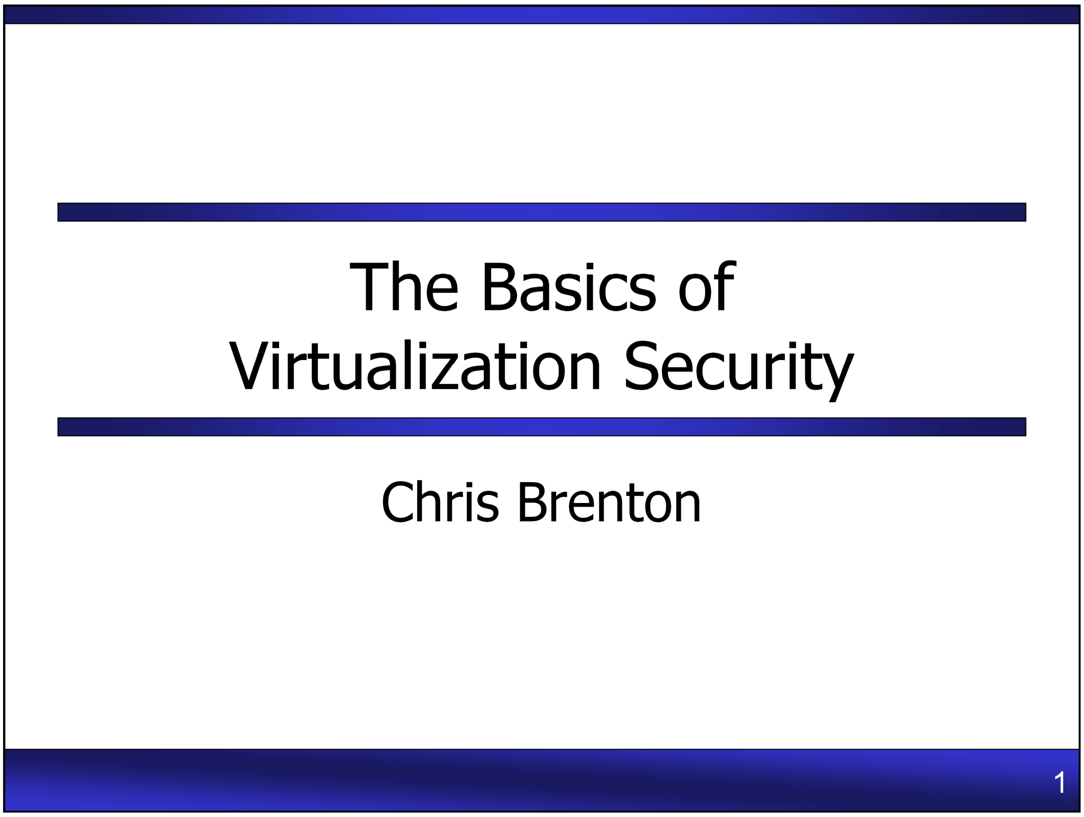

The Basics of Virtualization Security

Chris Brenton

1

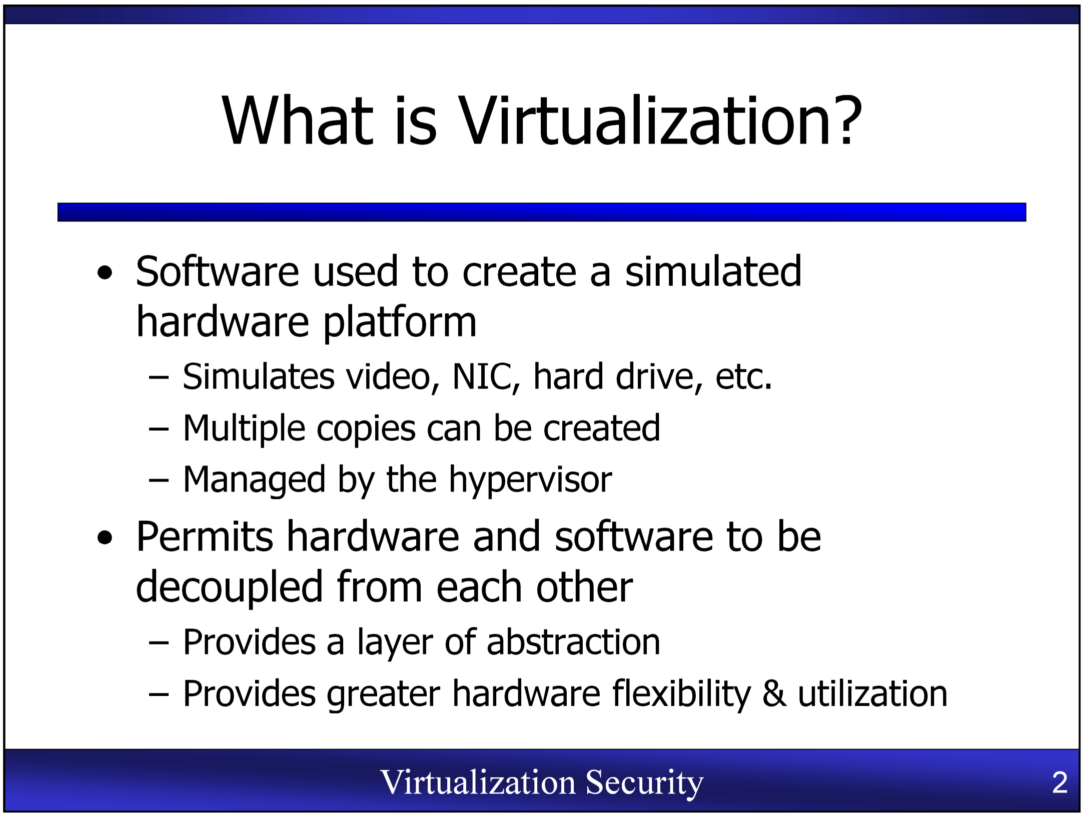

What is Virtualization?

