Introduction to Information Security
Prof. V. Kamakoti
Department of Computer Science and Engineering
Indian Institute of Technology, Madras

Lecture - 26

(Refer Slide Time: 00:33)

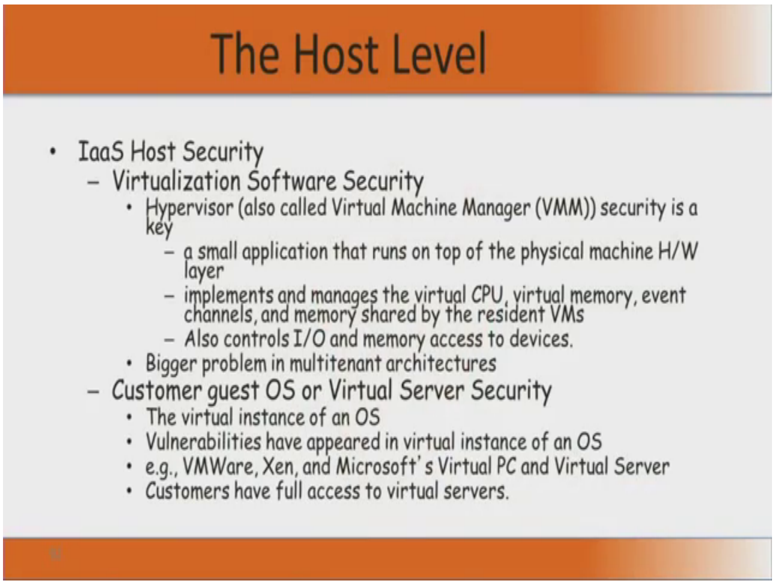

The Host Level

- IaaS Host Security
  - Virtualization Software Security
    - Hypervisor (also called Virtual Machine Manager (VMM)) security is a key
      - a small application that runs on top of the physical machine H/W layer
      - implements and manages the virtual CPU, virtual memory, event channels, and memory shared by the resident VMs
      - Also controls I/O and memory access to devices.
    - Bigger problem in multitenant architectures
  - Customer guest OS or Virtual Server Security
    - The virtual instance of an OS
    - Vulnerabilities have appeared in virtual instance of an OS
    - e.g., VMWare, Xen, and Microsoft's Virtual PC and Virtual Server
    - Customers have full access to virtual servers.

Next is the host level security. So, when the most important host level security is. So, the host can be used in three forms. I can use it just as an infrastructure. I can use it as a SAAS or a PAAS right software as a service or platform as a service or infrastructure as a service. Now, the security becomes extremely from a cloud service providers point of view the security becomes extremely complex, when we look at the infrastructure as a service, host security.

When I give an infrastructure as a service how will that whole infrastructure look like. First and foremost there will be the hardware, on top of this hardware there will be a virtualization software. This is called the hypervisor. Soon, we will see a block diagram there. The hypervisor is responsible for the virtualization. So, the hypervisor will provide several virtual machines to the operating to the higher layers.

So, the hypervisor is one who is responsible for the upper layers to look at the single physical machine as multiple machines. So, on this hypervisor which is also called the virtual machine manager you can actually instantiate several virtual machines. And each of this virtual machine can start executing different software or each of this virtual machine can run their own operating system on which the own software can be run.

Now, we have to look at the security at that level of the hypervisor level. And then comes what you call as the customer guest OS are the virtual server. And the custom in a IAAS type of model, the customer has complete control over the virtual server. Now, whatever I have told here could be best explained with a figure and this is the figure here.

