Introduction to Cloud Computing

Table of Contents

Introduction to Cloud Computing

Introduction to Cloud Computing

**001 Instructor: In this section, we will introduce the concept of Cloud Computing.

Definitions

**002 First, it's important to define cloud computing, because many vendors and many people use the phrase "cloud computing," and it's not always clear that cloud computing is what they think cloud computing is. Basically, we are storing and accessing data and resources remotely instead of locally. Typically, a cloud provider will provide some kind of service in some kind of metered way. And then a consumer will purchase that server -- server or data or whatever solution is being offered as the cloud, in some kind of measured way.

So they may pay per gigabyte. They may pay for the network bandwidth. Whatever the agreed-on rate is between the consumer and the provider, is what will be paid. So

we're sharing our resources for economy of scale that allows providers to spin up vast, bulk resources of processors, of hard drive space, of whatever resources that could be, whereas consumers, be it individuals or small businesses, can purchase a section of those resources for use without having to spend the resources and time to spin up their own solutions.

Cloud Characteristics

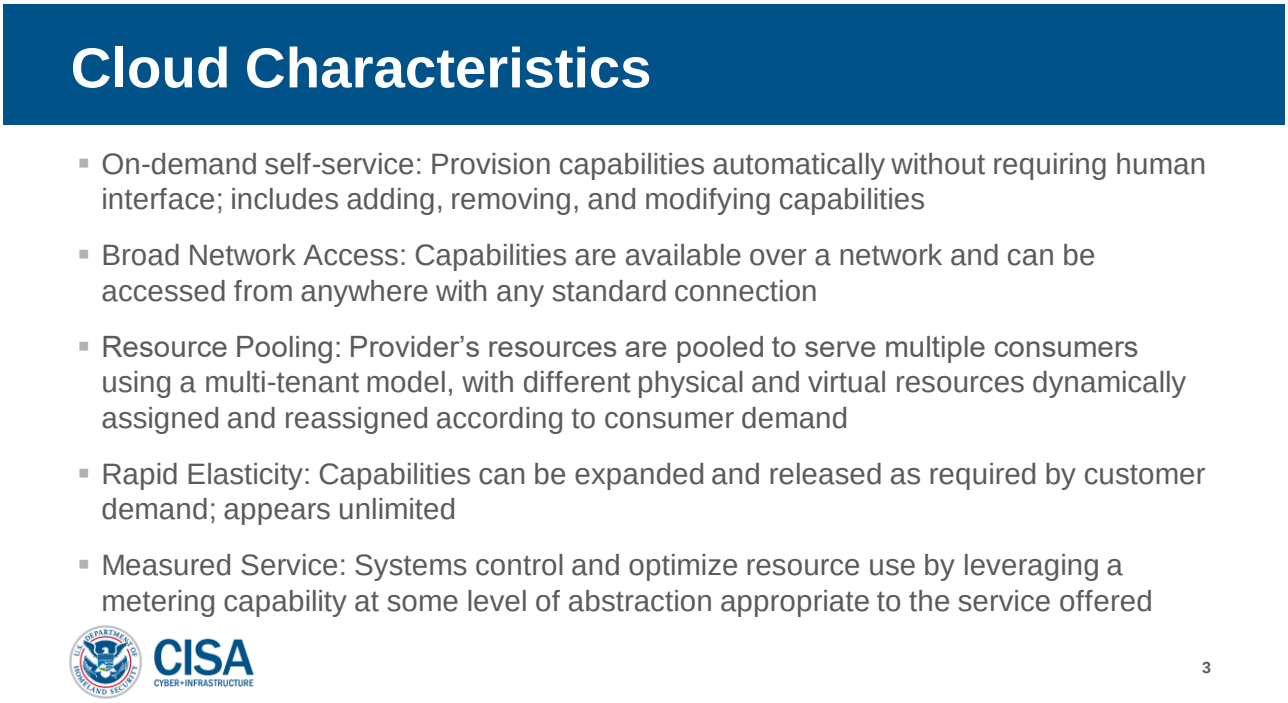

**003 The cloud has five basic characteristics including on-demand self-service. So users should be able to provision cloud resources on demand. They shouldn't have to go through a lengthy wait process to get additional resources or to return resources that they are not using. There should be broad network

access for the cloud. So in other words, the cloud has to be available wherever network access can be found with any kind of standard connection. So a cloud should not have to require a super expensive connection just to be able to connect. It should be available over standard residential connections. Resource pooling. And this is where the multi-tenant model comes into play because more than one tenant, much like an apartment building, will share the resources.

So everyone will share the heating bill. Everyone will share the air conditioning bill, the water bill. It becomes more like a group resource, rather than an individual resource. Rapid elasticity. That means that the capabilities that a consumer needs can be expanded and released as needed and on-demand. From the consumer side of things that resource pool should appear to be completely unlimited. From the cloud provider it will not be unlimited. There is a finite amount of resources but it will appear to the consumer to be completely unlimited. And finally, a cloud solution will be a measured service. So there will be some way to monitor and bill that consumer. Much like having a metered water bill or a metered electric or gas bill, same thing with the cloud.

