Network Security
Professor Gaurav S. Kasbekar
Department of Electrical Engineering
Indian Institute of Technology, Bombay
Week - 12
Lecture - 67

Cloud Security: Part 1

Hello, in this lecture and the next few lectures, we will discuss cloud security. First, we will discuss the basics of cloud computing, and then we will discuss its security aspects. So, what is meant by cloud computing? It is very popular these days, and all of us use it. So, there is an increasing trend in many organizations to move a substantial part or all of their information technology operations to an internet-connected infrastructure.

This is called enterprise cloud computing. And this is illustrated in this figure. This is a cloud, which is a collection of a lot of computing resources such as servers, storage, applications, and other services. These are present in a data center. So, this data center has a huge amount of servers, storage, and so on.

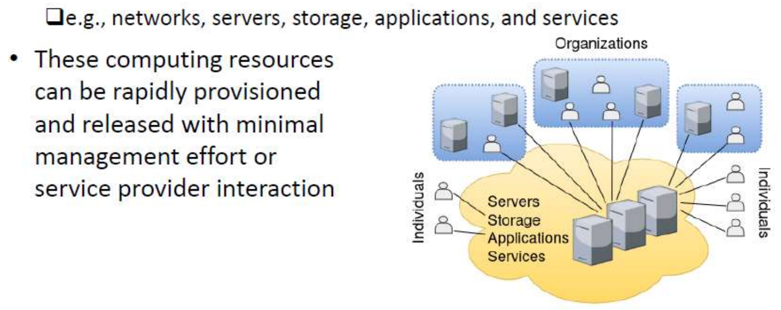

❑e.g., networks, servers, storage, applications, and services

- These computing resources can be rapidly provisioned and released with minimal management effort or service provider interaction

And users in other organizations connect to this data center and use these resources of the data center. And individuals like us can also connect to this data center and use its resources. So, examples are email services such as Gmail and Yahoo Mail, and editing services such as Overleaf and so on. So, here we connect to some data center and use the processing power, storage, etc., of that data center. This data center is the cloud.

These users use the resources of the cloud. Cloud computing is a model for enabling ubiquitous, convenient, and on-demand network access to a shared pool of configurable computing resources. So, these computing resources may be present in one data center or in multiple data centers. And users can connect from anywhere to one of these data centers. Hence, this is ubiquitous, and it is also convenient.

For example, one does not have to back up the copies of files from one's own computer to a flash drive, for example. All the data is present in the cloud, and it is also on demand. For example, if a user wants to perform a certain task, in that case, the user can acquire network access to the cloud on demand and use the resources of the cloud. Examples of such computing resources are networks, servers, storage, applications, and services. So, these computing resources can be rapidly provisioned and released with minimal management effort or service provider interaction.

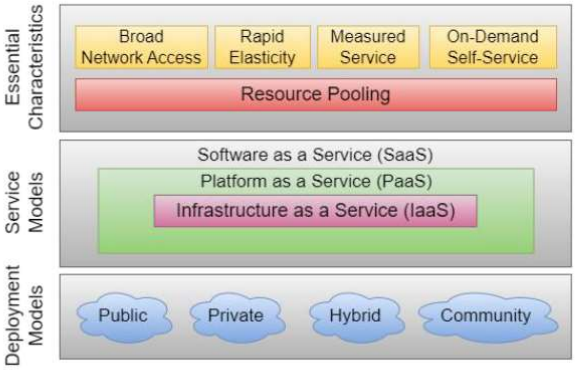

So, whenever a user requires some computing resources, these can be allocated to that user and once the user has finished their task, these resources can be released and they can be used by other users. We now discuss the cloud model. It consists of these different features. One is, there are five essential characteristics, which are shown in this upper part of the figure. So, these five characteristics are broad network access, rapid elasticity, measured service, on-demand self-service, and resource pooling.

There are three service models, which are shown in the middle part. These are Software as a Service, Platform as a Service, and Infrastructure as a Service. And there are four deployment models, which are shown here. These are public cloud, private cloud, hybrid cloud, and community cloud. So, we'll discuss these different aspects in this lecture and the next lecture.

We start with the essential characteristics. So, as we mentioned on the previous slide, these five are the essential characteristics of cloud computing. The first one is broad network access. Capabilities are available over the network, and anyone can use the resources of the cloud from anywhere on the internet; hence, the access is broad. These resources can be accessed through standard mechanisms, such as using any browser and connecting to the cloud.

