Chapter Three

Cloud Computing Architecture

3.1 Cloud Components

Cloud Components In a simple, topological sense, a cloud computing solution is made up of several elements: clients, the data center, and distributed servers. As shown in Figure 3-1, these components make up the three parts of a cloud computing solution. Each element has a purpose and plays a specific role in delivering a functional cloud based application, so let's take a closer look.

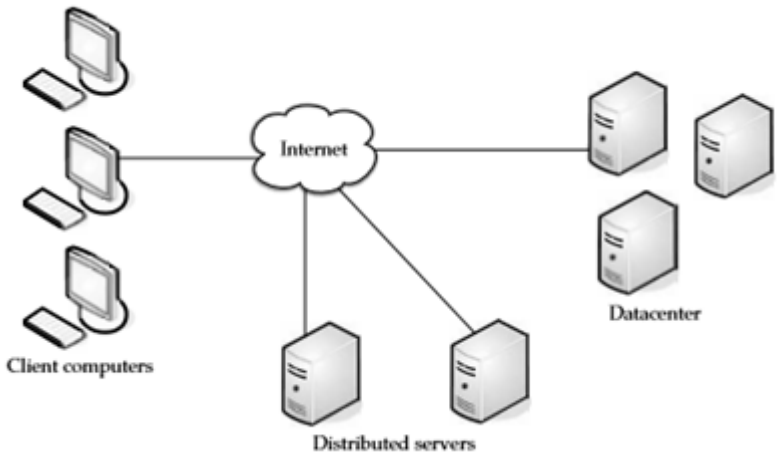

Figure 3.1 Thee components make up cloud computing solution.

1-Clients

Clients are, in a cloud computing architecture, the exact same things that they are in a plain, old, everyday local area network (LAN). They are, typically, the computers that just sit on your desk. But they might also be laptops, tablet computers, mobile phones, or PDAs—all big drivers for cloud computing because of their mobility. Anyway, clients are the devices that the end users interact with to manage their information on the cloud. Clients generally fall into three categories:

- **Mobile:**Mobile devices include PDAs or smartphones, like a Blackberry, Windows Mobile Smartphone, or an iPhone.
- **Thin Clients** are computers that do not have internal hard drives, but rather let the server do all the work, but then display the information.
- **Thick Clients** This type of client is a regular computer, using a web browser like Firefox or Internet Explorer to connect to the cloud.

Thin clients are becoming an increasingly popular solution, because of their price and effect on the environment. Some benefits to using thin clients include

2-Data Center

The data center is the collection of servers where the application to which you subscribe is housed. It could be a large room in the basement of your building or a room full of servers on the other side of the world that you access via the Internet. A growing trend in the IT world is virtualizing servers. That is, software can be installed allowing multiple instances of virtual servers to be used. In this way, you can have half a dozen virtual servers running on one physical server.

3-Distributed Servers

But the servers don't all have to be housed in the same location. Often, servers are in geographically disparate locations. But to you, the cloud subscriber, these servers act as if they're humming away right next to each other. This gives the service provider more flexibility in options and security. For instance, Amazon has their cloud solution in servers all over the world. If something were to happen at one site, causing a failure, the service would still be accessed through another site. Also, if the cloud needs more hardware, they need not throw more servers in the safe room—they can add them at another site and simply make it part of the cloud.

3.2 5-4-3 Principles of Cloud computing

The 5-4-3 principles put forth by NIST describe as shown in Figure 3.2 (a) the five essential characteristic features that promote cloud computing, (b) the four deployment models that are used to narrate the cloud computing opportunities for customers while looking at architectural models, and (c) the three important and basic service offering models of cloud computing.