Cloud Benefits

Cloud Benefits

- Additional agility to increase and decrease resources as required
  - Spin up more networking resources to combat a DDoS
  - Don’t have to purchase resources for 100% peak-load
- Cheaper infrastructure cost via resource sharing
  - Cheaper cost may come with less available features and less customization options
- Multi-user document collaboration and document versioning
- Infrastructure maintenance is streamlined
  - Applications not installed on workstations
  - Backups and failovers are centrally managed
- Physical space requirements are reduced

**004 There are some benefits to going to the cloud interface, as opposed to a more traditional interface. We have additional agility because we can increase and decrease resources as required. Small businesses won't have to purchase resources for 100% peak load when they don't use 100% peak load all the time. So, in other words, that small business will pay for what they actually use rather than paying to have a bunch of resources sitting there that aren't even being used. So that leads to cheaper infrastructure because we start sharing those resources. That cheaper infrastructure does come with some less available features because it's not as customizable because we're sharing it among multiple corporations. We're sharing

it among multiple small businesses and multiple individuals. The cloud also can offer multi-user document collaboration, as well as document versioning.

Much like SharePoint, multiple people can collaborate on the same document at once and all those changes are tracked and controlled through Microsoft Solution. The infrastructure maintenance is streamlined. Because there's not as much physical hardware space, there's not as much infrastructure maintenance to perform. The applications, those aren't stored on workstations. Those are stored on the servers. All the backups and the failovers, again those aren't stored on user workstations. They're stored on the cloud servers. So that's easier to centrally control and centrally manage, which also leads to a reduction in space requirements for the physical servers themselves.

Cloud Disadvantages

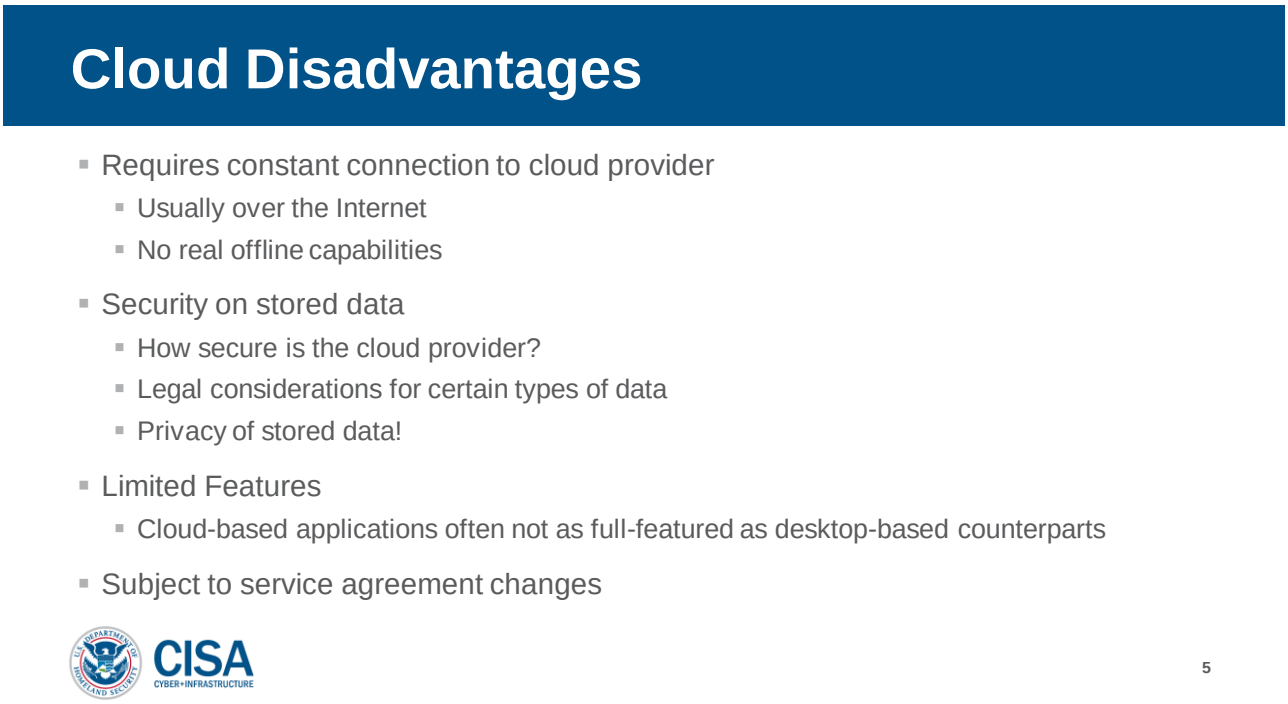

**005 But the other side of the coin is some disadvantages because the cloud does require a constant connection. There's no real offline capability here. You have to be connected to the cloud in order to use the resources in the cloud. Then there's the question of the data that is actually stored in the cloud. That data, how secure is it? Because even though a particular organization can still own the data it is being completely controlled by some kind of third party. So how secure is that data? What's the legal considerations for that data? If that data is protected health information, is that data still meeting HIPAA requirements by being stored in the cloud? There can be privacy issues. And again, we're sharing those resources among many different

tenants, much like an apartment building.