- Software used to create a simulated hardware platform
  - Simulates video, NIC, hard drive, etc.
  - Multiple copies can be created
  - Managed by the hypervisor
- Permits hardware and software to be decoupled from each other
  - Provides a layer of abstraction
  - Provides greater hardware flexibility & utilization

Before we can talk about virtualization security, we need to delineate the differences between the terms “virtualization” and “cloud”.

Virtualization, at its core, is the ability to emulate hardware via software. If we walk through the processes, some form of operating system still needs to be booted from the hardware. This may be a full blow OS such as Linux or Windows, or it may be a stripped down OS specifically designed to provide virtualization, such as VMware’s ESXi (which is a stripped down Linux operating system). In either case an operating system is first booted and then an emulation software stack is loaded which is referred to as “a hypervisor”.

The hypervisor is the component which is responsible for emulating specific hardware configurations to guest operating systems. When a guest is loaded into a virtual machine (VM), the hardware that gets detected is the simulated hardware via the hypervisor, not the actual hardware itself. The guest OS is abstracted from the true hardware, adding a component of versatility. The hypervisor is capable of creating multiple simulated environments, or multiple VMs, which permits us to run multiple operating systems that may have slightly different hardware requirements.

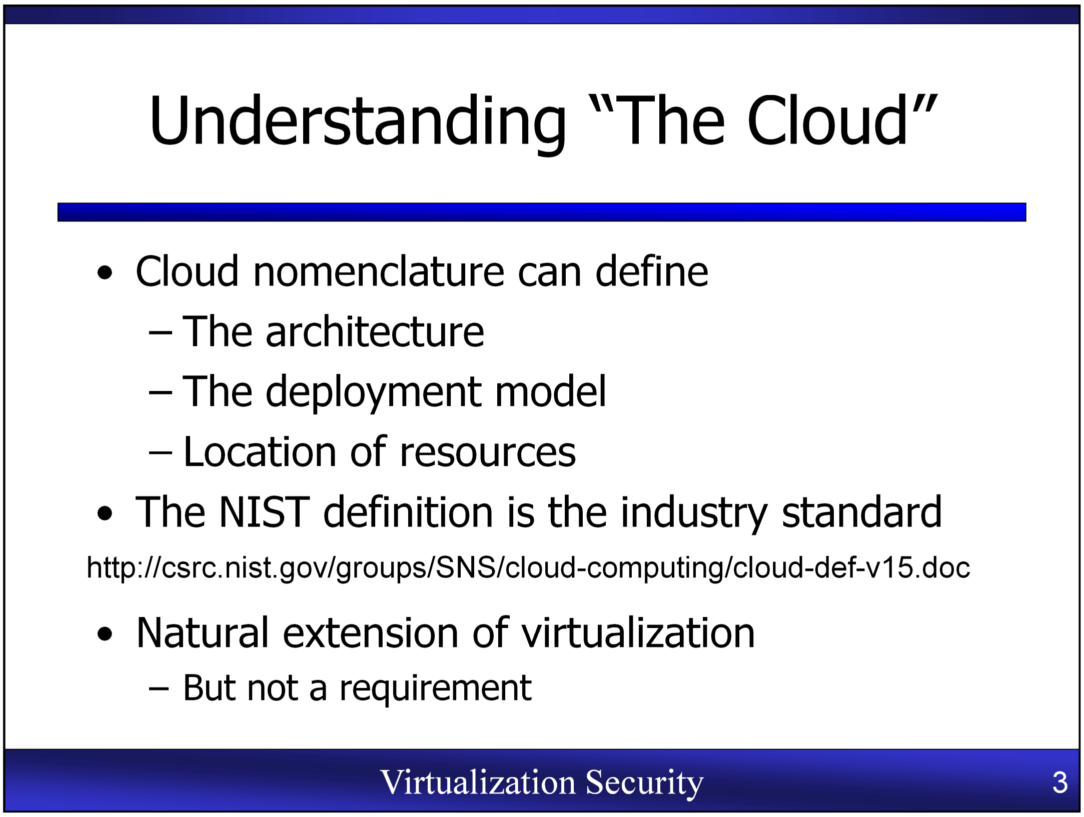

Understanding “The Cloud”

- Cloud nomenclature can define
  - The architecture
  - The deployment model
  - Location of resources
- The NIST definition is the industry standard

http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc

- Natural extension of virtualization
  - But not a requirement

Clearly, the term “cloud” has become so general purpose, it gives very little insight into the object being described. While NIST has an industry accepted definition of what cloud computing actually includes, this does not stop vendors from generating a steady stream of “aaS” offerings to try and add some sex appeal to their products. For the purposes of this course, we will stick with the terminology defined by NIST.