(Refer Slide Time: 02:33)

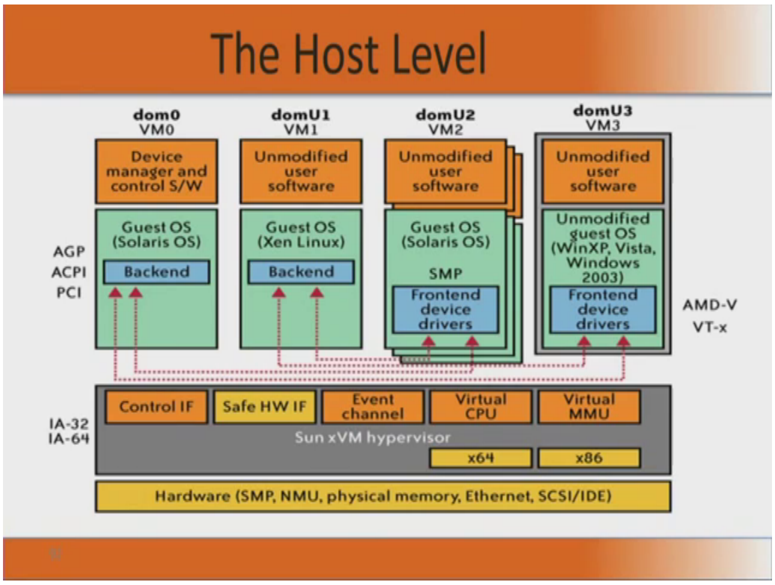

The Host Level

Now, you will see that there is one single hardware which is in the bottom most rectangle. That hardware can be a symmetric multiprocessor machine, multi core machine. It has physical lot of it has physical memory, it has all peripherals like a disk your ethernet and it can also have different varieties of things, like it can be a non-uniform memory access machine. So, it could several memory banks. So, thats an architecture.

On top of this hardware there is the actual what you call as the hypervisor. The hypervisor is call is also called the VMM which is the virtual machine manager. And what will this hyper vis hypervisor has lot of interfaces to the upper layers, there is a control interface, there is a safe hardware interface, there is a even channel, there is a virtual CPU, virtual memory management unit and then the hypervisor. So, that is to the higher layers that is software thats going to run on top of it.

The interface to the lower layer there are two interfaces one is the x64 and x86 which is your IA32 and IA64 bit architectures. So, the hypervisor may is an interface with the hardware and the upper layers. What is a role of this hypervisor? we will talk much more on the architecture of hypervisors in other level two or level three courses in great detail, but for our understanding now and hypervisor is one responsible for creating several virtual machines on the top.

Now, in this instance as you see in figure, there are four virtual machines that are basically generated or instantiated. So, you see there are the and these virtual machines can do different type of activities. All these virtual machines the one you see the virtual machine 0 can be just a device manager and a software controller. Then you have virtual machine 1 virtual machine 2 virtual machine 3 it can be for three users. So, the 1 virtual machine which is on left hand side is for the administrator and the remaining three could be for the user.

Each of these virtual machines. So, four virtual machines are running. So, the administrator and the three users will think that the server that hardware is exclusively for them. They will not even know the existence of the other virtual machines that are running. They are isolated from the other virtual machine. Now, your hardware is very fast. So, your hardware can actually run four software at the same time. That if I try and run these four software can or not on the same operating system.

One can be a windows, one can be a Solaris, another can be a Unix, but all these four software can run at the same time on the same server. And these four software would be using will be compatible to different operating systems. So, essentially what it means I have one physical machine on which four different operating systems are running. And each of these operating system there are four different software that is running. And none of them know the existence of another and all these are running concurrently.

So, in a sense we are giving a virtual machine to each of these software, along a virtual machine in the sense I am giving a hardware and then operating system to that. And all these virtual machines are isolated. So, this is the actual model. Now, when we look at it when we used is as an IAA that is IAAS that is infrastructure as a service, then the hardware is infrastructure, then there is a hypervisor we have to look at the security that. If there is a security lapse everything is gone. So, every the so user 3 in the in your figure the user three can read the data of user 2.

Who is providing that isolation the hypervisor is providing the isolation where hypervisor is compromised everything is gone. Now, the next thing is at a user level again the user has complete access to the server or the operating system and so, submission security need to be provided. So, that the operating system the user does not compromise the user cannot easily get or hack into to the operating system. And come and start doing something on the other virtual machines.

So, this is this is the actual challenge when we look at providing host level security at the IAAS level, when some when people when the people are using your cloud as an infrastructure. So, just summing up, there is an hypervisor level security and that security is very important when you look at multi-tenant architectures where more than one software runs on the same physical hardware. And then there is the next level, which is the customer guest OS or the virtual server and we have to look at security of them.