So if someone breaks into Apartment 2B, does that mean they're also breaking into Apartment 2C? Those are questions that are still being explored. And sometimes, a cloud-based application will have limited functionality and limited features, as opposed to a full desktop version. So it'll be more streamlined, more stripped down. It won't have all the bells and whistles as a desktop application. And finally, especially if a third-party cloud provider is used, the service provided is always subject to changes as defined by that cloud provider.

Conventional and Cloud Computing

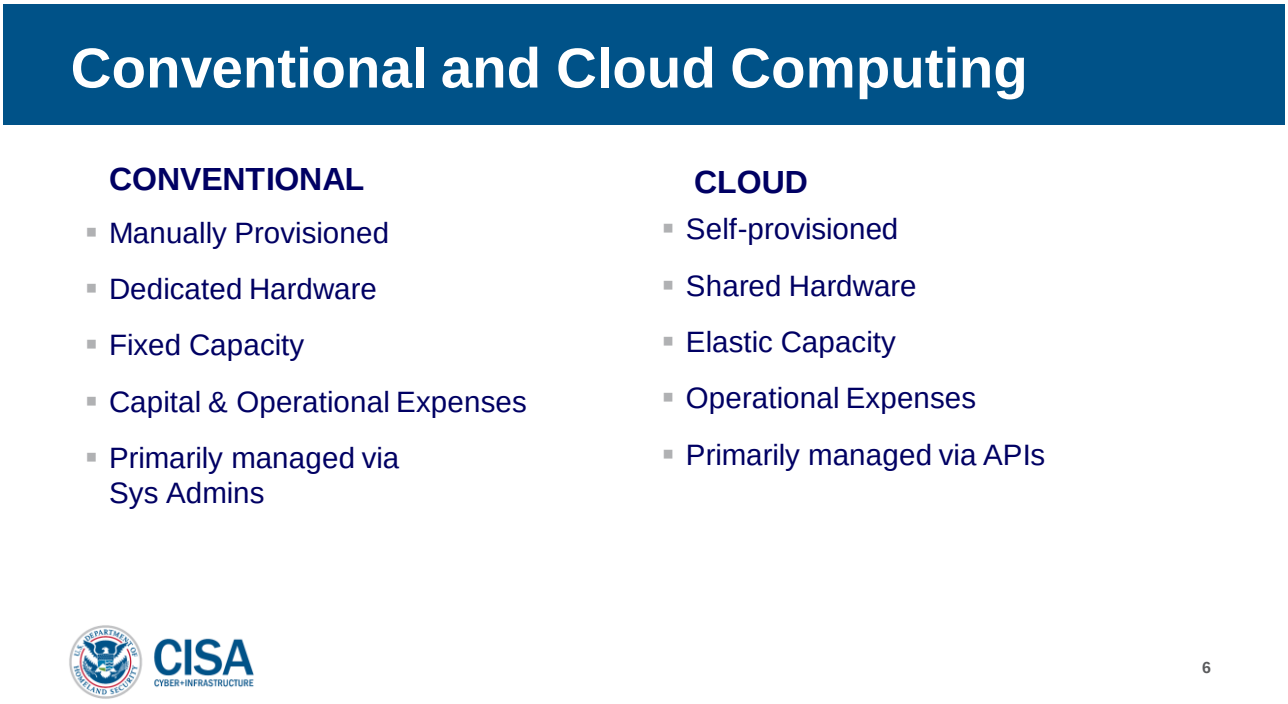

**006 Here is a more basic

comparison of conventional versus cloud computing. Sometimes the conventional way is the way to go because if you have any data security concerns or if you have specific legal ramifications surrounding the data, such as that HIPAA protected information, or maybe it's government classified information, then maybe the more conventional approach is still the correct approach, whereas the cloud, we have that shared hardware. We have that elastic capability. The management is slightly different. System administrators versus APIs. So not different or not better, not worse, just different.