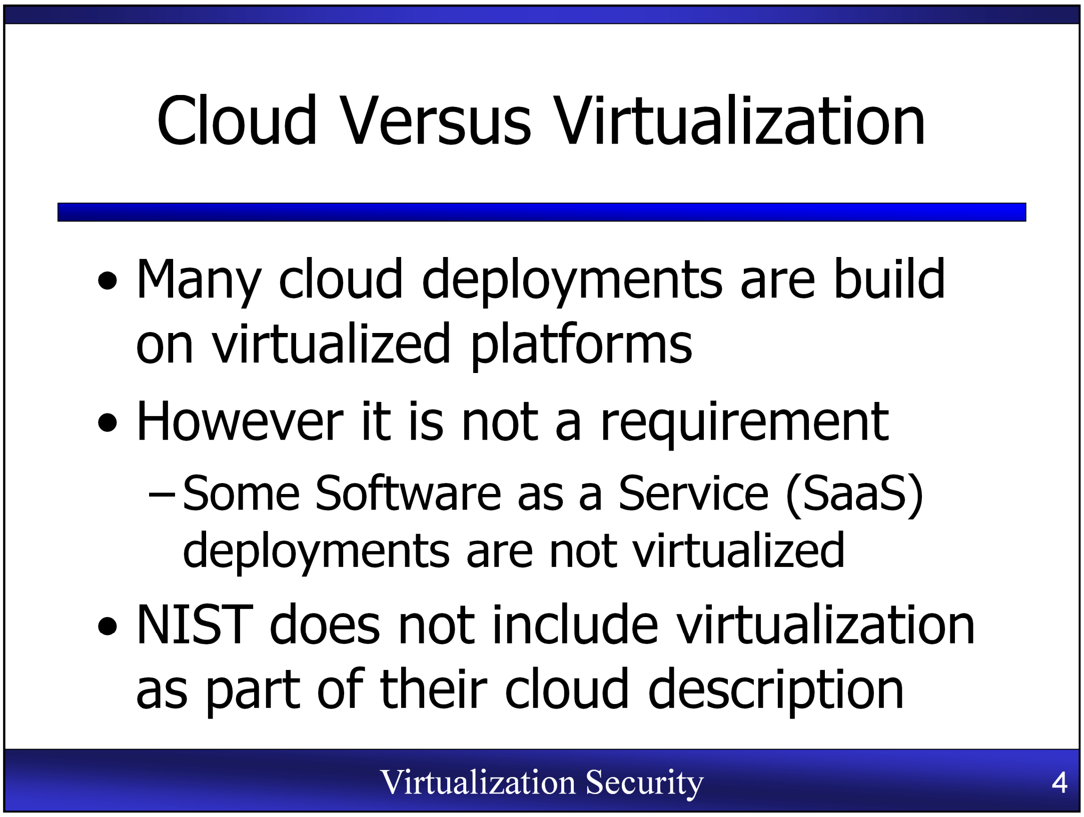

Cloud Versus Virtualization

- Many cloud deployments are build on virtualized platforms
- However it is not a requirement
  - Some Software as a Service (SaaS) deployments are not virtualized
- NIST does not include virtualization as part of their cloud description

Virtualization Security 4

Many cloud deployments include a virtualization component. While this is a common technique, it is not a strict requirement. For example the NIST cloud definition does not call out virtualization as a required component. There are a small subset of cloud deployments, mostly Software as a Service based, that do not rely on virtualization.

To help keep these two terms straight, consider the analogy of building a residential home. “Cloud” is analogous to the house itself, while virtualization would be the equivalent of a full basement. Most houses are built on full basements as they provide excellent structure and expansion capability. A full basement is not a requirement however, as houses can certainly be built on half basements, slabs or even stilts. So while we can say most houses (clouds) are built on full basements (virtualization), they are certainly not a requirement.

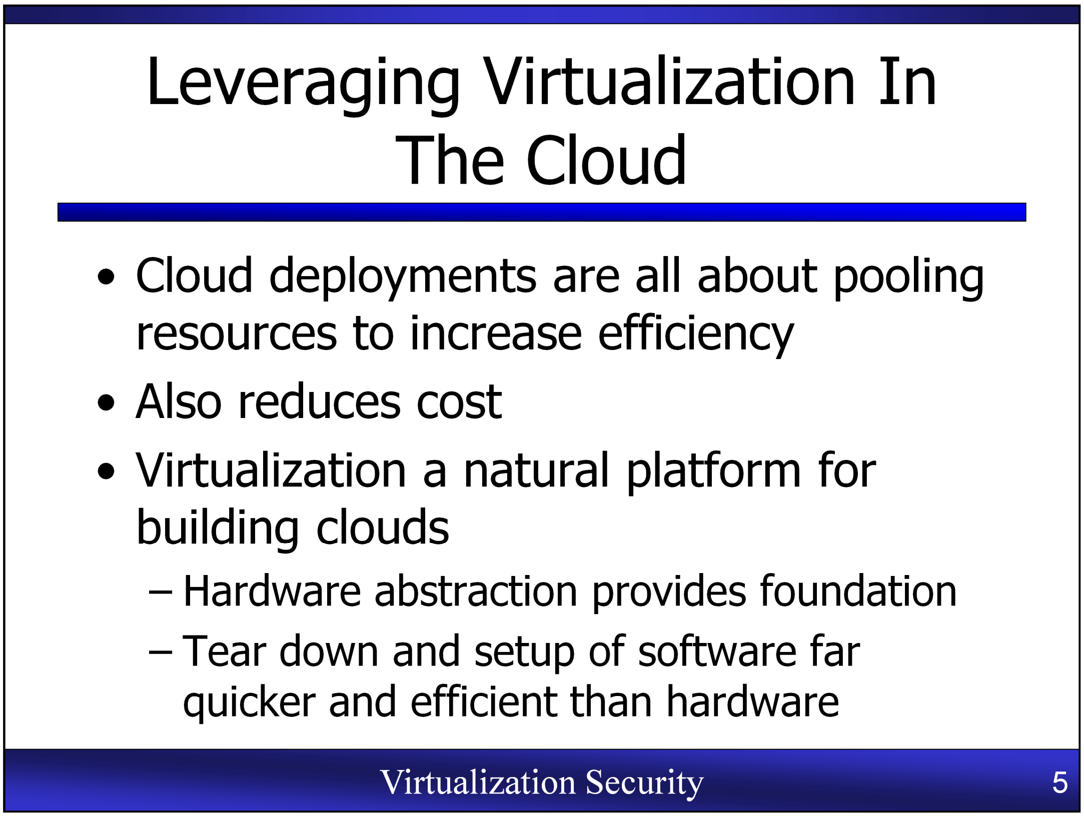

Leveraging Virtualization In The Cloud

- Cloud deployments are all about pooling resources to increase efficiency
- Also reduces cost
- Virtualization a natural platform for building clouds
  - Hardware abstraction provides foundation
  - Tear down and setup of software far quicker and efficient than hardware

So while virtualization is not a requirement of cloud computing, its ability to efficiently share resources makes it an excellent foundation.

