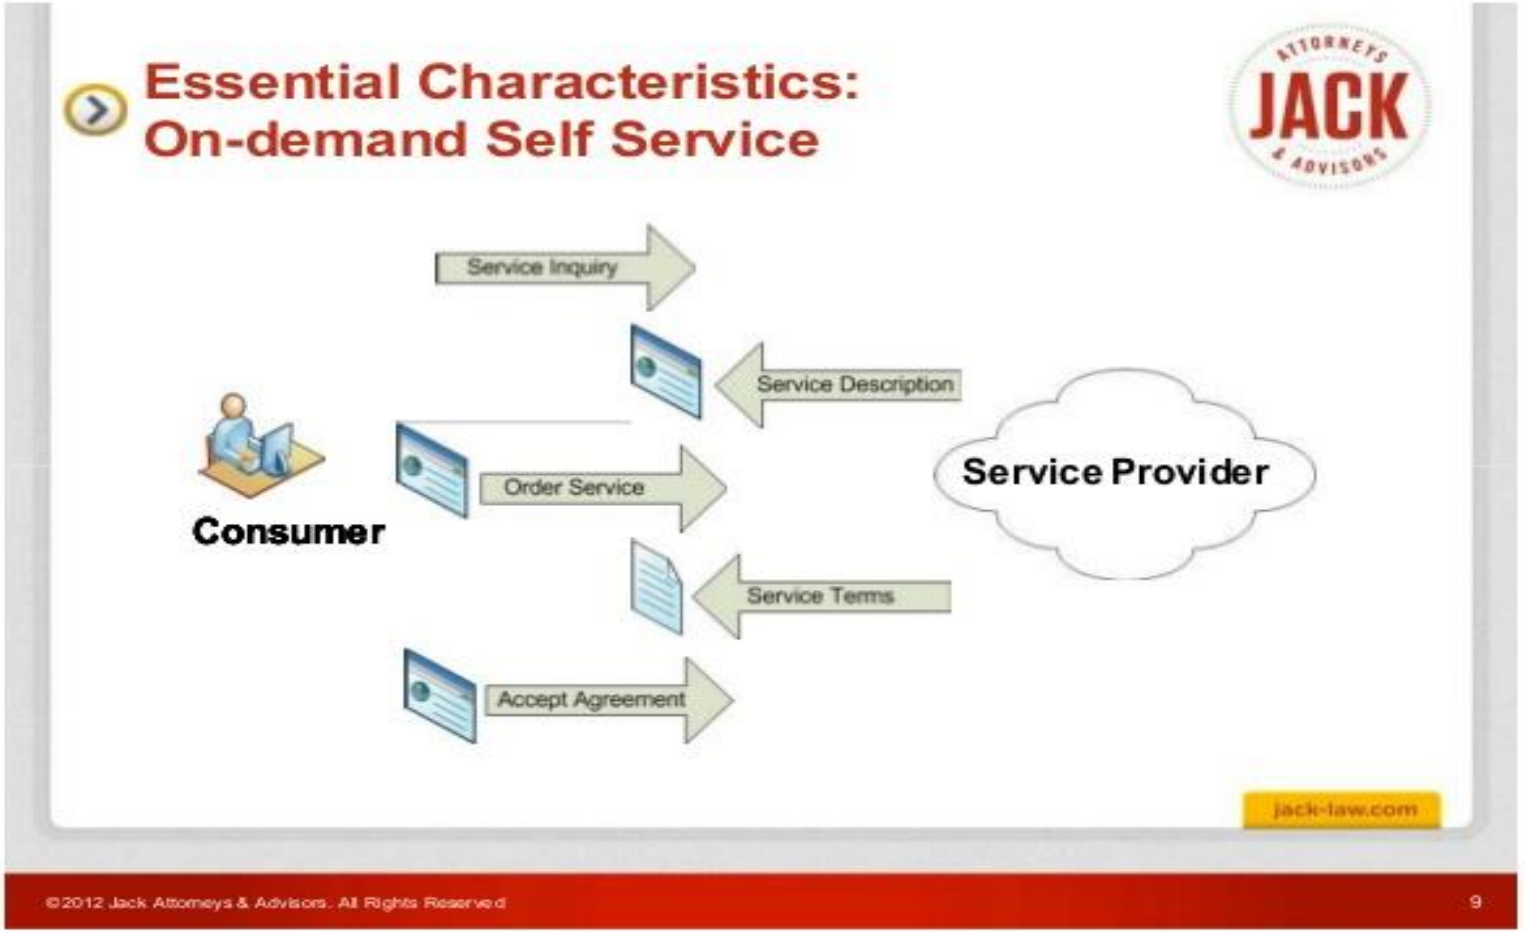

Broad network access/ Ubiquitous network access

Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).