Virtualization

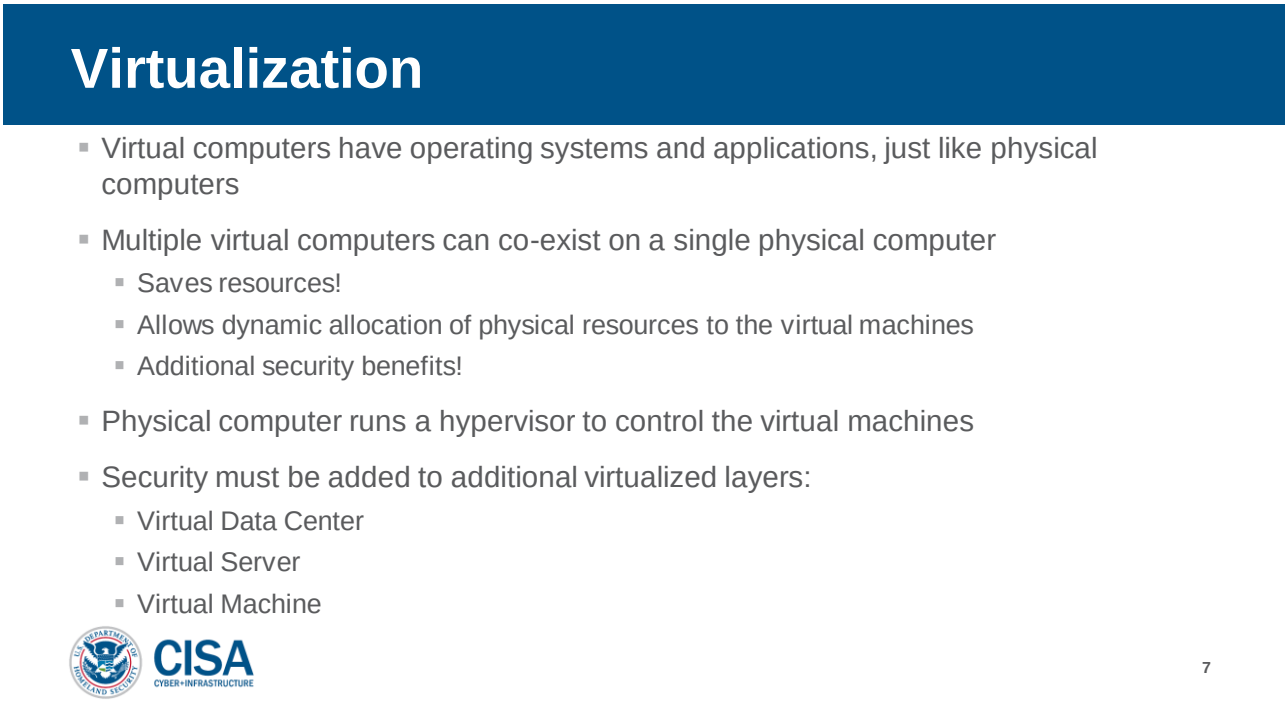

**007 The cloud relies on the concept of virtualization. And virtual

computers have operating systems and applications just like physical computers. The only catch is that there can be multiple, virtual computers on a shared, single, physical host. And that's where we get that shared resource because that one physical computer can have many different virtual computers. And those virtual computers can be sandboxed out so that they cannot communicate with each other except through standard network protocols. So this does have some additional security benefits because that shared physical space is reduced. So the physical computer is required to run a hypervisor to control the virtual machines. So that way the physical computer can power on the virtual machines, take snapshots, power them off, adjust the resources required.

However, with this virtualization we have to add in security to all of the virtualized layers. And that depends on how much of the infrastructure has been virtualized. It could just be the virtual machines. The virtual machines will have additional security that is required to address the fact that they are virtual machines, as well as a virtual server, as well as a virtual data center, can all be contained in a single, physical host.

Virtualization Models

- Type 1:
  - Runs directly on top of hardware with no other operating systems involved
  - “Bare metal hypervisors”
  - Control hardware and manage guest operating systems
  - Oracle VM Server, Citrix XenServer, Vmware ESX
- Type 2:
  - Runs as a process within an existing operating system
  - Abstracts guest operating system from the host
  - Vmware Workstation, Vmware Player, VirtualBox

**008 There are two basic virtualization models. The first type is more an enterprise level model. That's where a hypervisor runs directly on the physical hardware. Like Oracle's VM server. Like Vmware's ESXI servers.
The hypervisor controls the hardware directly and manages the guest operating systems. And as far as those guest operating systems are concerned, they are operating directly on that physical hardware.
The second type, Type 2, runs as a process within an existing operating system. So there's the layer of physical hardware and then an operating system and then the hypervisor and then the guest operating systems.