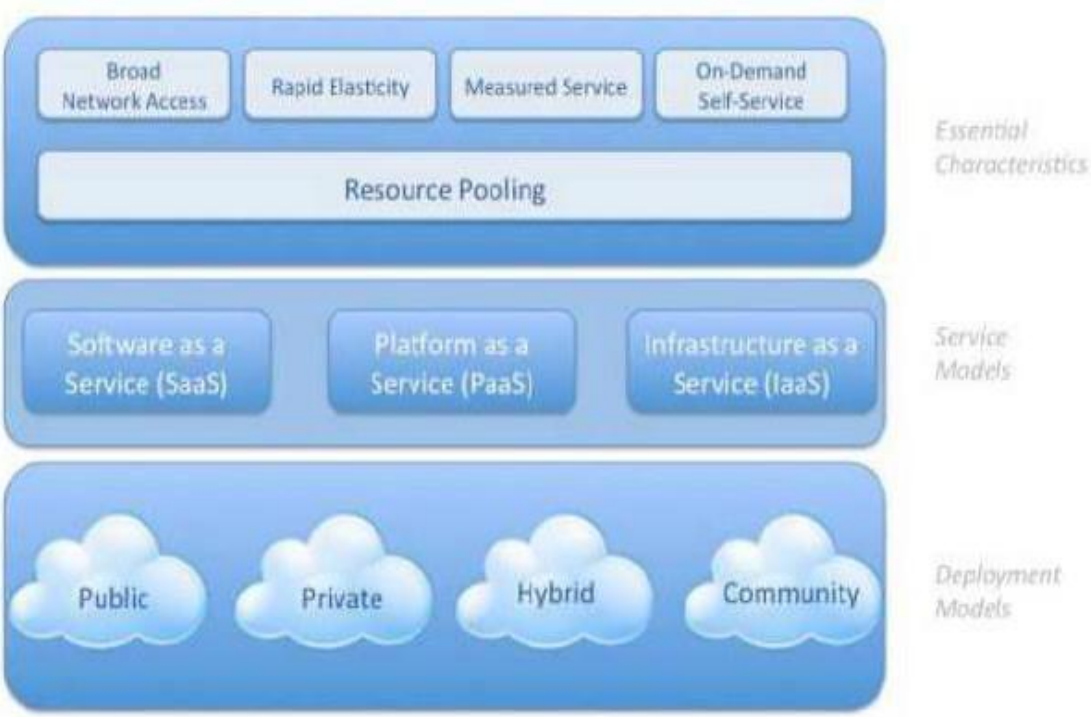

Figure 3.2 Architecture of cloud computing

3.3 Five Essential Characteristics

Cloud computing has five essential characteristics, which are shown in Figure 3.3. Readers can note the word essential, which means that if any of these characteristics is missing, then it is not cloud computing:

1. **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's provider.
2. **Broad 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, laptops, and personal digital assistants [PDAs])

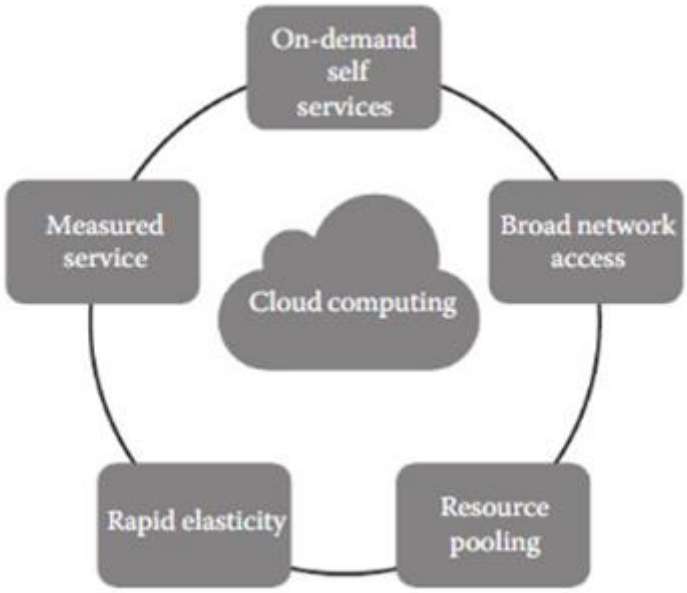

Figure 3.3 the essential characteristics of cloud computing

3. **Elastic 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 the location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.

4. **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.

5. **Measured service**: Cloud systems automatically control and optimize resource use by leveraging a metering capability 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.

3.4 Cloud Computing Models

Cloud computing is a model **that enables the end users to access the shared pool of resources such as compute, network, storage, database, and application as an on-demand service without the need to buy or own it**. The services are provided and managed by the service provider, reducing the management effort from the end user side.. The National Institute of Standards and Technology (NIST) defines three basic service models, namely, IaaS, PaaS, and SaaS, as shown in Figure 3.4 .The NIST definition of the three basic service models is given as follows:

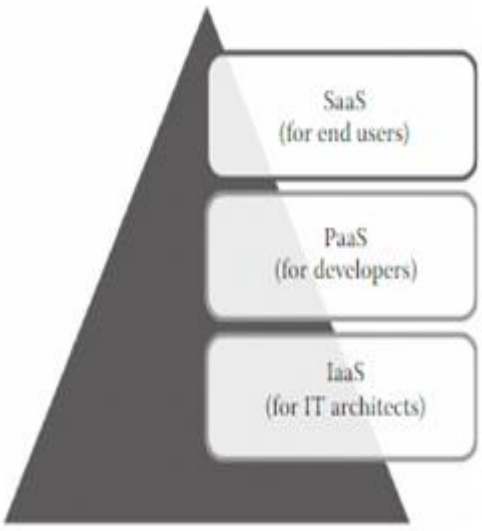

Figure 3.4 Basic cloud service models

1- lnfrastructure as a Service (IaaS):

The ability given to the infrastructure architects to deploy or run any software on the computing resources provided by the service provider. Here, the underlying infrastructures such as compute, network, and storage are managed by the service provider. Thus, the infrastructure architects are exempted from maintaining the data center or underlying infrastructure. The end users are responsible for managing applications that are running on top of the service provider cloud infrastructure. Generally, the IaaS services are provided from the service provider cloud data center.