So, when the customer has an access both to the operating system and the software. So, he what he installs inside the software what he installs how we configures the operating system? Is the operating system hardened enough? All these things are not under the control of the cloud service provider in this infrastructure as a service model. So, that makes the from the cloud service provider is point of view it makes security a larger concern because if he assumes that the customer can be malicious or the customer does not know much about security.

He just installed a very lite unhardened operating system, then somebody attacks from outside into the software or somebody internally attacks into software, then that is going to be a compromise. So, for a cloud service provider, the security becomes a major concern much more larger than a larger concern when the customer is going to use in the infrastructure as a service model.

(Refer Slide Time: 09:59)

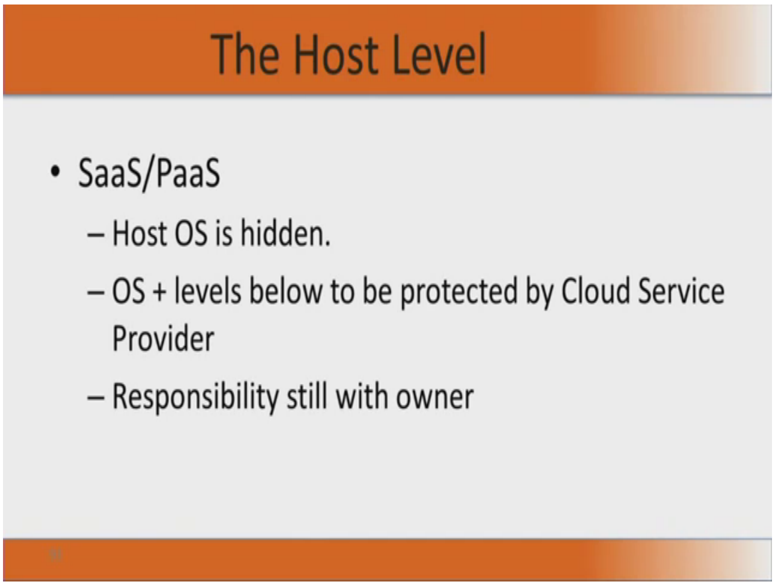

The Host Level

- SaaS/PaaS
  - Host OS is hidden.
  - OS + levels below to be protected by Cloud Service Provider
  - Responsibility still with owner

But if the customer is going to use the cloud service as a software as a service or platform as a service then the cloud provider as much much more control. If a software as a service is absolute control everything is his the software is his, the operating system is his, the guest operating system is his. Of course, the hypervisor and the hardware.

If even as a platform as a service then again the operating system is his, the server is his. So, the only the software is going to come from outside. So, it is easier from the cloud service providers point of view it becomes much more easier to provide security in this SAAS model and PAAS model rather than the IAAS model, but on all these things finally, the owner who is the customer. If he is doing if he is not a just come to the data center with malicious intent of trying to break into other virtual machines; genuinely trying to provide some service to the to his customers. So, let us say a bank bankers it is own customer. Now, this bank is a customer to a cloud service provider. Then from the from the banks customer point of view the bank is ultimately responsible for all these security. So, the bank believes in the cloud service provider to give the security right. So, this is that thing. So, imagine a cloud service provider with one using the cloud as an IAAS and another using a cloud as a PAAS.

So, in the case of IAAS the cloud service providers does not have the responsibility of you know configuring the hard the operating system and the software, but then he always has a fear that that can be very loose or less hardened people can enter if there is a vulnerability. And on the other hand in the same machine you may have another client who is using it as a software as a service and he has to provide absolute security because the customer now that another client believes in the cloud service provider for giving the necessary security. So, this is the major challenge when we look at cloud as an service cloud service from a security point of view.