And these resources can be used by heterogeneous thin or thick client platforms, for example, mobile phones, laptops, PDAs, and so on. So, thick client platforms are those that have a lot of resources, and thin platforms are those that have few resources. Possibly just a device with a simple interface, such as a browser on it. So, that's the first characteristic of cloud computing: broad network access. Then, the next characteristic is rapid elasticity.

Resources can be expanded and reduced according to the specific service requirements of the clients. Here's an example. A client may need a large number of server resources for the duration of a specific task. For example, the client may want to run some simulations for some experiment. In that case, the client requires a lot of server resources, so the cloud can allocate these large number of server resources to the client for carrying out this specific task, that is the simulations.

And once this task is complete, the resources can be released and they can be used by other users. So, this is the second essential characteristic, that is rapid elasticity. Then, the third essential characteristic is measured service. Cloud systems use a metering capability which measures the usage of the cloud resources by users. And this metering capability is appropriate to the type of service, for example, storage, processing, bandwidth, and so on.

So, for example, how much storage, how many MBs or GBs of data are being stored by a particular user, that is measured. Then how much CPU time is being used by users? That is also measured; that is processing. And bandwidth in terms of Mbps or Gbps that is being consumed, that is also measured. So, this metering capability is used to automatically control and optimize resource use. It provides transparency for the provider and consumer of the utilized service.

So, the provider knows exactly how much of its resources are being used by consumers, and consumers know how much resources they are using. So, this can help in planning the control and optimization of resources and billing—that is, how much the consumer has to pay the provider, and so on and so forth. So, all the usage that happens on the cloud is measured. So, that's the third essential characteristic of cloud computing. Then, the fourth one is on-demand self-service.

A Cloud Service Consumer (CSC), can unilaterally provision computing capabilities, such as server time and network storage. So, CSC is a term we'll use. A CSC is a user of the cloud—that is, the cloud service consumer. So, the CSC can unilaterally provision computing capabilities. And this is done as required automatically, without requiring human interaction with the service provider.

So, there is an automatic mechanism provided by the service provider, using which the CSC can unilaterally provision computing capabilities as required. The fifth and final essential characteristic of cloud computing is resource pooling. The provider's computing resources—for example, storage, processing, memory, network bandwidth, virtual machines, and so on—are pooled to serve multiple CSCs. For example, these computing resources from many locations can be pooled together to serve CSCs. Different physical and virtual resources are dynamically assigned and reassigned according to consumer demand.

The CSC generally has no control or knowledge of the exact location of the provided resources. For example, the resources that the user is using, which location those resources belong to. Do they belong to a data center in Mumbai or a data center in Bangalore? So, that is not known to the user of the resource. But the user may be able to specify the location at a higher level of abstraction.

For example, the user may say that I want to use a data center from a certain country or state or a particular data center. So, that may be for various reasons, such as security and so on. So, these are the essential characteristics of cloud computing. Now, we have discussed essential characteristics. Next, we discuss the different service models in cloud computing. So, the service models in cloud computing are shown here.

One is software as a service, then platform as a service, and infrastructure as a service. So, in software as a service, the provider of the cloud that provides not only the hardware but also the operating system and the applications, and the consumer uses the applications from the cloud. And in platform as a service, the provider provides the hardware as well as the operating system, and the user can run their own applications on top of the operating system provided by the provider. And in infrastructure as a service, the provider just provides the hardware resources, and the user runs their own operating system as well as the applications on top of the operating system using the infrastructure provided by the provider. So, let's discuss these three service models in some more detail.

So, we start with software as a service. So, this model provides service to customers in the form of application software running on and accessible in the cloud. So, there are some apps which are running in the cloud, and they are managed by the cloud provider, and customers can use these apps which run in the cloud. This enables the customer to use the cloud provider's applications running on the provider's cloud infrastructure. So, these applications are installed by the cloud provider, and they run on the cloud infrastructure.

The customer can use these applications. These applications are accessible from various client devices through a simple interface such as a web browser. So, regardless of what application it is, it can be run from a web browser. So, this model is quite convenient for the customer. Instead of obtaining desktop and server licenses for software products that the customer uses, an enterprise obtains the same functions from the cloud service.

So, the customer does not have to acquire licenses for software. These licenses are acquired by the cloud provider, and the customer uses the applications provided by the provider. The use of software as a service avoids the complexity of software installation, maintenance, upgrades, and patches. So, all this complexity is shifted to the cloud provider. The customer does not have to deal with these aspects.