The end users can access the services from their devices through web command line interface (CLI) or application programming interfaces (APIs)provided by the service providers. Some of the popular IaaS providers include Amazon Web Services (AWS), Google Compute Engine,OpenStack, and Eucalyptus.

Benefits of IaaS

allows the cloud provider to freely locate the infrastructure over the Internet in a cost-effective manner. Some of the key benefits of IaaS are listed below:

1. Full Control of the computing resources through Administrative Access to VMs.
2. Flexible and Efficient renting of Computer Hardware.
3. Portability, Interoperability with Legacy Applications.

2- Platform as a Service (PaaS):

The ability given to developers to develop and deploy an application on the development platform provided by the service provider. Thus, the developers are exempted from managing the development platform and underlying infrastructure. Here, the developers are responsible for managing the deployed application and configuring the development environment. Generally, PaaS services are provided by the service provider on an on-premise or dedicated or hosted cloud infrastructure.

The developers can access the development platform over the Internet through web CLI, web user interface (UI), and integrated development environments (IDEs). Some

of the popular PaaS providers include Google App Engine, Force.com, RedHat OpenShift, Heroku, and Engine Yard.

Benefits of the PaaS model include:

1. **Lower administrative overhead:** Consumer need not to bother much about the administration because it's the responsibility of cloud provider.
2. **Lower total cost of ownership:** Consumer need not purchase expensive hardware, servers, power and data storage.
3. **Scalable solutions:** It is very easy to scale up or down utomatically based on application resource demands.
4. **More current system software:**It is the responsibility of the cloud provider to maintain software versions and patch installations

3-Software as a Service (SaaS):

The ability given to the end users to access an application over the Internet that is hosted and managed by the service provider. Thus, the end users are exempted from managing or controlling an application, the development platform, and the underlying infrastructure. Generally, SaaS services are hosted in service provider–managed or service provider–hosted cloud infrastructure.

The end users can access the services from any thin clients or web browsers. Some of the popular SaaS providers include Saleforce.com, Google Apps, and Microsoft office 365.

Benefits of the SaaS model include:

1. easier administration
2. automatic updates and patch management
3. compatibility: All users will have the same version of software.
4. easier collaboration, for the same reason
5. global accessibility.

The different cloud service models target different audiences. For example, **IaaS** model targets the information technology (IT) architects**, PaaS** targets the developers**, SaaS** targets the end users**.**

Based on the services subscribed, the responsibility of the targeted audience may vary as shown in Figure 3.5.

In IaaS, the end users are responsible for maintaining the development platform and the application running on top of the underlying infrastructure. The IaaS providers are responsible for maintaining the underlying hardwareas shown in Figure 3.5a.

In PaaS, the end users are responsible for managing the application that they have developed. The underlying infrastructure will be maintained by the infrastructure provider as shown in Figure 3.5b.

In SaaS, the end user is free from maintaining the infrastructure, development platform, and application that they are using. All the maintenance will be carried out by the SaaS providers as shown Figure 3.5c.

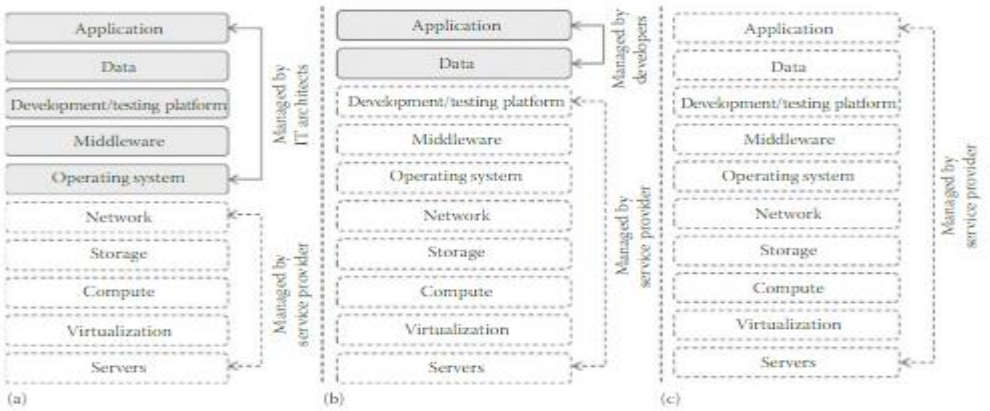

Figure 3.5 User and service provider responsibilities of cloud service models: (a) IaaS, (b) PaaS, and (c) SaaS.