Virtualization Models

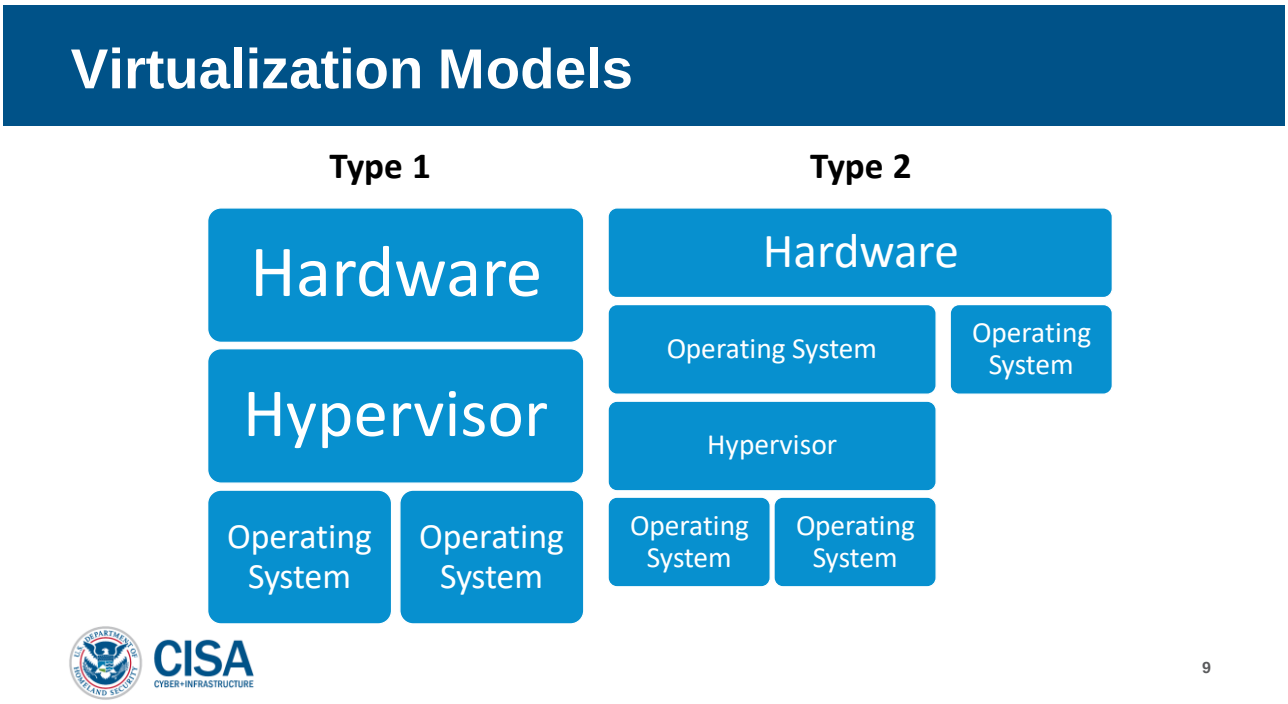

**009 To put this in a diagram, here's the Type 1 and this is going to be your more enterprise level model. So you have the bare hardware followed by the hypervisor followed by the operating system. Whereas with Type 2 we have the hardware, we have the operating system. We could add multiple operating systems. It depends on how the boot setup is ordered. And then on top of the operating system we have the hypervisor and then we have the additional guest operating systems. Again, not better, not worse, just different. It depends on how the resources need to send.