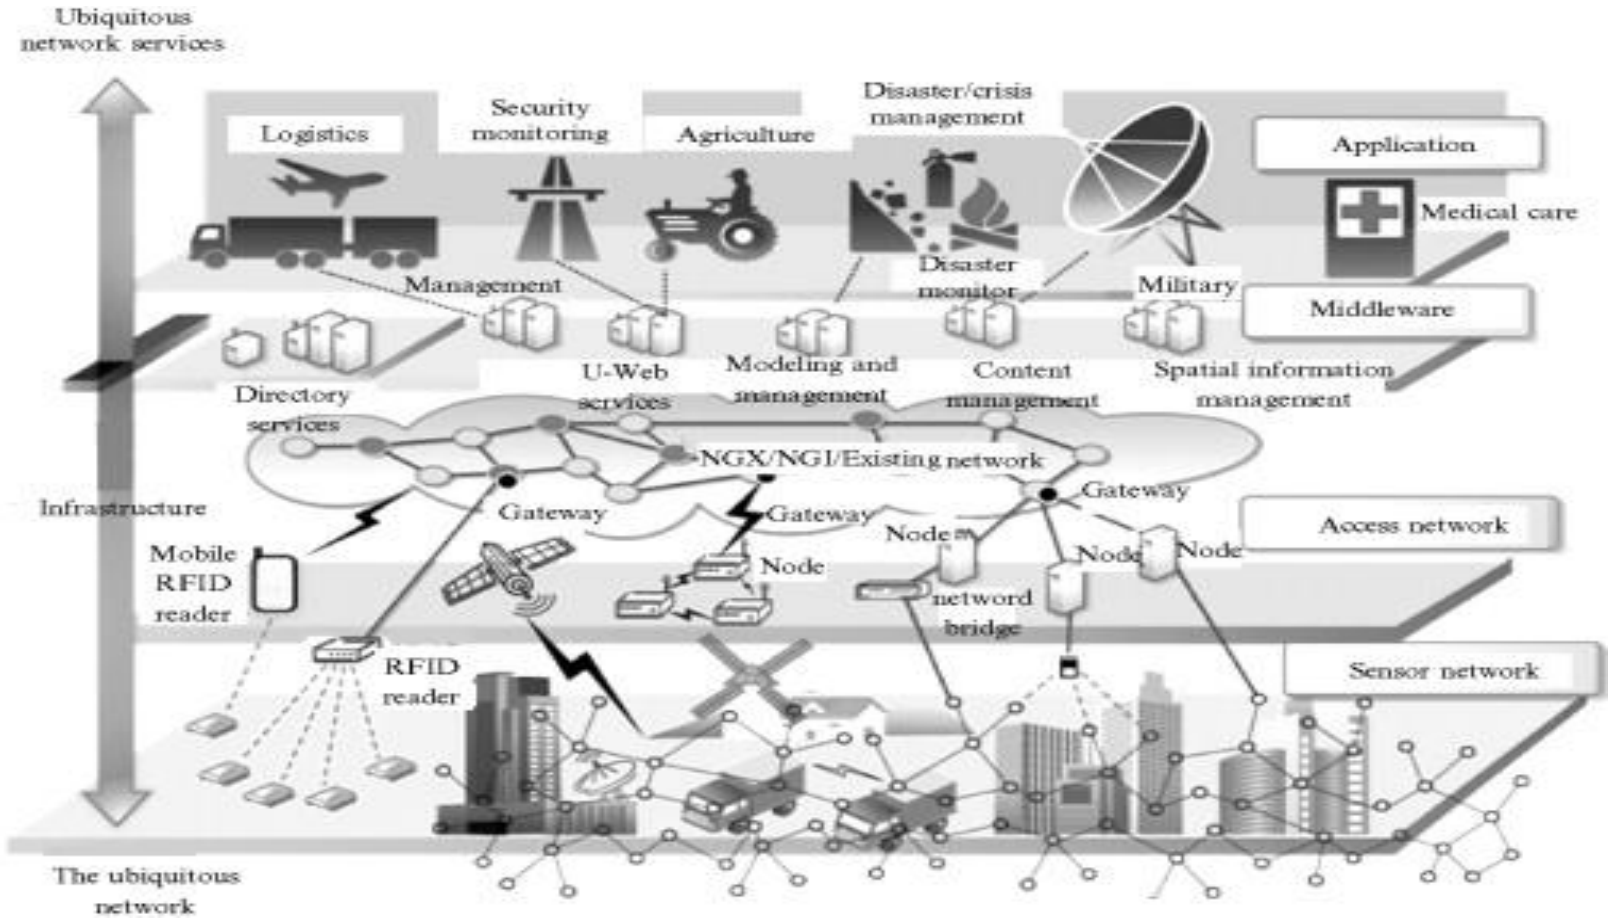

Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.

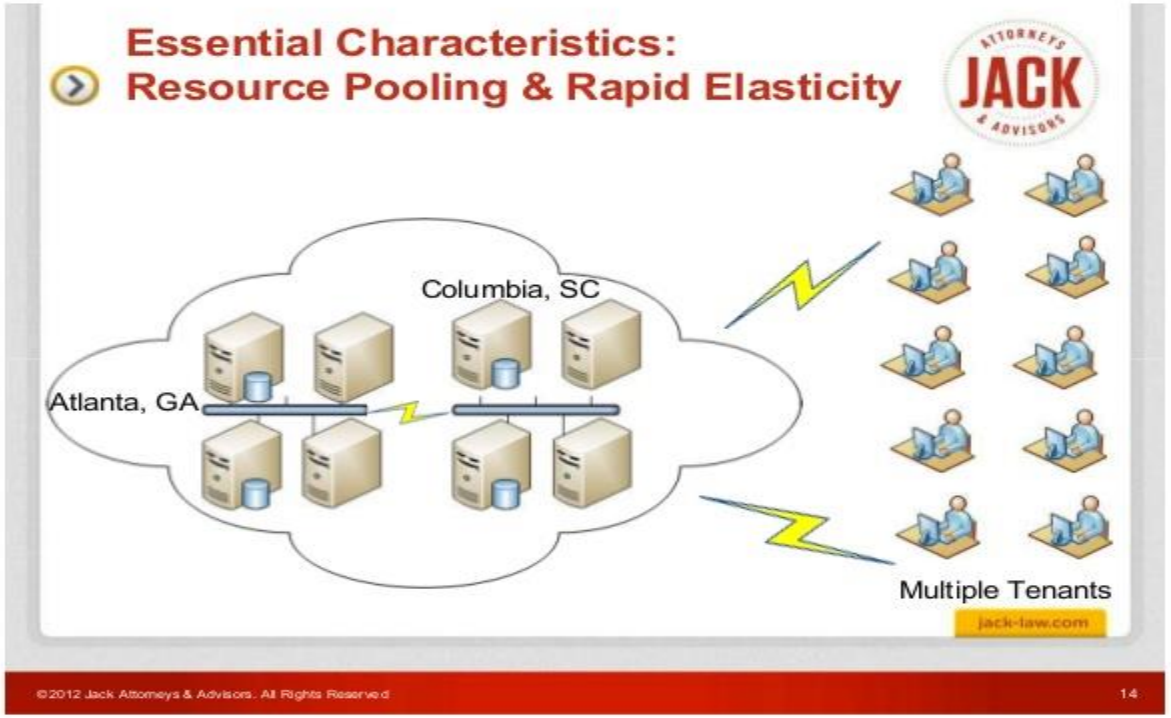

Rapid elasticity. Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

Cloud Computing Characteristics: *Rapid Elasticity*

Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out, and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

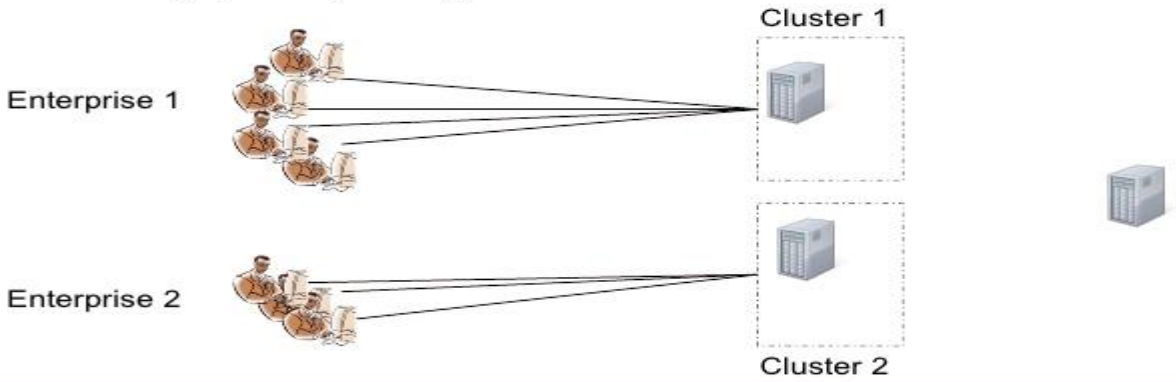

15

Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability1 at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