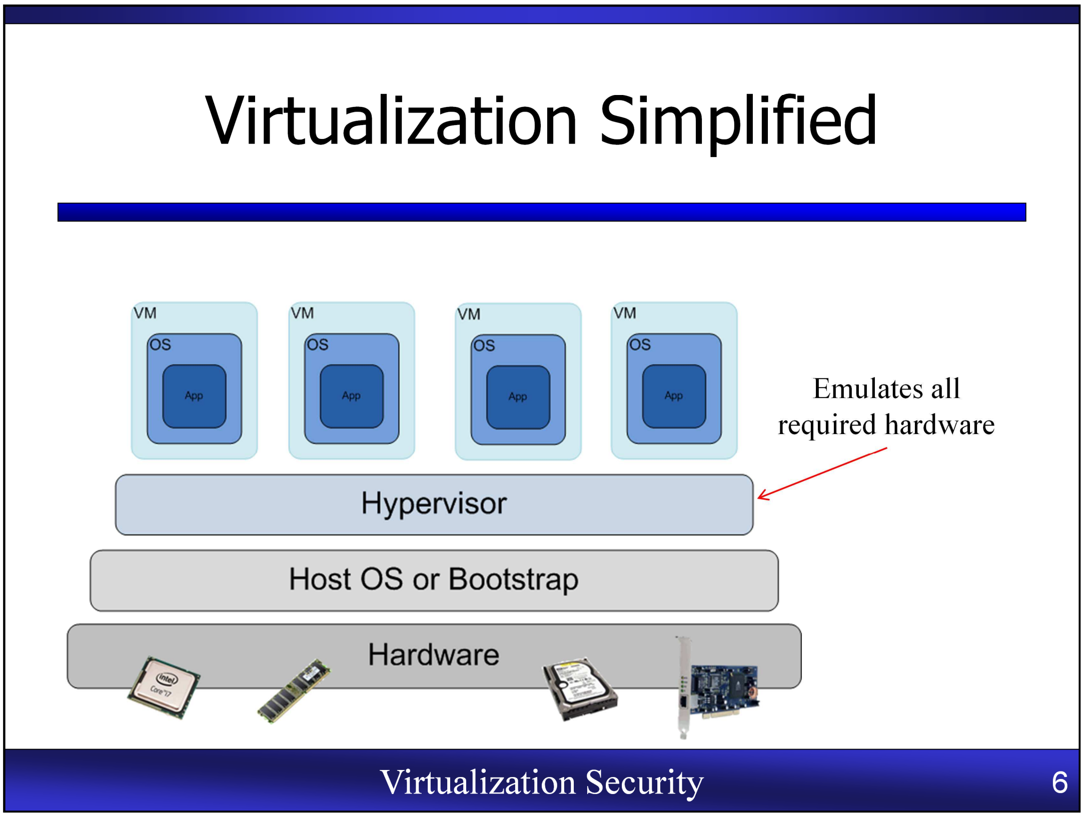

Virtualization Simplified

VM

OS

App

VM

OS

App

VM

OS

App

VM

OS

App

Emulates all required hardware

Hypervisor

Host OS or Bootstrap

Hardware

Virtualization Security 6

The above diagram graphically represents the layout of a virtualized platform. The hypervisor abstracts the VMs from the actual hardware by emulating these components. This can create a potential security bridge between the VMs, which we will discuss in greater detail later.