Some examples of software as a service are Google, Gmail, Microsoft 365, Salesforce, Citrix, GoToMeeting, and Cisco WebEx. So, these are all examples of this model, software as a service. Next, we discuss platform as a service, where the operating system and hardware are provided by the cloud provider, and the customer runs their own applications on top of this operating system. This model provides services to customers in the form of a platform on which the customer's applications can run. This platform is the operating system.

It enables customers to deploy onto the cloud infrastructure customer-created or acquired applications. The customer can create some applications or acquire these applications from third-party software sellers. They can deploy these applications onto the cloud infrastructure. In particular, on the operating system provided by the cloud provider. PaaS is an operating system in the cloud.

Customers can run their own applications or acquired applications over this operating system. Some examples of services that use this model are the following. One is Google App Engine, Engine Yard, Heroku, Microsoft Azure Cloud Services, and Apache Stratos. These are all platform as a service instances. Next, we discuss infrastructure as a service, which is the third service model.

So, here the hardware resources are provided by the cloud provider, but the OS as well as the applications on it are run by the customer. So, the customer has access to the hardware resources of the underlying cloud infrastructure. The customer has control over the operating systems as well as the deployed applications running over the operating systems. So, IaaS provides virtual machines and other virtualized hardware. So, these virtual machines run over the physical hardware, and there is isolation between different virtual machines.

Similarly, there can be other virtualized hardware that can run on the physical resources of the cloud provider. IaaS offers the customer processing, storage networks, and other fundamental computing resources so that the customer is able to deploy and run arbitrary software, which can include operating systems as well as applications. Some examples of services that use this model are the following. Amazon Elastic Compute Cloud, or Amazon EC2, Microsoft Azure, Google Compute Engine, and Rackspace. So, this concludes our discussion of the three service models.

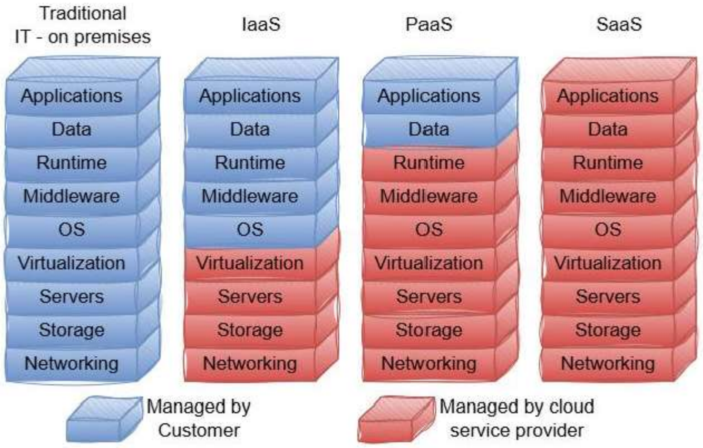

This figure compares the functions implemented by the cloud service provider for the three service models. This leftmost figure shows the case where there is a traditional IT on-premises. So, the blue coloring shows that it is managed by the customer, and red coloring shows that it is managed by the cloud service provider. So, in the traditional model, the entire stack is managed by the customer. This includes applications, data, runtime, middleware, OS, virtualization, servers, storage, and networking.

So, all of this needs to be managed by the customer. Typically, it is managed by the system administrators of the enterprise. This involves a lot of work from the customer. Then, the next model is IaaS, Infrastructure as a Service. So, here the hardware resources, which can include networking, storage, servers, and virtualization, which divide these physical resources into different virtual machines.

So, all these hardware resources are managed by the cloud service provider in IaaS. And the OS, as well as the applications, middleware, runtime, and data. So, these are managed by the customer. Then, the next model is Platform as a Service. So, here we can see that all the hardware resources up to virtualization, as well as the OS, middleware, and runtime, these are all managed by the cloud service provider, and applications and data are managed by the customer.

And in the final model, Software as a Service, the entire stack is managed by the cloud service provider, including the applications and data. So, the applications are installed by the cloud service provider, and they can be used by customers. So, there is a gradation among these different service models. We can see that from the customer's point of view, it is easiest to use Software as a Service and most difficult to use the traditional IT model. But the cost in terms of the payment to the cloud provider, that is the most in the case of software as a service, and there is no payment to the cloud provider in the case of traditional IT.

But everything has to be managed by the customer. So, in summary, we introduced cloud computing, we discussed the essential characteristics of cloud computing, and we discussed the three different service models that are used in cloud computing. We'll continue our discussion on cloud computing in the next lecture.