3.5 Cloud Deployment Models

Deployment models can be defined as the different ways in which the cloud can be deployed. These models are fully user centric, that is, these depend on users' requirement and convenience. A user selects a model based on his or her requirement and needs. The NIST defines four different types of cloud deployment models in the cloud:

1. **Private cloud**
2. **Public cloud**
3. **Community cloud**
4. **Hybrid cloud**

The service delivery of cloud services through different deployment models is shown in Figure 3.6.

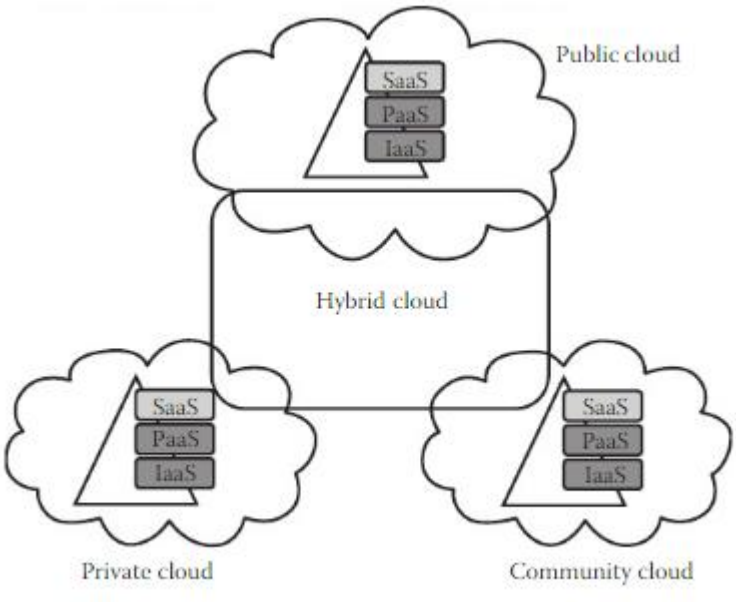

Figure 3.6 Deployment and delivery of different cloud service delivery models.

The classification of the cloud is based on several parameters such as the size of the cloud (number of resources), type of service provider, location, type of users, security, and other issues. The smallest in size is the private cloud (Figure 3.7).

The private cloud is the most basic deployment model that can be deployed by a single organization for its personal use. It is not shared by other organizations, and it is not allowed for public use. The private cloud is to serve the people of an organization. It is usually on premise but can be outsourced also. The next one is the community cloud, which is an extension of the private cloud. Here, the cloud is the same as the private cloud but is shared by several organizations. The community cloud is established for a common cause.

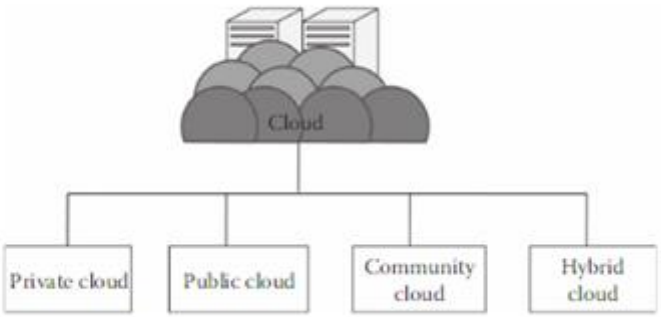

Figure 3.7 Cloud deployment models.

The cause can be anything, but usually it leads to mutual benefits among the participating organizations. The next is the public cloud, which is the opposite of the private cloud. This cloud allows access from any place in the world and is open to the public. This cloud is biggest in size among all other deployment models. The public cloud model is one of the most popular deployment models. The public cloud service provider charges the users on an hourly basis and serve the users according to the service-level agreements (SLAs), which are discussed in the succeeding sections. The next one is the hybrid cloud, which is a combination of other deployments. Usually, it consists of the private and public clouds combined. Several properties of the private cloud are used with the properties of the public cloud. This cloud is one of the upcoming cloud models growing in the industry. All four types of cloud deployments are discussed in detail in subsequent sections.

3.5.1 Private Cloud

In this section, the private cloud deployment model is discussed. According to the National Institute of Standards and Technology (NIST), **private cloud can be defined as the cloud infrastructure that is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units).** It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises .The private cloud in simple terms is the cloud environment created for a single organization. It is usually private to the organization but can be managed by the organization or any other third party. Private cloud can be deployed using Open source tools such as Open stack, Eucalyptus. The private cloud is small in size as compared to other cloud models. Here, the cloud is deployed and maintained by the organizations itself.