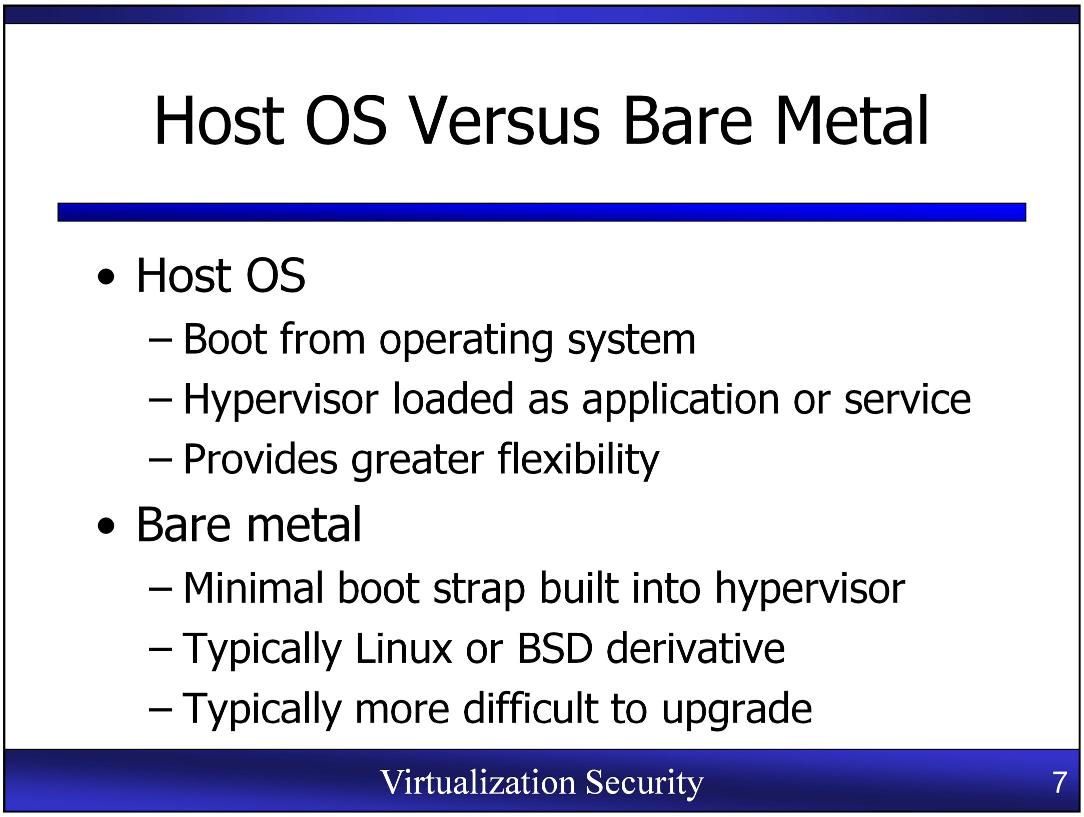

Host OS Versus Bare Metal

- Host OS
  - Boot from operating system
  - Hypervisor loaded as application or service
  - Provides greater flexibility
- Bare metal
  - Minimal boot strap built into hypervisor
  - Typically Linux or BSD derivative
  - Typically more difficult to upgrade

Virtualization Security 7

Virtualization is available in two different flavors, host OS based or bare metal. When virtualization is run on a host OS system, it is run like any other application. This permits the administrator to level any tools that are capable of running on the host OS while managing the environment. On the down side, the host OS increases the amount of code being executed on the system, and thus increases the surface area of potential attacks.

A bare metal system still uses an operating system, but the OS has been stripped down to only support the virtualized environment. This reduces the number of available tools, but also decrease the amount of code that can be potentially exploited. It can also be argued that bare metal system can be more difficult to patch and upgrade.

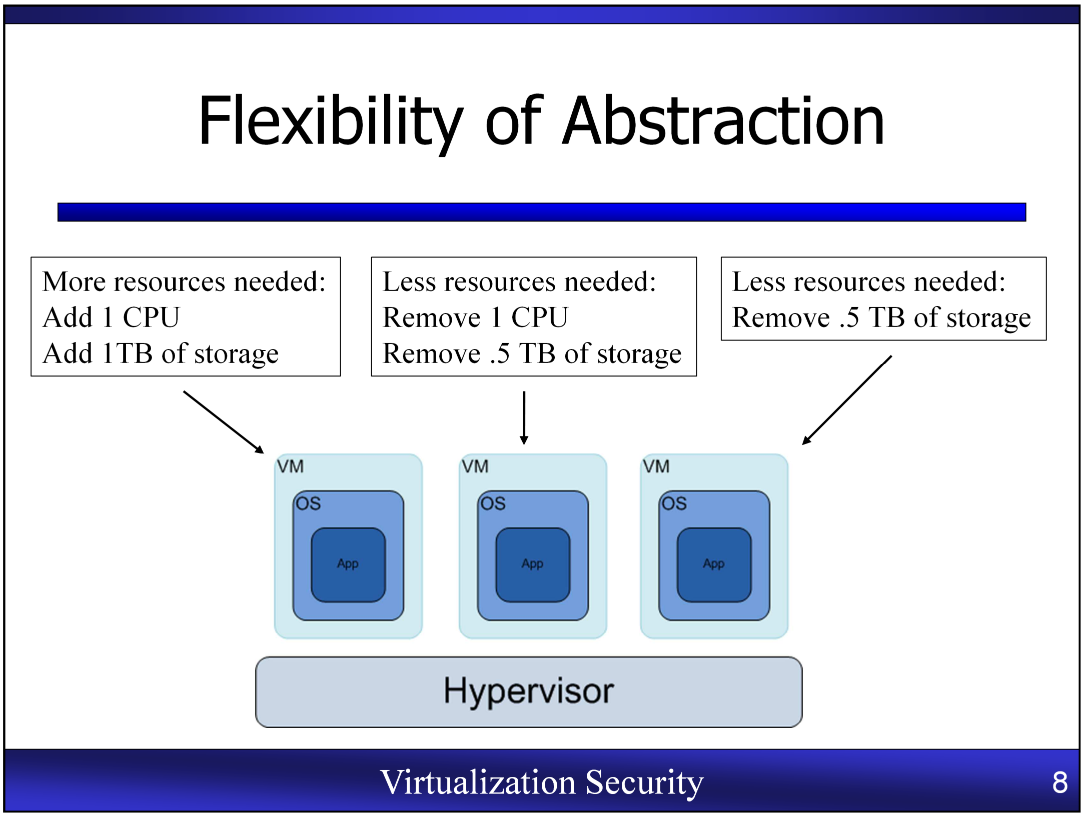

Flexibility of Abstraction

More resources needed:
Add 1 CPU
Add 1TB of storage

Less resources needed:
Remove 1 CPU
Remove .5 TB of storage

Less resources needed:
Remove .5 TB of storage

Virtualization Security 8

One of the benefits of virtualization is that system resources can be re-allocated as needed. This permits the administrator to better optimize the environment.