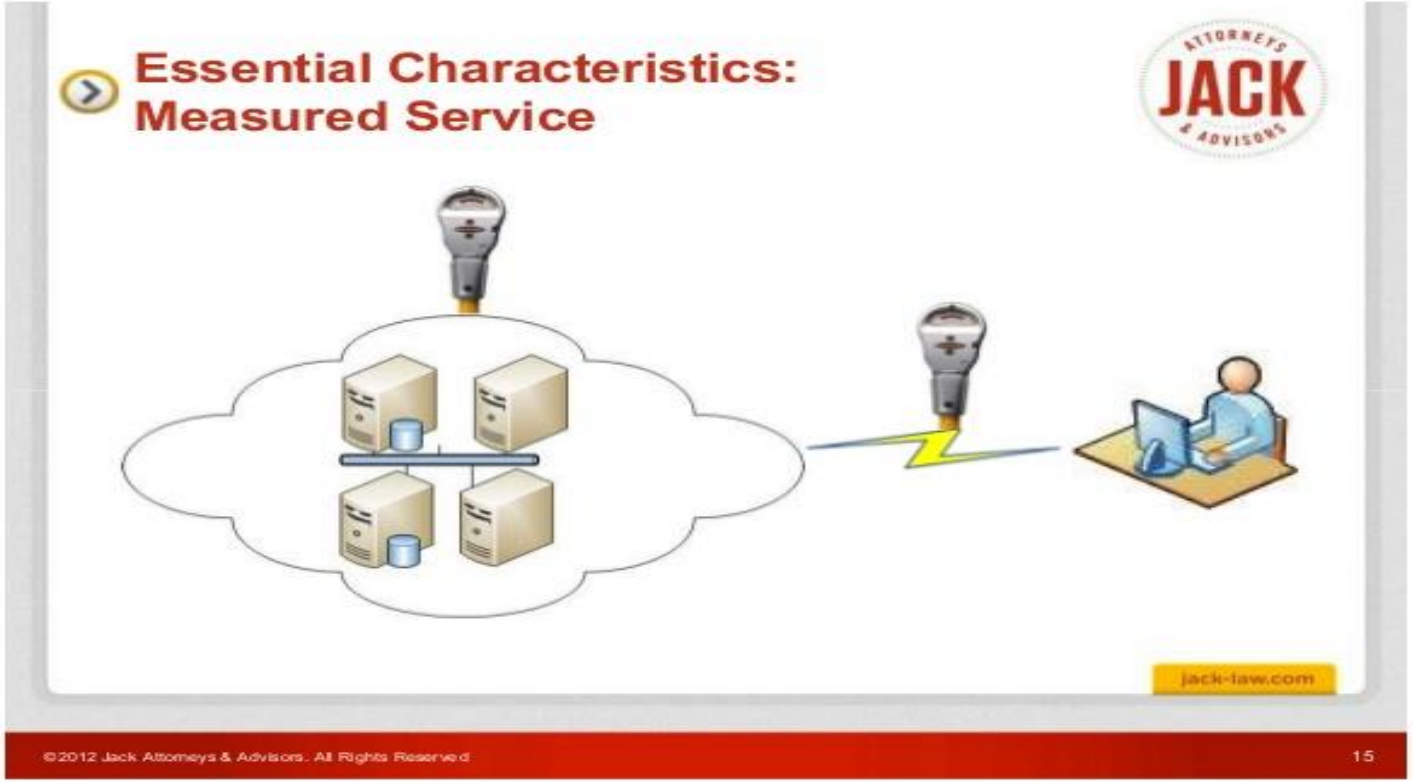

Service Models

Software as a Service (SaaS). The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure2. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Examples of SaaS are Salesforce, Google Docs, Office 365, Basecamp etc

Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Examples of PaaS are Google App Engine, Cloud Foundry, Engine Yard Etc.