Service Models

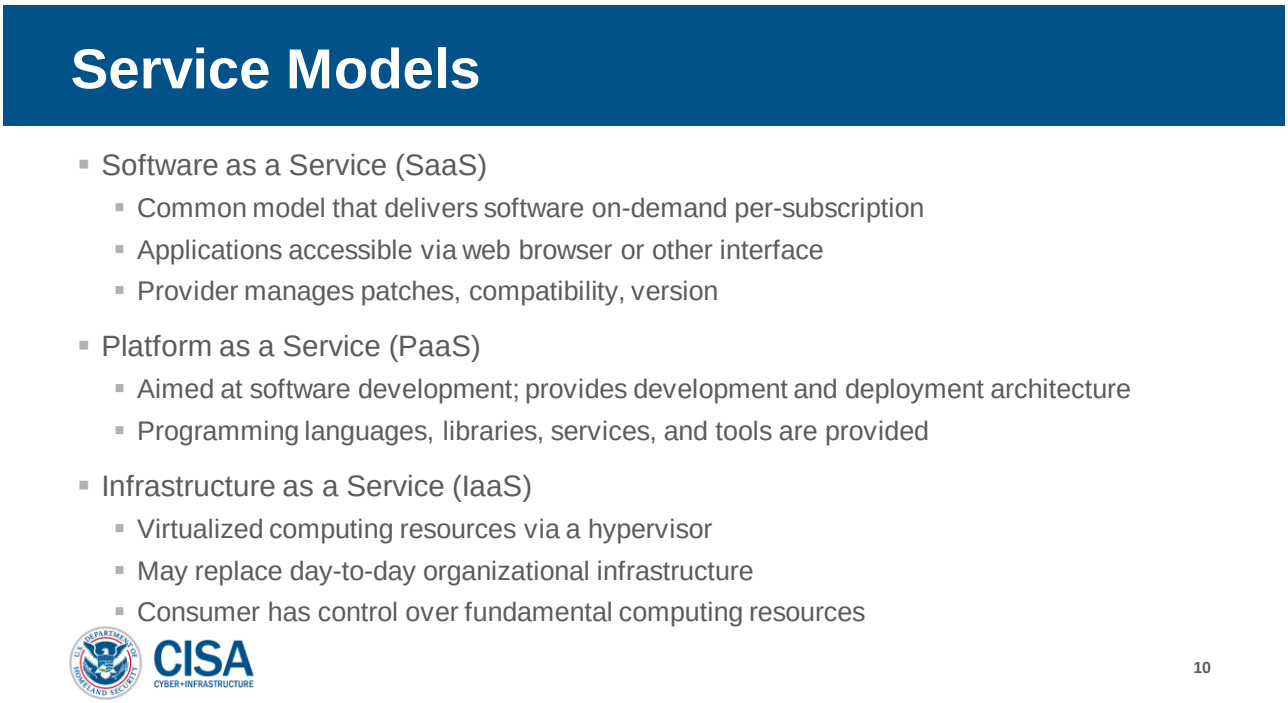

**010 There are three common service models. The first is Software as a Service. And this model is extremely common. This is where all of your web-based mail applications come from. That is software as a service. This is where Microsoft Office 365 comes from if it's software as a service. So you pay for a subscription to a particular piece of software and then you have access to that software's functionality. So the provider, be it Google or Yahoo or Microsoft, manages all the patches of that software. They manage the compatibility, the versioning. It will get updated according to that provider. Here, the application is not installed on the user's desktop itself. You have to access that application through a web browser. Next model is Platform as a Service. And this is

really aimed at software development. So the cloud provider will provide the infrastructure, the operating system, and the development environment. And will essentially rent that out to any organization or any individual that wants that service.

So the PaaS provider will control the programming languages available, the libraries, the services, and all of the tools will be provided as part of that cloud infrastructure. The last model is Infrastructure as a Service. And this is where the virtualized computing resources are completely controlled through a hypervisor. And this is where it can replace day-to-day operational infrastructure. So instead of standing up your own network or your own organizational network, you can rent a cloud-based infrastructure as a service and you have all of your servers, you have all of your network infrastructure, and then you control the webserver. You control the database server. But it's actually physically sitting at the cloud provider. All the cloud provider does is control the hypervisor and controls the bare metal, physical hardware.

Service Model Responsibilities

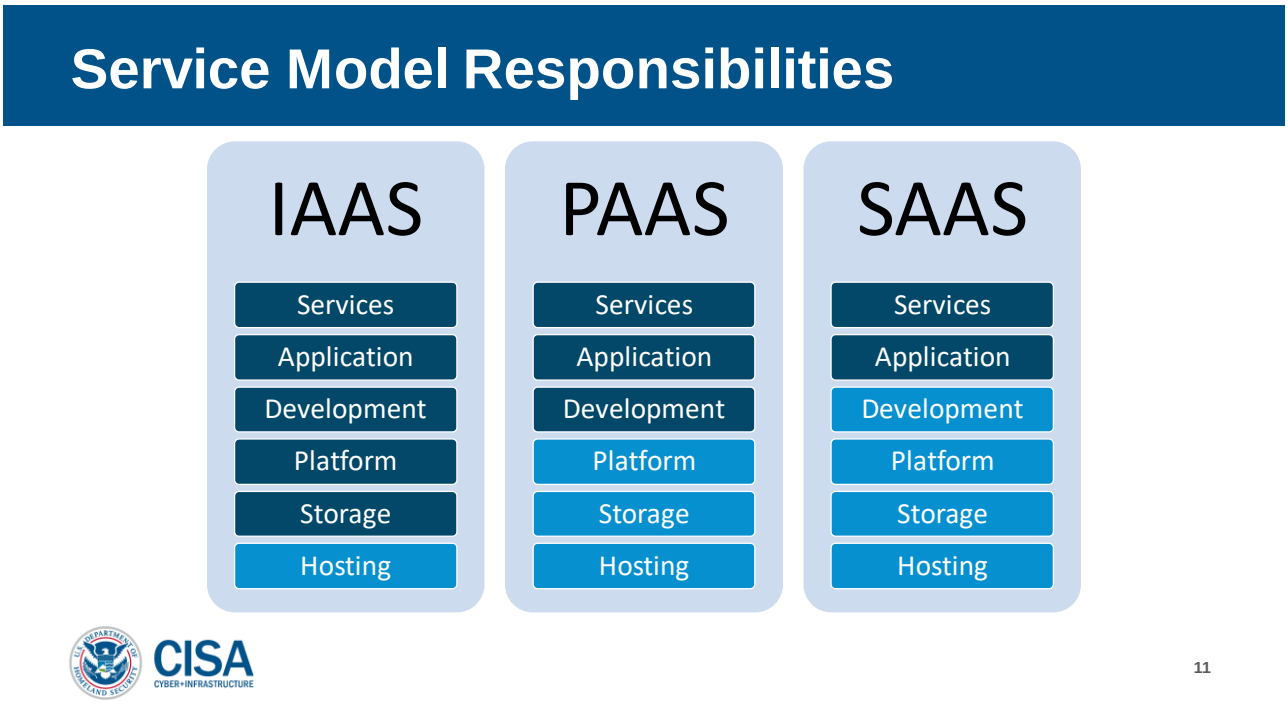

**011 So here in dark blue, is what the cloud consumer will have control over. So for infrastructure as a service, the cloud provider controls the hosting. The cloud consumer controls the storage, the platform, the development, the application, and the services. This model is most intense for the cloud consumer because the cloud provider doesn't really control much. Just the hosting. Versus platform as a service, the cloud consumer will have control over the development, application, and services on a particular server but will not have any control over the platform, the storage, or the hosting. All of that will be provided to the consumer from the provider. And then finally, with software as a service, all the consumer will be able to control are the services and the

application. Everything will be -- everything else is controlled by the cloud provider.