(Refer Slide Time: 12:41)

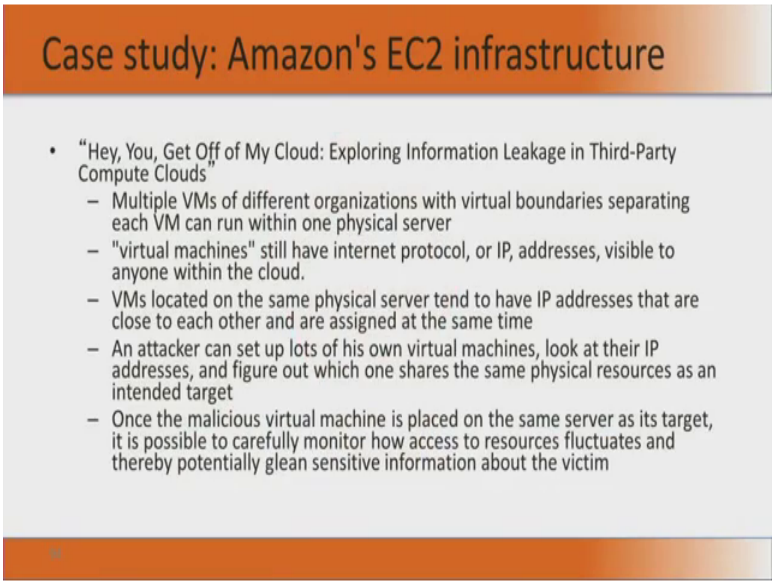

Case study: Amazon's EC2 infrastructure

- "Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds"
  - Multiple VMs of different organizations with virtual boundaries separating each VM can run within one physical server
  - "virtual machines" still have internet protocol, or IP, addresses, visible to anyone within the cloud.
  - VMs located on the same physical server tend to have IP addresses that are close to each other and are assigned at the same time
  - An attacker can set up lots of his own virtual machines, look at their IP addresses, and figure out which one shares the same physical resources as an intended target
  - Once the malicious virtual machine is placed on the same server as its target, it is possible to carefully monitor how access to resources fluctuates and thereby potentially glean sensitive information about the victim

Now, this is one very interesting case study it is called Amazon's a EC2 case study. So, here this is the this is the you can just search google or on hey you get off my cloud exploring information leakage in third party compute clouds. And you will find lot more things here. So, this is a very interesting. So, just some up what is there in the slide you can read the slide, but I will give you a sum up here. So, the ultimate objective of an attacker is to go and; so, this is a multitenant model, is to go and instantiate a virtual machine. So, I am targeting one software say software X of some customer A. This software X is running on one particular physical server as a virtual machine. And that physical server as a multi-tenant model.

From outside I would like to go and instantiate another virtual machine on the same physical server in which X is running and try and start attacking it. So, one of the interesting thing one can do is to keep on spawning many virtual machines into the cloud and then find out an IP address and find out the patterns. If two virtual machines which are installed on the same physical server, every virtual machine will get an IP address they may get the same very close IP address. So, one will be say 192 18 16 14, another will be 192 18 16 15. Like that we would like to do some cartography here and try and find out which exact physical server is the software X is running. Which virtual machine in which X is running. And I would like to instantiate a virtual machine on the same server. So, that I could now go and attack X in a very good physical proximity. So, this is what this case study is all about. So, you can read about this case study in much more detail in the in the website. Just google for that and will get it.

(Refer Slide Time: 15:10)