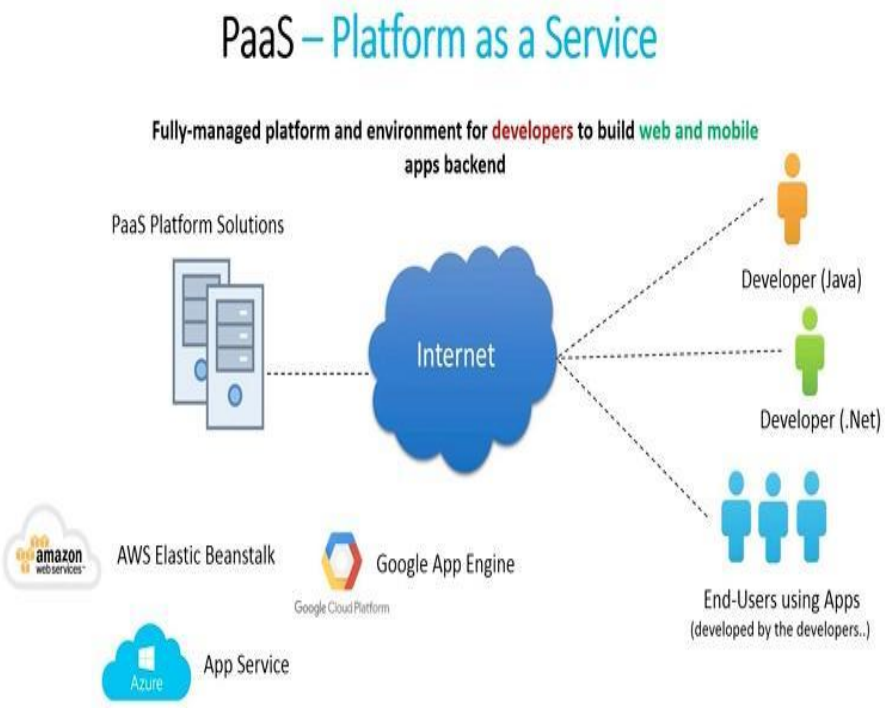

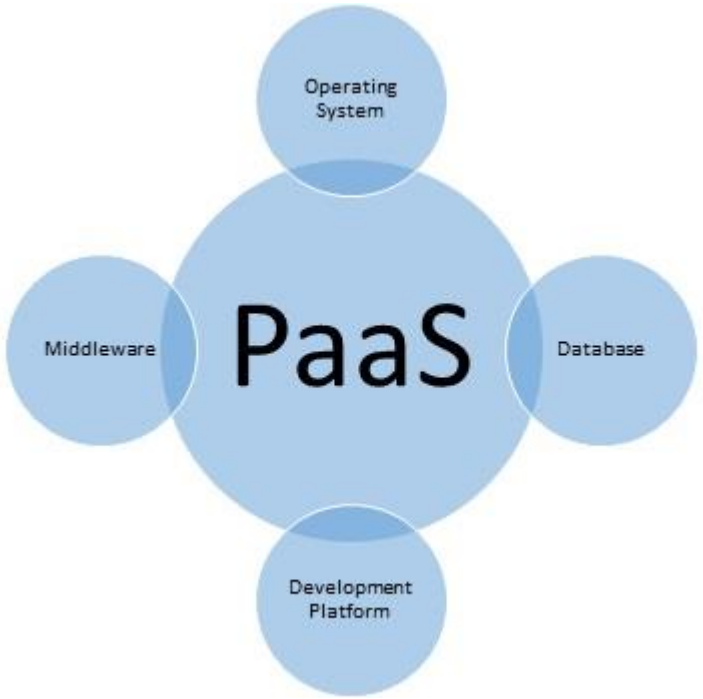

Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

Examples of IaaS are Amazon EC2, Rackspace, Google Compute Engine etc.

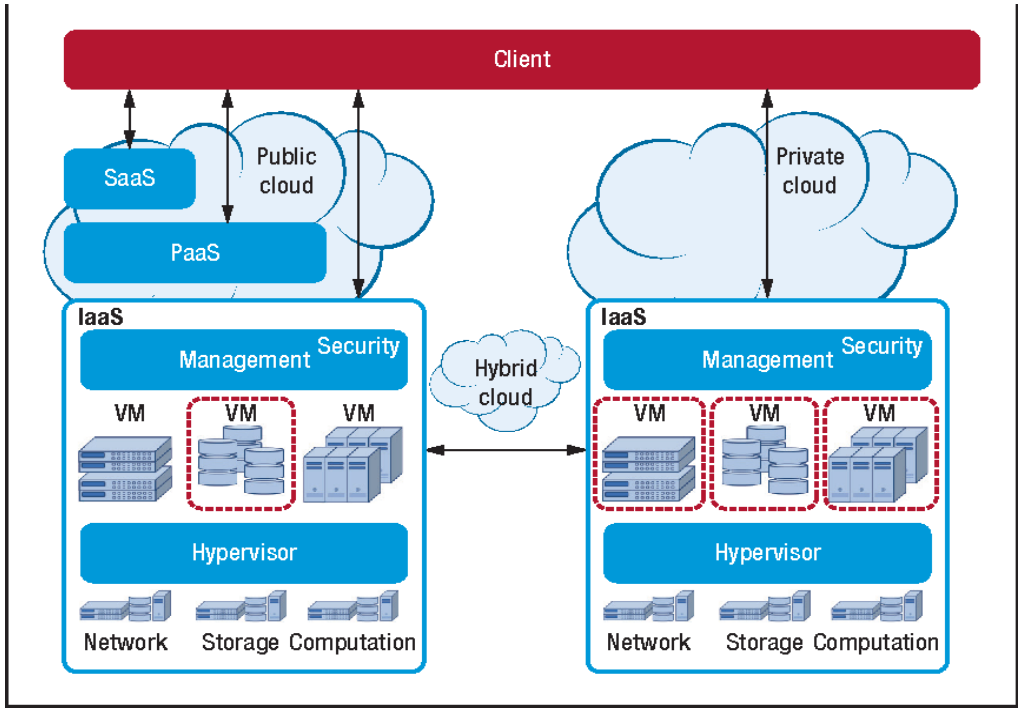

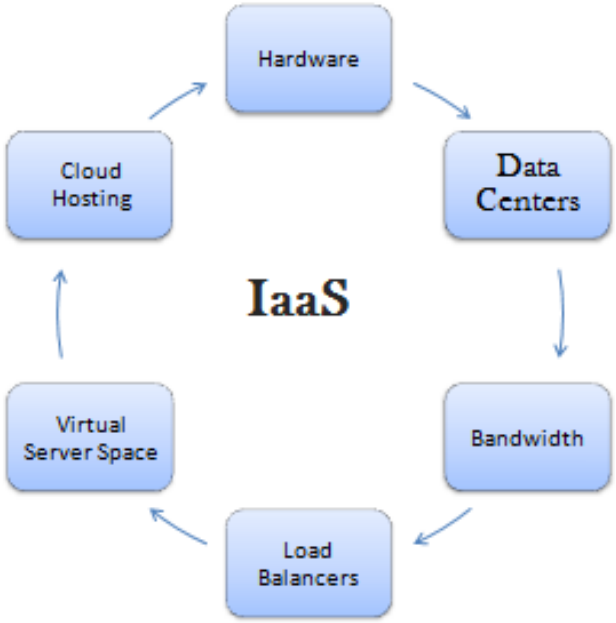

A **cloud deployment model** is a specific configuration of environment parameters such as the accessibility and proprietorship of the deployment infrastructure and storage size. It means that deployment types vary depending on who controls the infrastructure and where it resides.

- To make the most use of this computing type, a company should opt for a model that suits it the most.
- To choose one, consider computing, networking and storage requirements, available resources and business goals, as well as the pros and cons of cloud deployment models.

  - Public Cloud
  - Private Cloud
  - Community Cloud
  - Hybrid Cloud

Public Cloud

The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

- Server infrastructure belongs to service providers that manage them and administer pool resources, which is why there is no need for user companies to buy and maintain their hardware.

- Provider companies offer resources as a service both free of charge or on a pay-per-use basis via the Internet connection. Users can scale resources when required.

When it comes to popular public cloud deployment models, examples are Amazon Elastic Compute Cloud (Amazon EC2) — the top service provider, Microsoft Azure, Google App Engine, IBM Cloud, Salesforce Heroku and others.