1.Characteristics

Certain characteristics of the private cloud are as follows:

1. S**ecure**: The private cloud is secure. This is because usually the private cloud is deployed and managed by the organization itself, and hence there is least chance of data being leaked out of the cloud. In the case of outsourced cloud, the service provider may view the cloud (though governed by SLAs), but there is no other risk from anybody else as all the users belong to the same organization.
2. **Central control**: The organization mostly has full control over the cloud as usually the private cloud is managed by the organization itself. Thus, when managed by the organization itself, there is no need for the organization to rely on anybody.

3. **Weak SLAs:** Formal SLAs may or may not exist in a private cloud. But if they exist they are weak as it is between the organization and the users of the same organization. Thus, high availability and good service may or may not be available. This depends on the organization that is controlling the cloud.

2. Suitability

Suitability refers to the instances where this cloud model can be used. It also signifies the most suitable conditions and environment where this cloud model can be used, such as the following:

1. The organizations or enterprises that require a separate cloud for their personal or official use.
2. The organizations or enterprises that have a sufficient amount of funds as managing and maintaining a cloud is a costly affair
3. The organizations or enterprises that consider data security to be important
4. The organizations that want autonomy and complete control over the cloud
5. The organizations that have a less number of users.
6. The organizations that have prebuilt infrastructure for deploying the cloud and are ready for timely maintenance of the cloud for efficient functioning.
7. Special care needs to be taken and resources should be available for troubleshooting.

2. The private cloud platform is not suitable for the following:

- The organizations that have high user base
- The organizations that have financial constraints
- The organizations that do not have prebuilt infrastructure

3. Advantages

- The cloud is small in size and is easy to maintain.
- It provides a high level of security and privacy to the user.
- It is controlled by the organization.

4. Disadvantages

- For the private cloud, budget is a constraint.
- The private clouds have loose SLAs.

3.5.2 Public Cloud

According to NIST, the **public cloud is the cloud infrastructure that 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. The typical public cloud is depicted in Figure 3.8. Public cloud consists of users from all over the world. A user can simply purchase resources on an hourly basis and work with the resources. There is no need of any prebuilt infrastructure for using the public cloud. These resources are available in the cloud provider's premises. Usually, cloud providers accept all the requests , and hence, the resources in the service providers' end are considered infinite in one aspect. Some of the well-known examples of the public cloud are Amazon AWS, Microsoft Azure, etc.

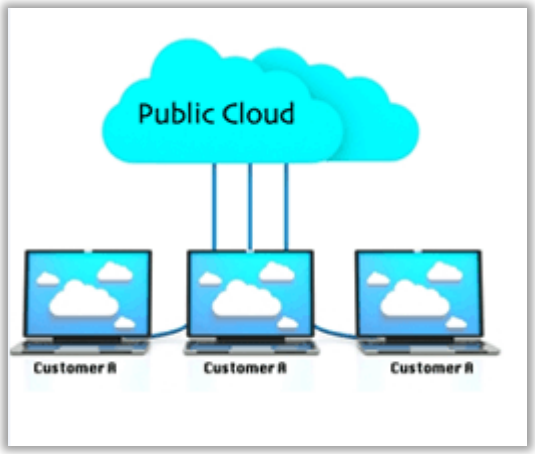

1. Characteristics

1. **Highly scalable:** The public cloud is highly scalable. The resources in the public cloud are large in number and the service providers make sure that all the requests are granted. Hence, the public cloud is considered to be scalable.
2. **Affordable:** The public cloud is offered to the public on a pay-as-you-go basis; hence, the user has to pay only for what he or she is using (usually on a per-hour basis). And, this does not involve any cost related to the deployment.