Software as a Service

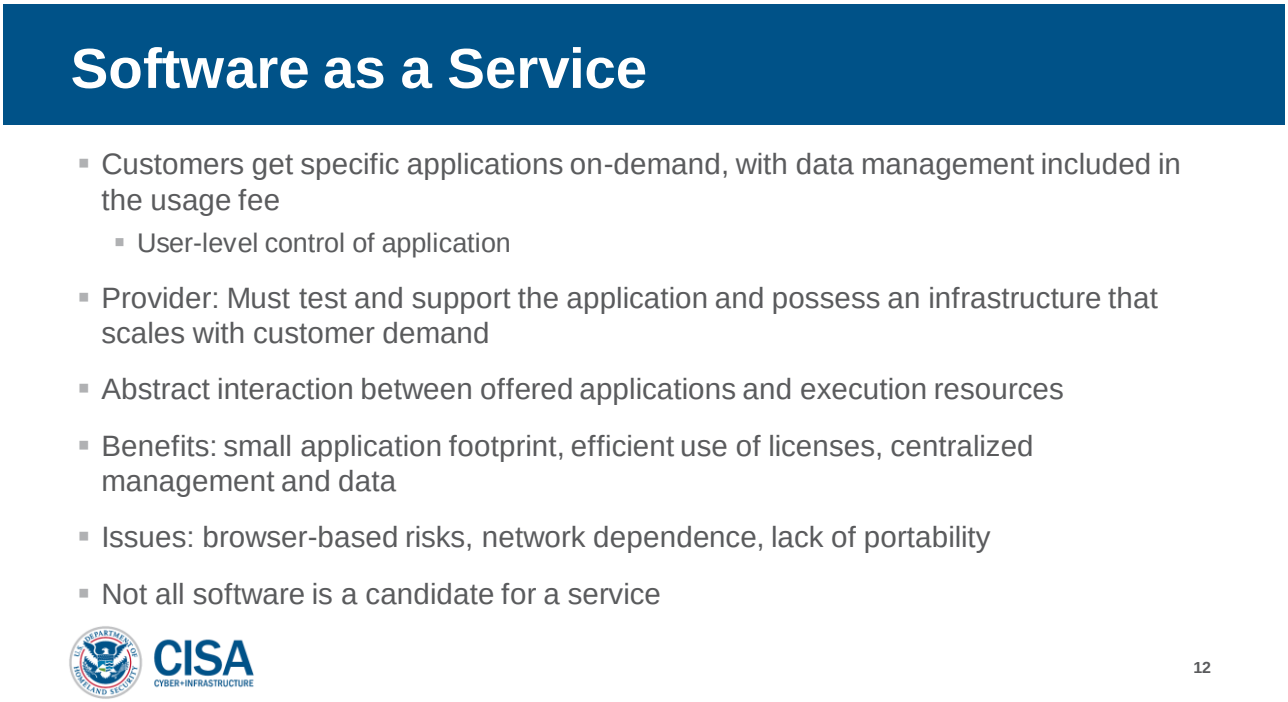

**012 So to go into a little bit more detail about each model. With software as a server, the customers get an on-demand application like a word processor, like a spreadsheet. So, the data management, all of the files, are typically included in those usage fees. So users can create and modify and update files to their heart's content based on their usage fees. So they pay their monthly subscription or their yearly subscription fee. And then they get that storage. They get that software and they're able to store those files. So the provider has to test and support that application. Anytime

a customer has an issue, the provider has to support that application. So if the customer's files completely go missing, it's on the provider to perform some kind of customer service to replace or otherwise restore those files. And it's also on the provider to scale their infrastructure to meet customer demand.

So the benefits here is that the consumer does not have to install software on their machine. They can use licenses very efficiently. If we talk a small business, maybe the business only has the money to purchase 12 licenses for a particular piece of software. But they have way more than 12 employees. Maybe they have 100 employees and they all work different shifts. But essentially, they can work and share those licenses across those shifts based on the software of the service model. So you end up purchasing less software licenses under this model. However, there are some issues, such as network dependence. You have to be connected to a network in order for -- to use this service. You can't just call up a file on your local desktop and work on that file with no network connection. And there are always browser-based risks. You can have Man-in-the-middle attacks. You can have Man-in-the-Browser attacks. Anything that can attack a browser will also affect the software as a service. And finally, not every piece of software is a candidate to be a service. Some software packages are more intense

than others and require more resources than others. So not everything makes sense to offer as a service.