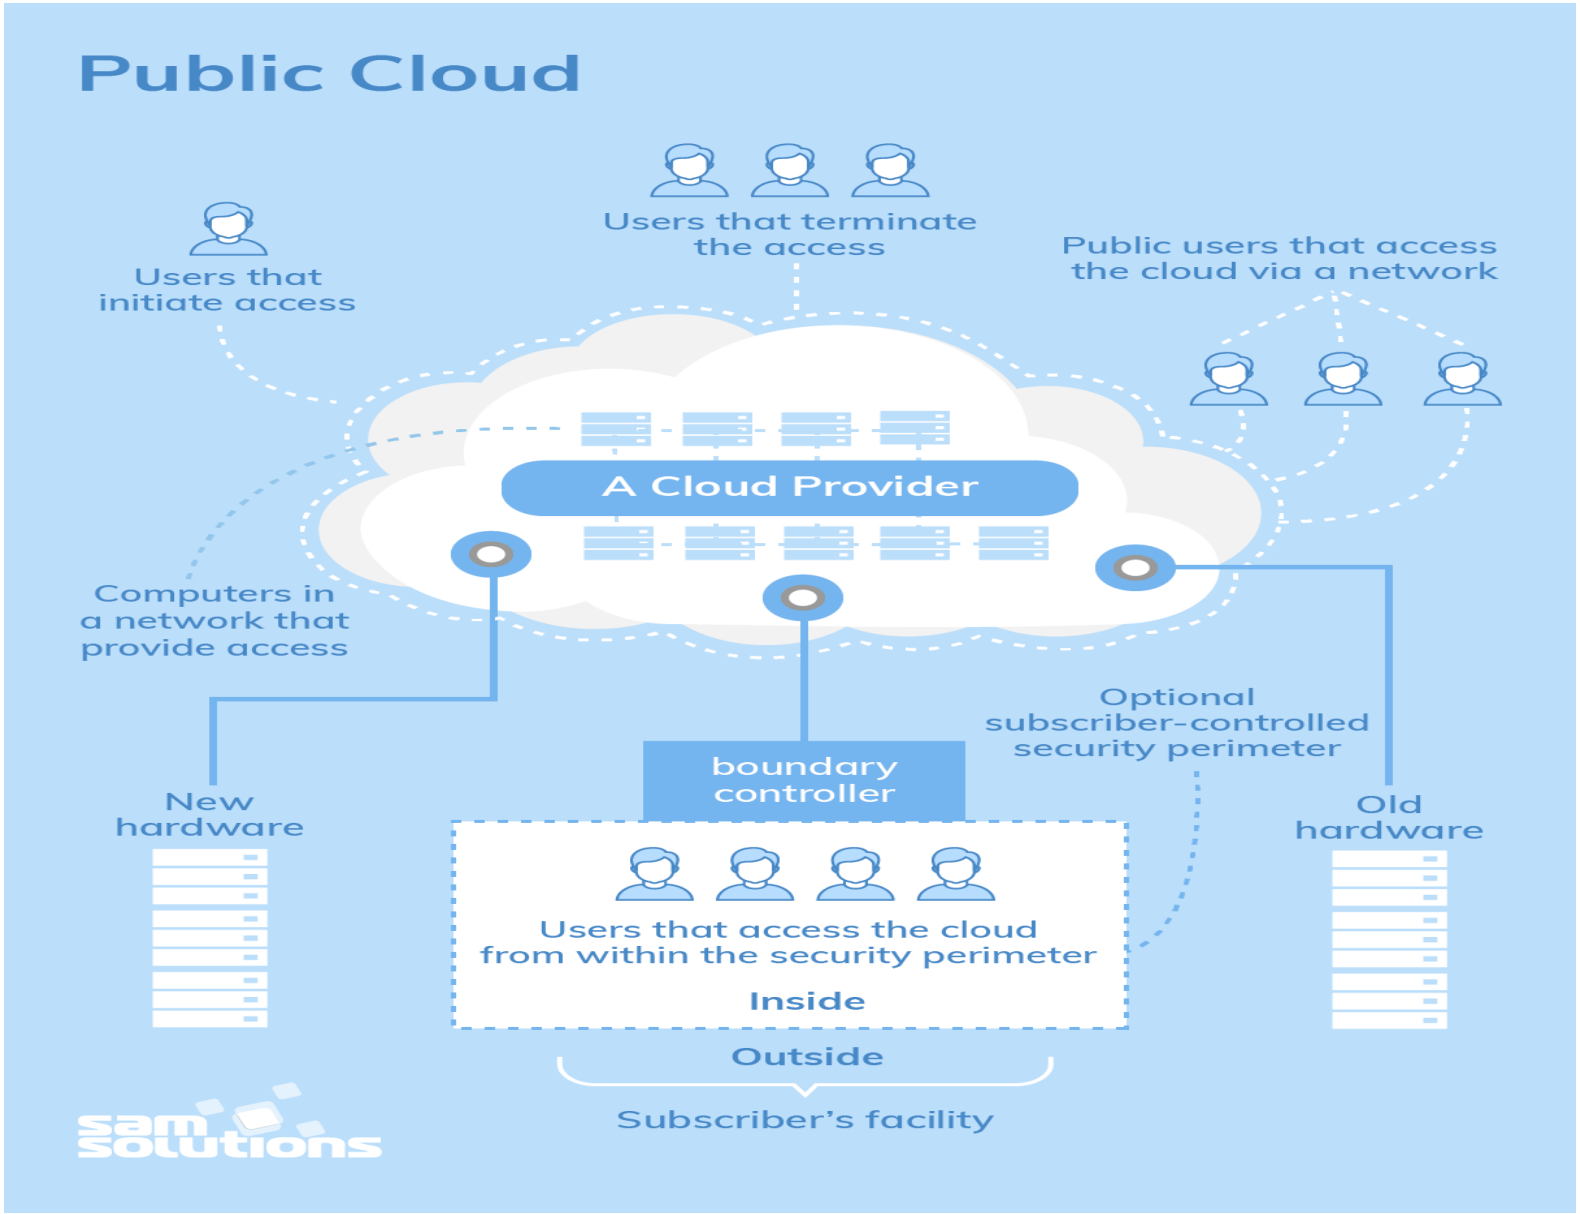
Public Cloud
Users that terminate the access
Users that initiate access
Public users that access the cloud via a network
A Cloud Provider
Computers in a network that provide access
Optional subscriber-controlled security perimeter
boundary controller
New hardware
Old hardware
Users that access the cloud from within the security perimeter
Inside
Outside
Subscriber's facility
sam solutions

- **Hassle-free infrastructure management**. Having a third party running your cloud infrastructure is convenient: you do not need to develop and maintain your software as the service provider does it for you. Also, the infrastructure setup and use are unsophisticated.
- **High scalability**. You can easily extend the available capacity as your company requirements increase.
- **Reduced costs**. You pay only for the service you use, no need to invest in hardware or software.
- **24/7 uptime**. An extensive network of the provider's servers ensures the constant availability of your infrastructure and its improved operation time

Compromised reliability.
often enough, public clouds experience outages and malfunction, as it was in the case of the Salesforce CRM disruption in 2016 that caused a storage collapse.
Data security and privacy give rise to concern. Although access to data is easy, a public deployment model deprives users of knowing where their information is kept and who has access to it.