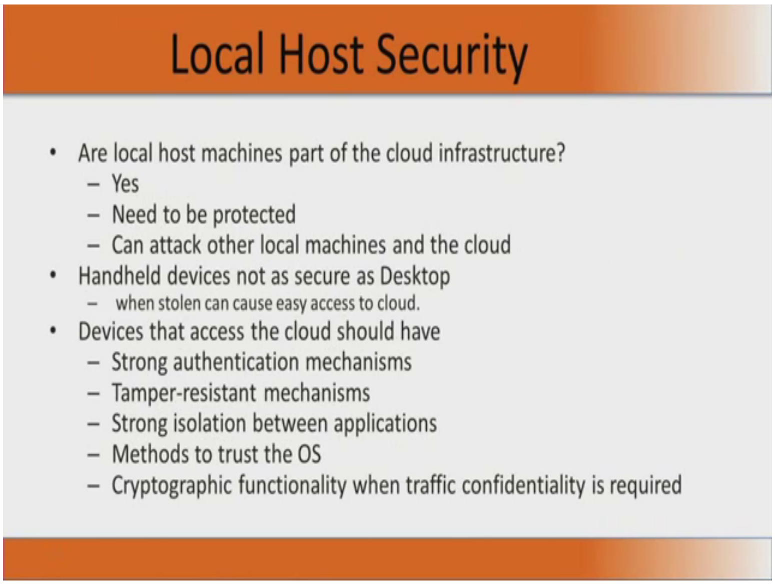

Local Host Security

- Are local host machines part of the cloud infrastructure?
  - Yes
  - Need to be protected
  - Can attack other local machines and the cloud
- Handheld devices not as secure as Desktop
  - when stolen can cause easy access to cloud.
- Devices that access the cloud should have
  - Strong authentication mechanisms
  - Tamper-resistant mechanisms
  - Strong isolation between applications
  - Methods to trust the OS
  - Cryptographic functionality when traffic confidentiality is required

The next thing that we have to look at is of course, local host right. So, many of us are accessing. So, icloud is accessed through ipad for example. So, there are all the clouds are public clouds are access through different machines. Those local machine security are the important or not important. Are they part of the cloud infrastructure, the answer is yes because through them somebody can access the cloud. And do they need to be protected yes they need to be protected. Can they attack other local machines on the cloud, yes they can because electronically there is and there is a channel through this local host into the cloud.

That is enough to go and start exploring or crawling into this floor. So, and very important to note that today desktops, desktops are basically [FL]. There are lot of local host which are connected to the cloud. Are these local host part of the cloud infrastructure? Should they be protected? If they are not protected can somebody attack through those local machines into the cloud and other local machines which are connected to the cloud.

Answers to all this is yes because the local host for example, your ipad which connects to your icloud. All your hand-held devices which I connect; your point of say terminal which connects to some cloud. So, all your machines are access points to your cloud. And if somebody can enter through that into the cloud it paves way and moment they enter into the cloud then the there is a possibility for them to go and attack the cloud, attack the attack the other machines connected to the cloud. And very interestingly you note that hand held machines, hand held devices are not as secure as a desktop. So, it is very very so because the simple thing is when it is stolen and you actually made all your password, you make the handle devices remember the password etctera.

Then obviously, the moment it is stolen the they can get access to your cloud. The person who steals it can get access to your cloud, get access to lot of your credentials. So, this is very very important that your handle devices are kept very safe. Now, so how do we go and handle this. So, when we look at local host today either point of sale terminal or tablets look at secure tablet etcetera we should have some strong authentication mechanisms here. So, when you want to login there should be multi factor authentication like your finger print or NFC card extra.

Then there should be some tamper resistance mechanisms. Nobody can open when your open some secret keys that are used for security should be erased etcetera. And then there should also be a strong isolation at the server end between applications. So, somebody comes through one local machine and as start accessing one application, it should not go and affect another application. And the OS at both the host end and local host end and the cloud should be trusted.

There should be lot of cryptography functionality especially when traffic confidentiality is required. So, all these things can make your cloud very secure. So, to sum up what you have done. So, far on the host. We have looked at the host at the cloud, we have looked at the server hypervisor infra hardware hypervisor operating system and software infrastructure the cloud side. And we are also looked at the local host side which should be accessing the cloud from outside. And we looked at some of the security concerns here. The last two just do here is the application level. What do you mean by an application level security?

(Refer Slide Time: 21:07)