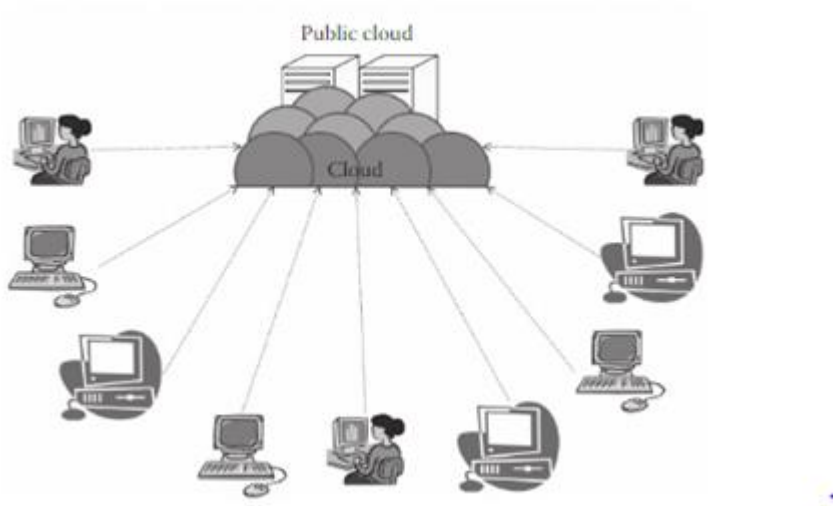

Figure 3.8 public cloud

3. **Less secure:** The public cloud is less secure out of all the four deployment models. This is because the public cloud is offered by a third party and they have full control over the cloud. Though the SLAs ensure privacy, still there is a high risk of data being leaked.
4. **Highly available:** The public cloud is highly available because any body from any part of the world can access the public cloud with proper permission, and this is not possible in other models as geographical or other access restrictions might be there.
5. **Stringent SLAs:** SLA is very stringent in the case of the public cloud. As the service provider's business reputation and customer strength are totally dependent on the cloud services, they follow the SLA strictly and violations are avoided. These SLAs are very competitive.

2. Suitability

There are several occasions and environments where the public cloud is suitable. Thus, the suitability of the public cloud is described. The public cloud can be used whenever the following applies:

1. The requirement for resources is large, that is, there is large user base.
2. The requirement for resources is varying.
3. There is no physical infrastructure available.
4. An organization has financial constraints.

3. The public cloud is not suitable, where the following applies:

- Security is very important.
- Organization expects autonomy.
- Third-party reliability is not preferred.

3. Issues

Several issues pertaining to the public cloud are as follows:

1. **SLA:** Unlike the private cloud, here the number of users is more and so are the numbers of service agreements. The service provider is answerable to all the users. The users here are diverse. The SLA will cover all the users from all parts of the world. The service provider has to guarantee all the users a fair share without any priority. Having the same SLA for all users is what is usually expected, but it depends on the service provider to have the same SLA for all the users irrespective of the place they are.
2. **Network:** The network plays a major role in the public cloud. Each and every user getting the services of the cloud gets it through the Internet. The services are accessed through the Internet by all the users, and hence, the service delivery wholly depends on the network. Unlike the private cloud where the organization takes responsibility for the network, here the service provider is not responsible for the network. The service provider is responsible for providing proper service to the customer, and once the services are given from the service provider, it goes

on in transit to the user. The user will be charged for even if he or she has problem due to the network.The network usually consists of a high bandwidth and has a low latency. This is because the connection is only inside the organization. Network management is easier in this case.

3. **Performance:** As mentioned, the performance of a cloud delivery model primarily depends on the network and the resources. The service provider has to adequately manage the resources and the network. As the number of users increases, it is a challenging task for the service providers to give good performance.
4. **Multitenancy:** The resources are shared, that is, multiple users share the resources, hence the term multitenant. Due to this property, there is a high risk of data being leaked or a possible unprivileged access.
5. **Location:** The location of the public cloud is an issue. As the public cloud is fragmented and is located in different regions, the access to these clouds involves a lot of data transfers through the Internet. There are several issues related to the location. For example, a user from India might be using the public cloud and he might have to access his personal resources from other countries. This is not good as the data are being stored in some other country.
6. **Security and data privacy :** Security and data privacy are the biggest challenges in the public cloud. As data are stored in different places around the globe, data security is a very big issue. A user storing the data outside his or her country has a risk of the data being viewed by other people as that does not come under the jurisdiction of the user's country. Though this might not always be true, but it may happen.
7. **Laws and conflicts:** The data are stored in different places of the world in different countries. Hence, data centers are bound to laws of the country in which they are located. This creates many conflicts and problems for the service providers and the users.
8. **Cloud management:** Here, the number of users is more, and so the management is difficult. The jobs here are time critical, and as the number of users increases, it becomes more difficult. Inefficient management of resources will lead to resource shortage, and user service might be affected. It has a direct impact on SLA and may cause SLA violation.

9. **Maintenance:** Maintaining the whole cloud is another task. This involves continuous check of the resources, network, and other such parameters for long-lasting efficient delivery of the service. There source provider has to continuously change the resource components from time to time. The task of maintenance is very crucial in the public cloud. The good the cloud is maintained, the better is the quality of service. Here, the cloud data center is where the maintenance happens; continuously, the disks are replaced from time to time.

The issues discussed earlier will help to understand the public cloud. Before using the public cloud, one has to choose a cloud service provider. One can choose the public cloud based on certain parameters like SLA violations, security, and cost of resources. Thus, a cloud's quality is determined by the SLA violation it does. The less the SLA violation it does, the better the cloud is. This is one way of selecting the public cloud; another way is by cost. If the job for which the resources are used is not time sensitive, then the service provider who offers the least cost is selected. There following are several advantages and disadvantages of public clouds.