The lack of a bespoke service. Service providers have only standardized service options, which is why they will probably fail to satisfy their requirements if they are unusual.

Private Cloud

There is little to no difference between a public and a private model from the technical point of view, as their architectures are very similar. However, opposed to a public cloud that is available to the general public, **only one specific company owns a private one. That is why it is also called an *internal* or *corporate***.

Only one organization uses this deployment model to run its workloads, and the server can be hosted externally or on the premises of the user company.

Regardless of their physical location, these infrastructures are maintained on a designated private network and use software and hardware that are intended for a specific company.

- Private model is especially suitable for companies that seek to safeguard their mission-critical operations or for businesses with changing requirements.

- Multiple public cloud service providers — including Amazon, IBM, Cisco, Dell and Red Hat — also provide private solutions along with public ones.

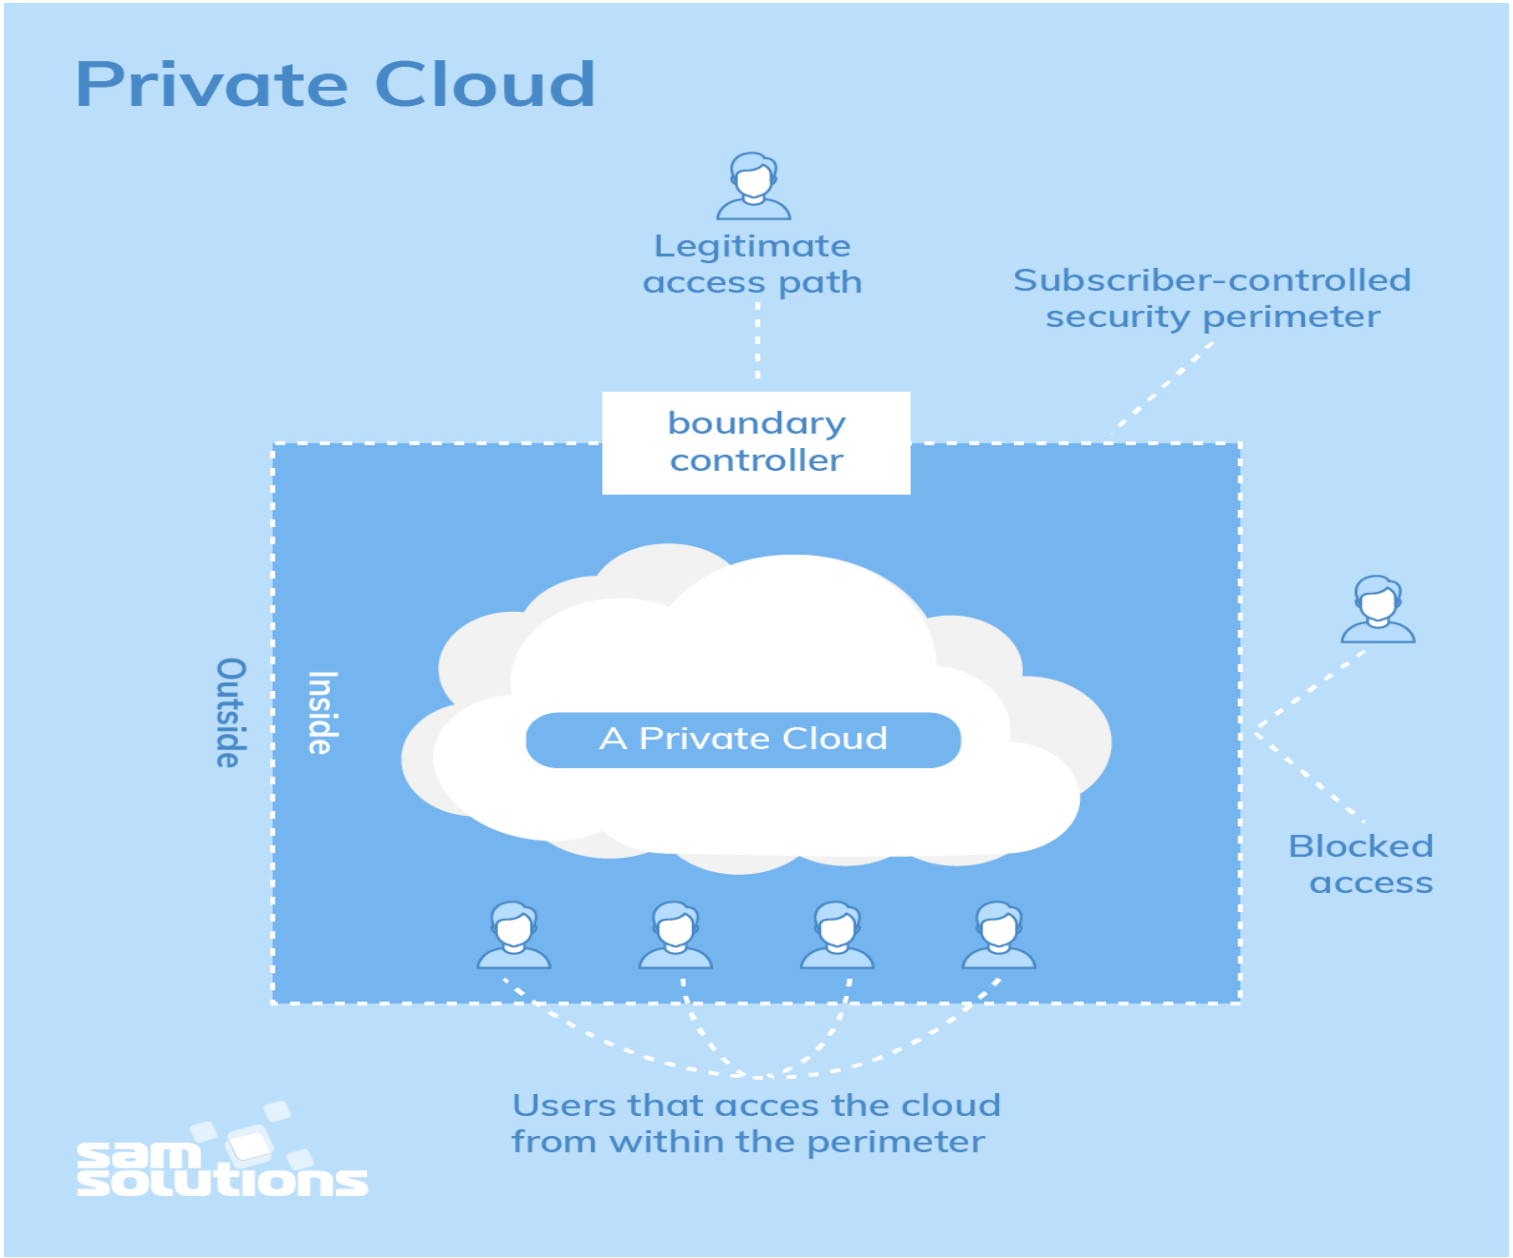
Private Cloud
Legitimate access path
Subscriber-controlled security perimeter
boundary controller
Outside
Inside
A Private Cloud
Blocked access
Users that acces the cloud from within the perimeter
sam solutions

- **Bespoke and flexible development and high scalability, which allows companies to customize their infrastructures in accordance with their requirements**

- **High security, privacy and reliability as only authorized persons can access resources**

- **The major disadvantage of the private cloud deployment model is its cost, as it requires considerable expenses on hardware, software and staff training.**

3. Community Cloud

- A community deployment model largely resembles a private one; the only difference is the set of users. While a private type implies that only one company owns the server, in the case of a community one, several organizations with similar backgrounds share the infrastructure and related resources.