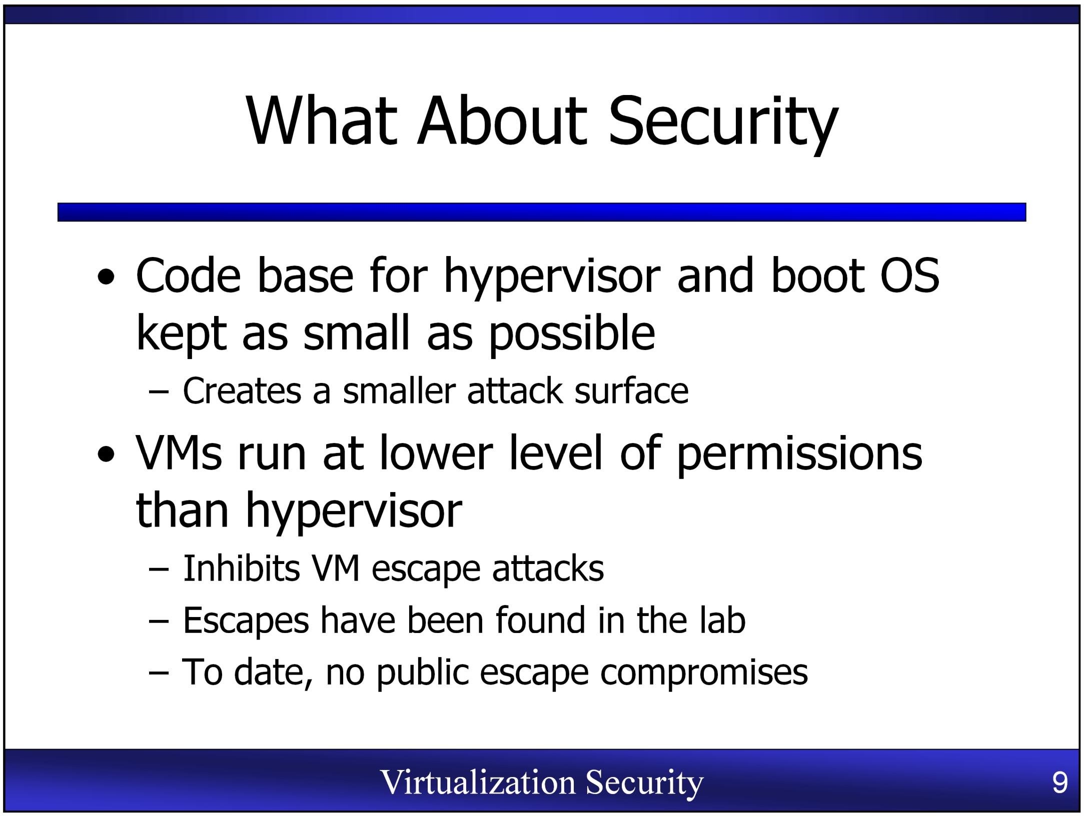

What About Security

- Code base for hypervisor and boot OS kept as small as possible
  - Creates a smaller attack surface
- VMs run at lower level of permissions than hypervisor
  - Inhibits VM escape attacks
  - Escapes have been found in the lab
  - To date, no public escape compromises

As mentioned previously, reducing the amount of code being executed on the virtualization host reduces chances of an exploit being introduced.

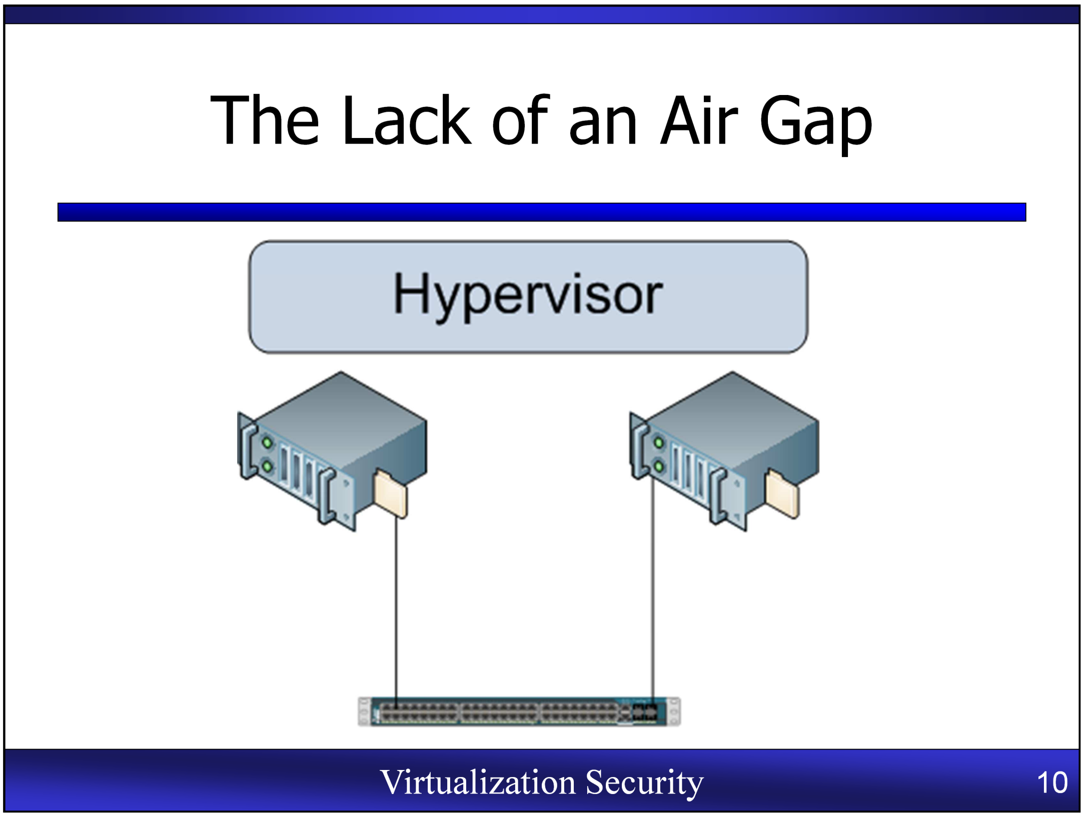

The Lack of an Air Gap

Virtualization Security 10

In a legacy network, some semblance of an air gap exists between operating systems. For example two systems connected to the same Ethernet network can only communicate with each other via the Ethernet network. If that network is disconnected or firewalled, the systems will be unable to communicate with each other.