Platform as a Service

**013 Next up, we have Platform as a Service. And this is where, like I said, customers get the tools and the execution resources to develop and test and deploy applications. So this could be part of a software as a service cloud, but consumers have a little bit more control and they're able to develop their own applications. They're not dependent on the applications provided to them by the provider. So the providers have to maintain that application inventory. They have to maintain the development tools. They have to patch the development

environments. And they have to patch the execution environments so that they can be used by the customers.

Infrastructure as a Service

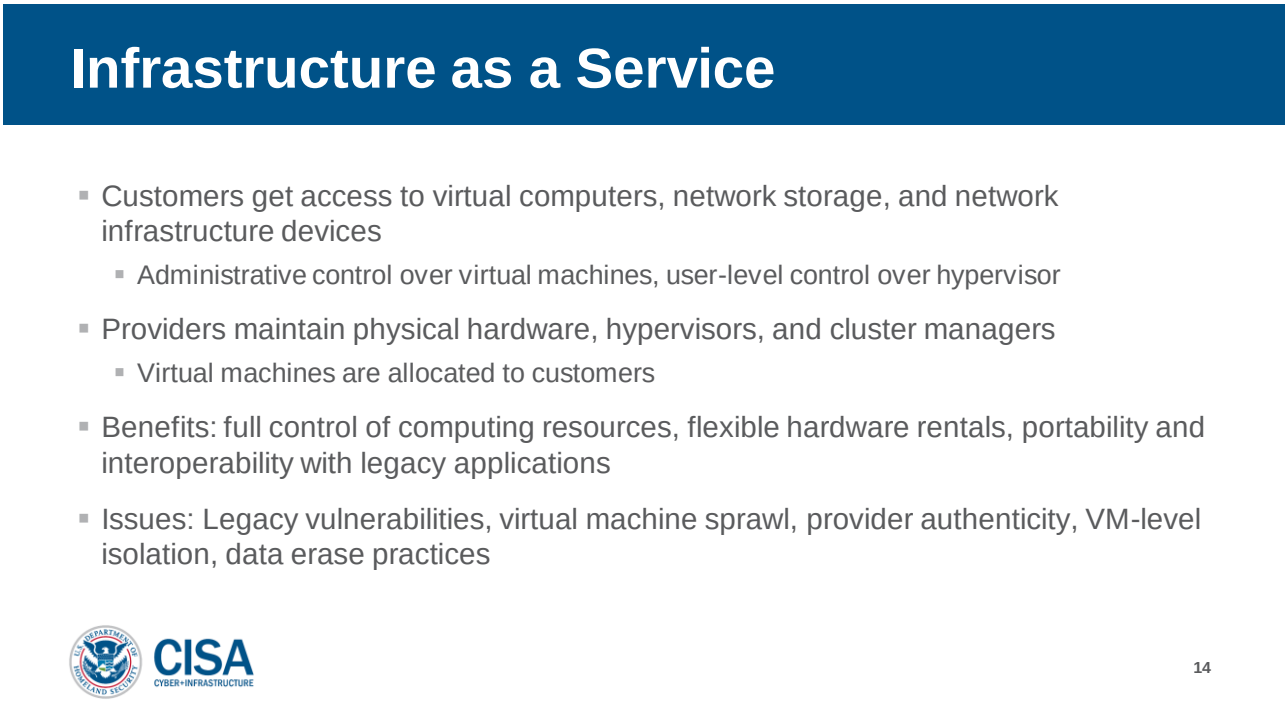

**014 Then we have Infrastructure as a Service. And this is where consumers have to do the most work because they get access to the virtual computers, the virtual network storage, and the virtual infrastructure devices. All the configuration of those items is completely left up to the consumer. The provider maintains control over the physical hardware and maintains control over how the resources are assigned but once the resources are assigned to the customer, it's under the customer's control and so how to actually configure and utilize those resources.

So under infrastructure as a service, the consumer has full control over those computing resources, only they have flexible hardware rentals.

So when they first spin up a web server, if you only need say 20 gigabytes of memory and 4 gigabytes of RAM, then that can be what you purchase. But in six months when a business grows and you need say a terabyte of memory and 16 gigabytes of RAM, instead of having to go out and purchase that and upgrade the infrastructure, you can just rent more hardware. So you have more control there and you have that hardware flexibility and you can scale your hardware as your business scales. There are some issues here primarily considering data and data erasure practices. Again, what kind of laws and regulations are governing a particular business's data? With HIPAA or with financial information, there may be strict rules and regulations over how the data is stored and what must happen to a hard drive once a particular brand of data has been stored on it. Some laws require that a hard drive be destroyed when the data is destroyed so that the data can be no longer accessible. So you have to rely on the hardware provider, on the cloud provider, to help you as the business to perform those data erasure practices.

Notices

Notices

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM-0004104

**015 Thank you.