- If the organizations have uniform security, privacy and performance requirements, this multi-tenant data center architecture helps companies achieve their business-specific objectives.

- That is why a community model is particularly suited for organizations that work on joint projects.

Community Cloud

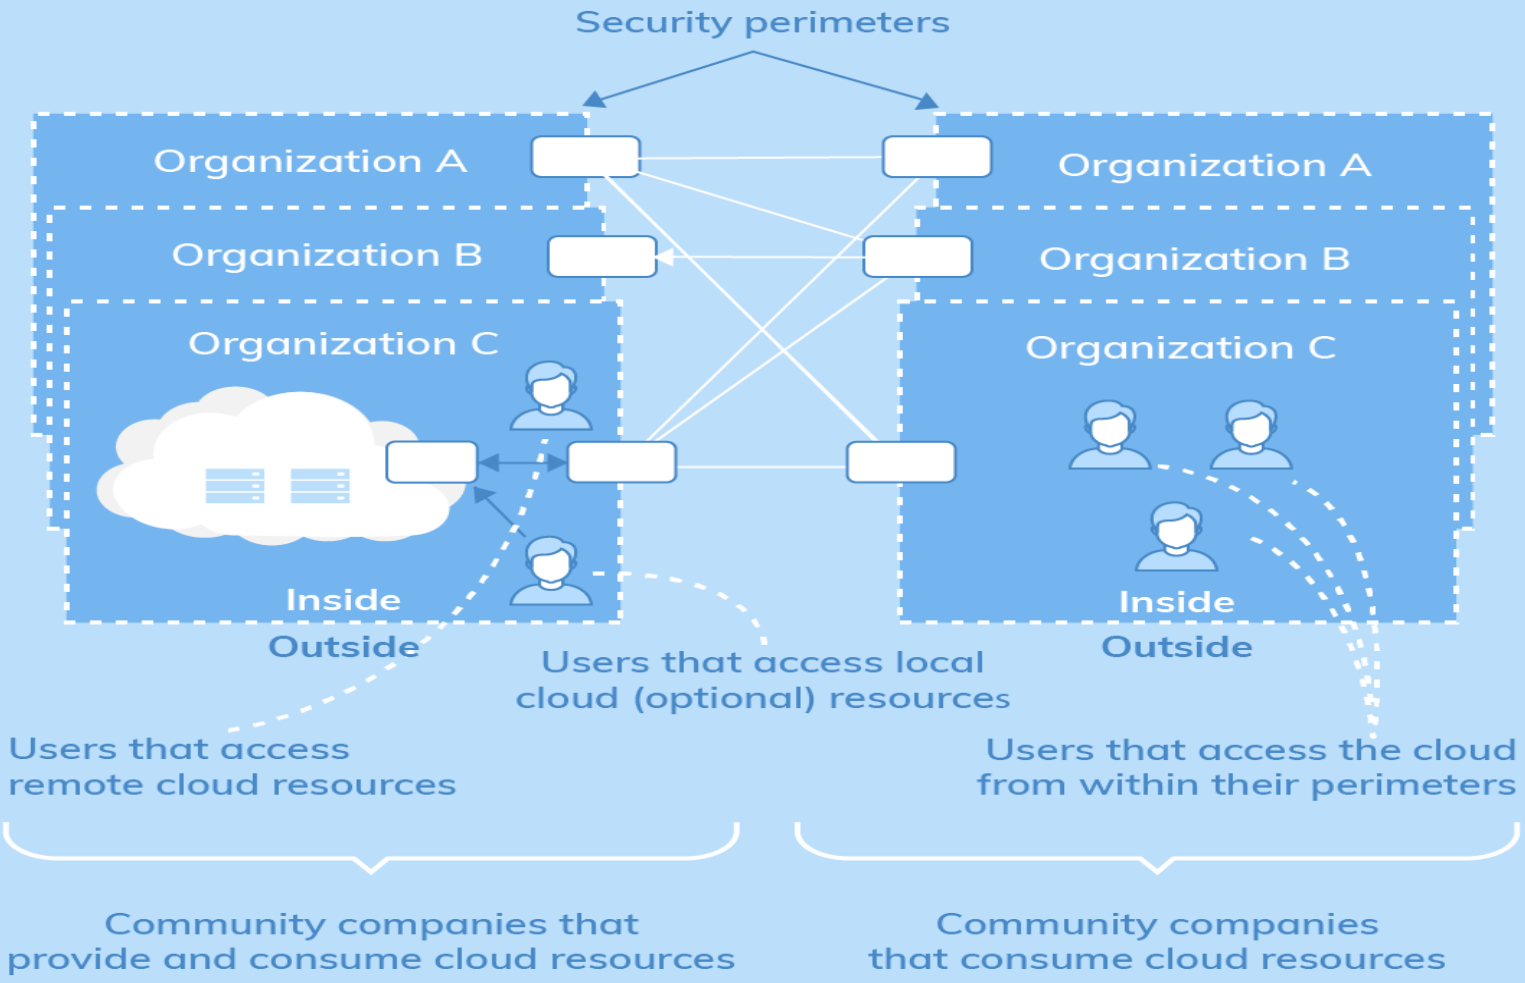

- Cost reduction
- Improved security, privacy and reliability
- Ease of data sharing and collaboration

The Shortcomings of a Community Cloud

- High cost if compared to a public deployment model
- Sharing of fixed storage and bandwidth capacity
- It is not widespread so far

4. Hybrid Cloud

As it is usually the case with any hybrid phenomenon, a hybrid cloud encompasses the best features of the above-mentioned deployment models — public, private and community ones. It allows companies to mix and match the facets of all three types that best suit their requirements.

As an example, a company can balance its load by locating mission-critical workloads on a secure private cloud and deploying less sensitive ones to a public one.