In a virtualized environment however, the hypervisor always creates a software connection between systems. There is no way to completely isolate one operating system from another, without migrating one of the operating systems to a different hardware platform. It is this persistent software connection that has lead many to feel that virtualization can never be configured as securely as a legacy network.

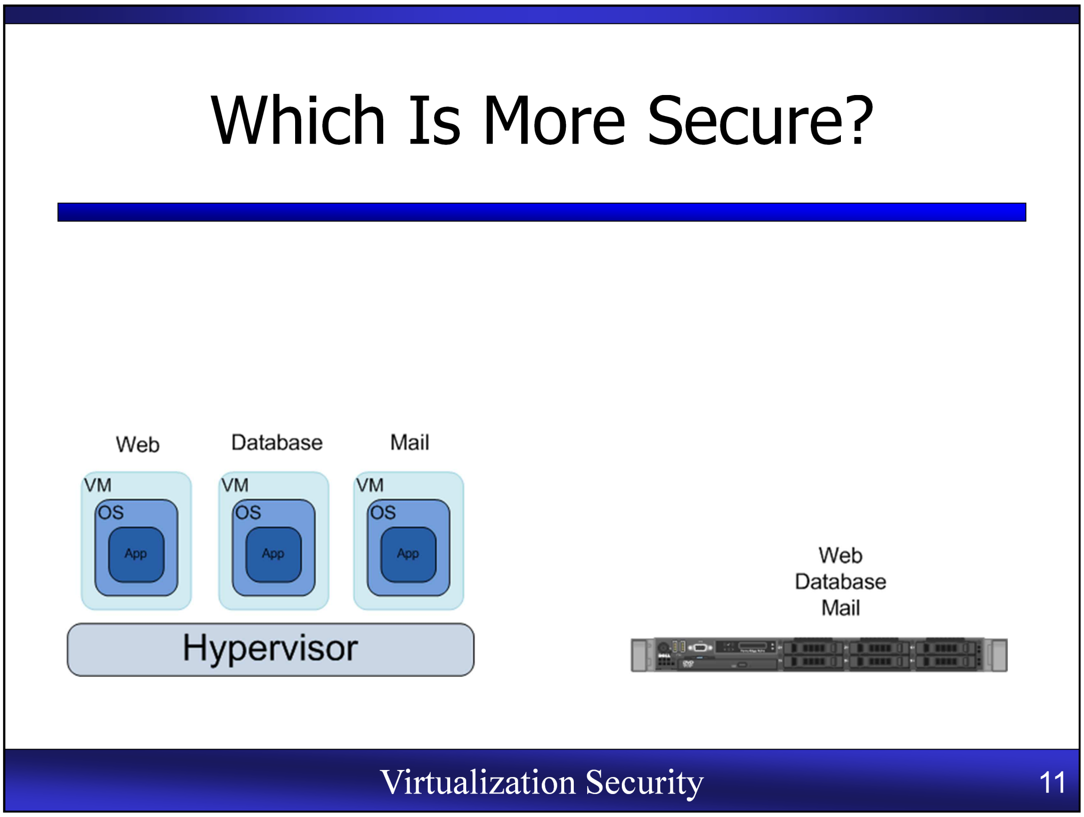

Which Is More Secure?

Virtualization Security 11

Consider the two deployments shown in the above slide. The system on the right represents how we have historically deployed servers. While we would ideally like to isolate each service to its own server, the reality is hardware availability and cost typically force us to combine multiple services onto the same system. This means that if any one of these services are compromised, the other services will quickly fall as well.

Now consider the virtualization deployment on the left. Because virtualization permits us to run multiple operating systems on the same hardware platform, it becomes far more cost effective to give each service its own dedicated operating system. So while its true each operating system is connected via a software bridge in the hypervisor, the fact that each service is running within its own VM container reduces the risk of compromise below the level of the legacy deployment.