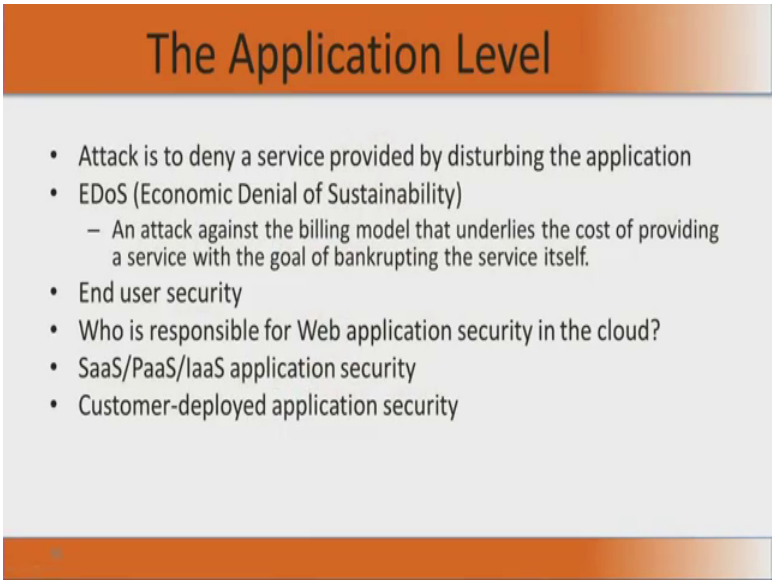

The Application Level

- Attack is to deny a service provided by disturbing the application
- EDoS (Economic Denial of Sustainability)
  - An attack against the billing model that underlies the cost of providing a service with the goal of bankrupting the service itself.
- End user security
- Who is responsible for Web application security in the cloud?
- SaaS/PaaS/IaaS application security
- Customer-deployed application security

The application should be secure enough such that, they through then there should not be a denial of service which will be disturbing the application. So, the application should be available. And then there is also something like an economic denial of sustainability. What happens is an attack against the billing model that underlies the cost of proving a service can go and kill the service itself. So, for example, say I am looking at electronic metering. So, I go and keep resetting the meter value then whole electricity department can go for a loss right.

So, that is anther application. So, I am running applications here and those applications the data they are getting and the data based on which they are making certain decisions all these to be protected. And then of course, the end user security. So, end user logs into your cloud, it is again an application, log in is an application as for as we are concerned and that also should be because should be lot of authentication and protection that is credentials or not stolen.

So, the question now comes is who is responsible for web application security in the cloud? When you are looking at SaaS, PaaS and IAAS. In all these models the application security rest with whom. In SaaS, it can be a cloud is it not with a user or the customer. In PaaS is it with a cloud or with the customer. In IaaS is it with a cloud or the customer. And if you have customer deployed application who is responsible for the security. Now, these are all very important things. I go and put an application running on a cloud or I high in the IaaS model or PaaS model or I run an application which is given by the cloud provider. Then finally, in all these three context who is responsible for that security. So, these are some of the things that we need to understand in great depth.

(Refer Slide Time: 23:23)

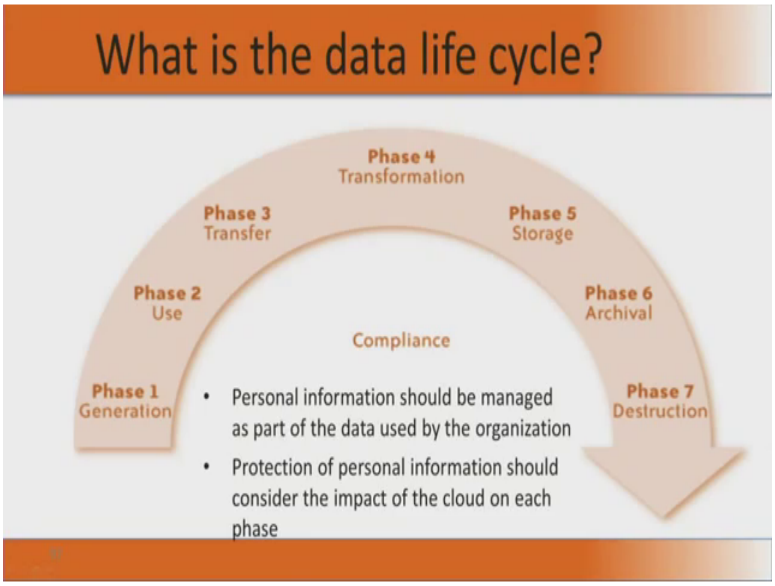

What is the data life cycle?

- Personal information should be managed as part of the data used by the organization
- Protection of personal information should consider the impact of the cloud on each phase

So, for us to appreciate application level security we may have to go and look at the data life cycle in great detail where it will give us pointer about security at the application level. That we will see in the next session.