Hybrid cloud deployment model not only safeguards and controls strategically important assets but does so in the most cost- and resource-effective way possible for each specific case

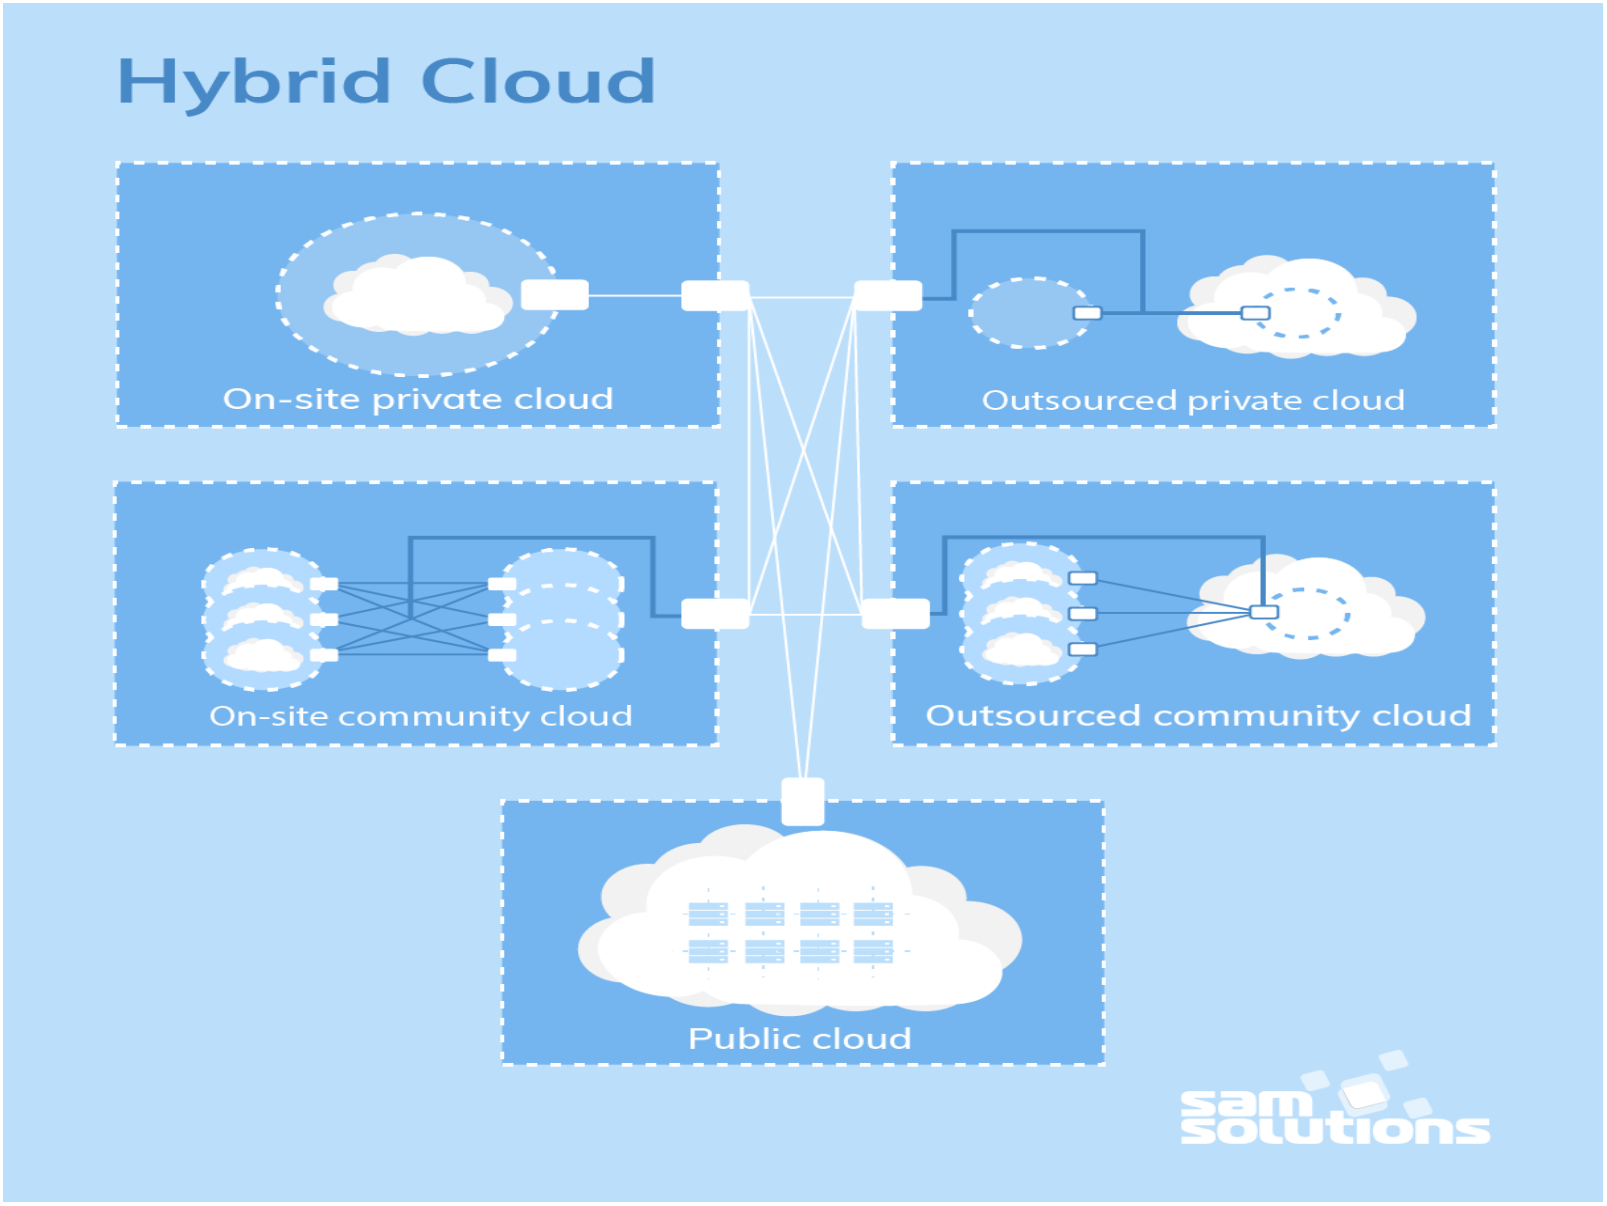
Hybrid Cloud
On-site private cloud
Outsourced private cloud
On-site community cloud
Outsourced community cloud
Public cloud
sam solutions

- Improved security and privacy
- Enhanced scalability and flexibility
- Reasonable price