4. Advantages

- There is no need of establishing infrastructure for setting up a cloud.
- There is no need for maintaining the cloud.
- They are comparatively less costly than other cloud models.
- Strict SLAs are followed.
- There is no limit for the number of users.
- The public cloud is highly scalable.

5. Disadvantages

- Security is an issue
- Privacy and organizational autonomy are not possible.

3.5.3 Community Cloud

According to NIST, the community cloud is the cloud infrastructure that is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises. It is a further extension of the private cloud. Here, a private cloud is shared between several organizations. Either the organizations or a single organization may collectively maintain the cloud.

The main advantage of the public cloud is that the organizations are able to share the resources among themselves based on specific concerns. Thus, here the organizations are able to extract the power of the cloud, which is much bigger than the private cloud, and at the same time, they are able touse it at a usually less cost. The community is formed based on any commoncause, but eventually, all the members of the community are benefitted.This model is very suitable for organizations that cannot afford a private cloud and cannot rely on the public cloud either. Figure 4.5 describes the community cloud.

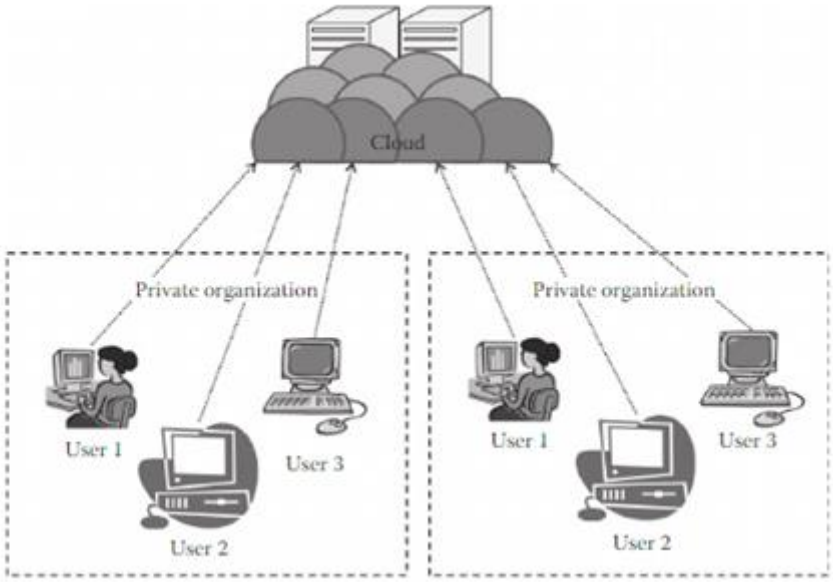

Figure 3.9 community cloud

1. Characteristics

1. **Collaborative and distributive maintenance:** The community cloud is wholly collaborative, and usually no single party has full control over the whole cloud (in some cases, it may be controlled by one party). This is usually distributive, and hence, better cooperation gives better results. Even though it may be outsourced, collaboration based on purpose always proves to be beneficial.
2. **Partially secure:** Partially secure refers to the property of the com-munity cloud where few organizations share the cloud, so there is a possibility that the data can be leaked from one organization to another, though it is safe from the outside world.
3. **Cost effective:** The community cloud is cost effective as the whole cloud is being shared by several organizations or a community. Usually, not only cost but every other sharable responsibilities are also shared or divided among the groups.

2. Suitability

This kind of cloud is suitable for organizations that

- Want to establish a private cloud but have financial constraint
- Do not want to complete maintenance responsibility of the cloud
- Want to establish the cloud in order to collaborate with other clouds
- Want to have a collaborative cloud with more security features than the public cloud

3. This cloud is not suitable for organizations that

- Prefer autonomy and control over the cloud
- Does not want to collaborate with other organizations

4. Advantages

- It allows establishing a low-cost private cloud.
- It allows collaborative work on the cloud
- It allows sharing of responsibilities among the organization.
- It has better security than the public cloud.

4. Disadvantages

- Autonomy of an organization is lost
- Security features are not as good as the private cloud
- It is not suitable if there is no collaboration.

3.5.4 Hybrid Cloud

According to NIST, the hybrid cloud can be defined as the cloud infrastructure that is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability.

The hybrid cloud usually is a combination of both public and private clouds. This is aimed at combining the advantages of private and public clouds. The usual method of using the hybrid cloud is to have a private cloud initially, and then for additional resources, the public cloud is used. There are several advantages of the hybrid cloud. The hybrid cloud can be regarded as a private cloud extended to the public cloud.

This aims at utilizing the power of the public cloud by retaining the properties of the private cloud. One of the popular examples forthe hybrid cloud is Eucalyptus . Eucalyptus was initially designed for the private cloud and is basically a private cloud, but now it also supports hybrid cloud. Figure 3.10 shows the hybrid cloud. The hybrid cloud can be further extended into a vast area of federated clouds that is discussed in subsequent chapters.

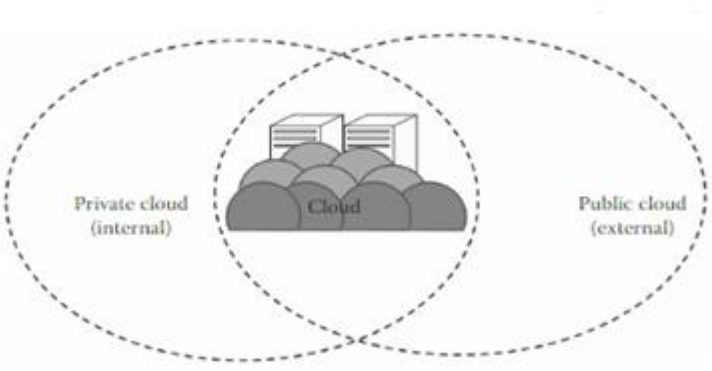

Figure 3.10 Hybird cloud

1. Characteristics

1. **Scalable :** The hybrid cloud is a combination of one or more deployment models. Usually, the private with public cloud gives hybrid cloud. The main reason of having a hybrid cloud is to use the property of a public cloud with a private cloud environment. The public cloud is used whenever needed; hence, as the public cloud is scalable, the hybrid cloud with the help of its public counter part is also scalable.
2. **Partially secure:** The hybrid cloud usually is a combination of public andprivate. The private cloud is considered to be secured, but as the hybridcloud also uses the public cloud, there is high risk of security breach. Thus, it cannot be fully termed as secure but as partially secure.
3. **Stringent SLAs:** As the hybrid cloud involved a public cloud intervention, the SLAs are stringent and might as per the public cloud service provider. But overall, the SLAs are more stringent than the private cloud.
4. **Complex cloud management:** Cloud management is complex and is a difficult task in the hybrid cloud as it involves more than one type of deployment models and also the numbers of users are high.

2 Suitability

The hybrid cloud environment is suitable for

- Organizations that want the private cloud environment with the scalability of the public cloud
- Organizations that require more security than the public cloud

3. The hybrid cloud **is not suitable for**

- Organizations that consider security as a prime objective
- Organizations that will not be able to handle hybrid cloud management

3 Issues

The cloud can be analyzed in the following aspects:

1. **SLA:** SLA is one of the important aspects of the hybrid cloud as both private and public are involved. There is a right combination of SLAs between the clouds. The private cloud does not have stringent agreements, whereas the public cloud has certain strict rules to be covered. The SLAs to be covered under each purview are clearly defined, and it wholly depends on the service provider (private cloud) to provide efficient services to the customers.
2. **Network:** The network is usually a private network, and when everthere is a necessity, the public cloud is used through the Internet. Unlike the public cloud, here there is a private network also. Thus, a considerable amount of effort is required to maintain the network. The organization takes the responsibility from the network.
3. **Performance:** The hybrid cloud is a special type of cloud in which the private environment is maintained with access to the public cloud whenever required. Thus, here again a feel of an infinite resource is restored. The cloud provider (private cloud) is responsible for providing the cloud.
4. **Multitenancy:** Multitenancy is an issue in the hybrid cloud as it involves the public cloud in addition to the private cloud. Thus, this property can be misused and the breaches will have adverse effects as some parts of the cloud go public.
5. **Location:** Like a private cloud, the location of these clouds can be on premise or off premise and they can be outsourced. They will have all the issues related to the private cloud; in addition to that, issues related to the public cloud will also come into picture whenever there is intermittent access to the public cloud.
6. **Security and privacy:** Whenever the user is provided services using the public cloud, security and privacy become more stringent. As it is the public cloud, the threat of data being lost is high.
7. **Laws and conflicts:** Several laws of other countries come under the purview as the public cloud is involved, and usually these public clouds are situated outside the country's boundaries.
8. **Cloud management:** Here, everything is managed by the private cloud service provider.

9. **Cloud maintenance:** Cloud maintenance is of the same complexity as the private cloud; here, only the resources under the purview of the private cloud need to be maintained. It involves a high cost of maintenance.

The hybrid cloud is one of the fastest growing deployment models, which is now being discussed because of its characteristics as discussed earlier. The issues discussed provide an overview about the difference between the other cloud models and the hybrid cloud model. There is another part of the cloud called as federated cloud that is described in the subsequent chapter. There are several advantages and disadvantages of the hybrid cloud.

4. Advantages

- It gives the power of both the private and public clouds.
- It is highly scalable.
- It provides better security than the public cloud

5. Disadvantages

- The security features are not as good as the public cloud.
- Managing a hybrid cloud is complex.
- It has stringent SLAs.