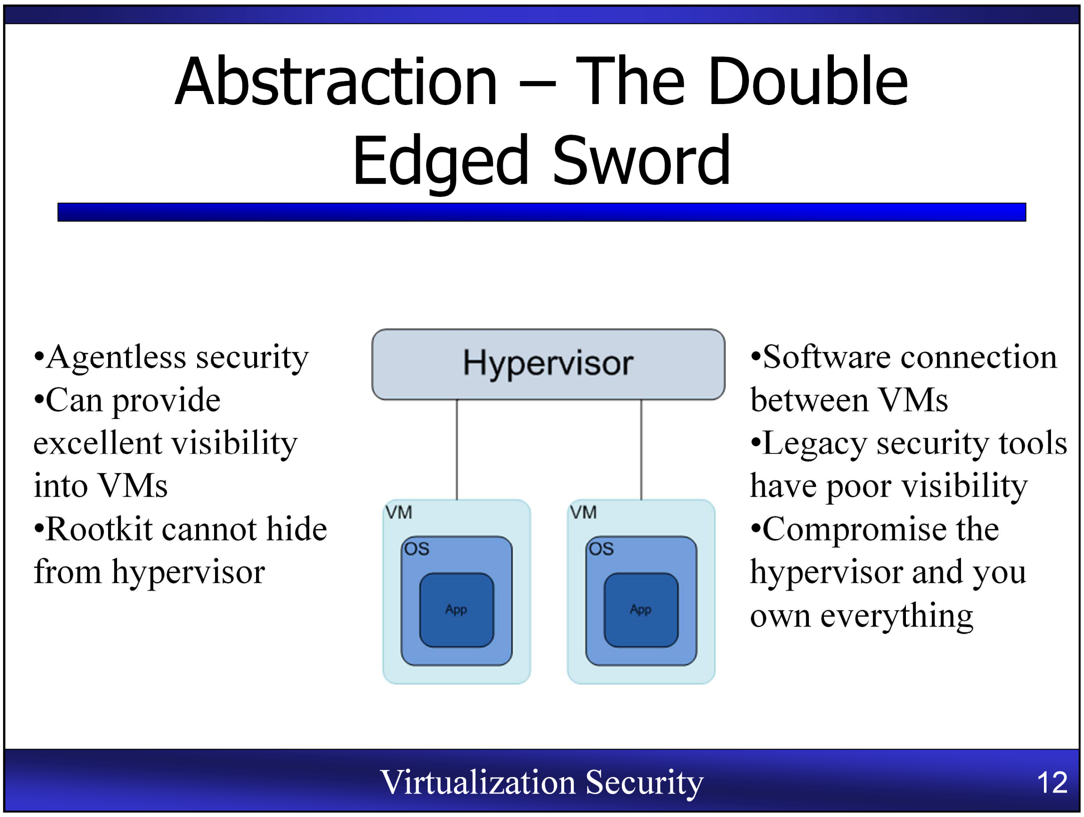

Abstraction – The Double Edged Sword

- Agentless security
- Can provide excellent visibility into VMs
- Rootkit cannot hide from hypervisor

- Software connection between VMs
- Legacy security tools have poor visibility
- Compromise the hypervisor and you own everything

Virtualization Security 12

The above slide shows some of the security gains and loses experienced when moving to virtualization. The trick is to leverage the new capabilities in order to augment the deficiencies.

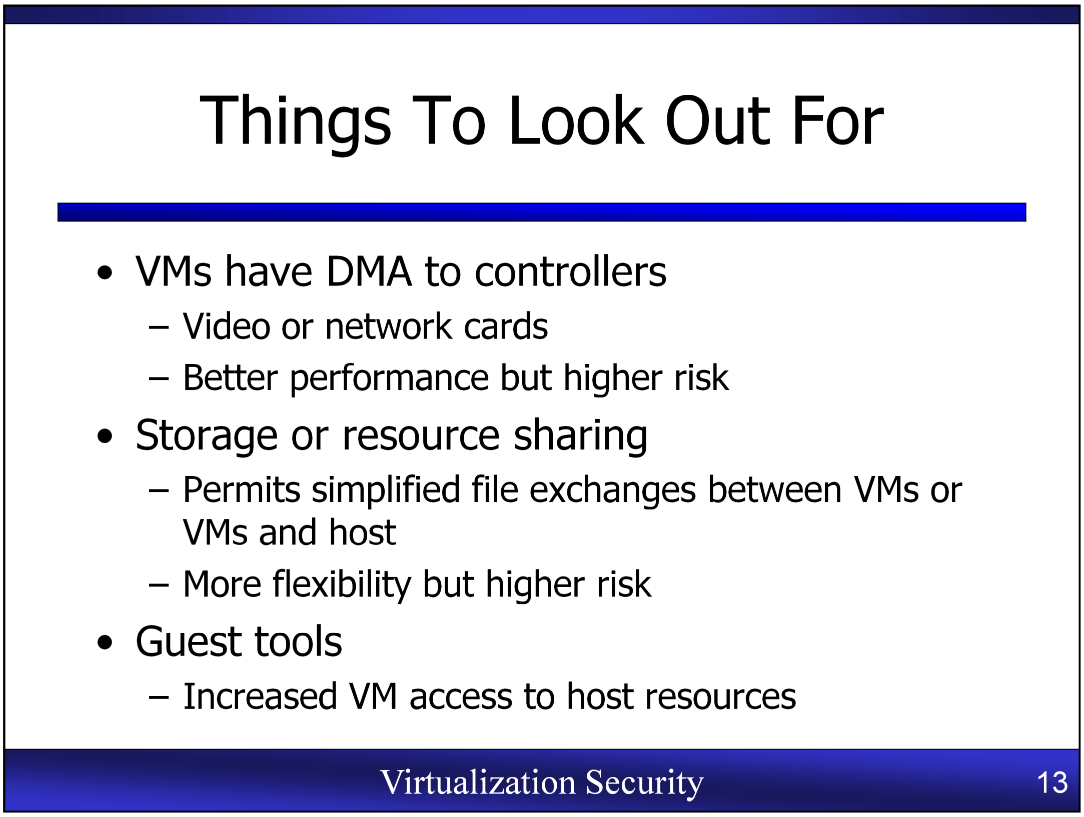

Things To Look Out For

- VMs have DMA to controllers
  - Video or network cards
  - Better performance but higher risk
- Storage or resource sharing
  - Permits simplified file exchanges between VMs or VMs and host
  - More flexibility but higher risk
- Guest tools
  - Increased VM access to host resources

Virtualization Security 13

The above slide shows virtualization settings that can be problematic from a security perspective. Direct Memory Access (DMA) can potentially permit an attack to use storage on peripheral devices in order to move code off of the VM. Features like clipboard sharing can increase functionality, but can also open up similar security holes.

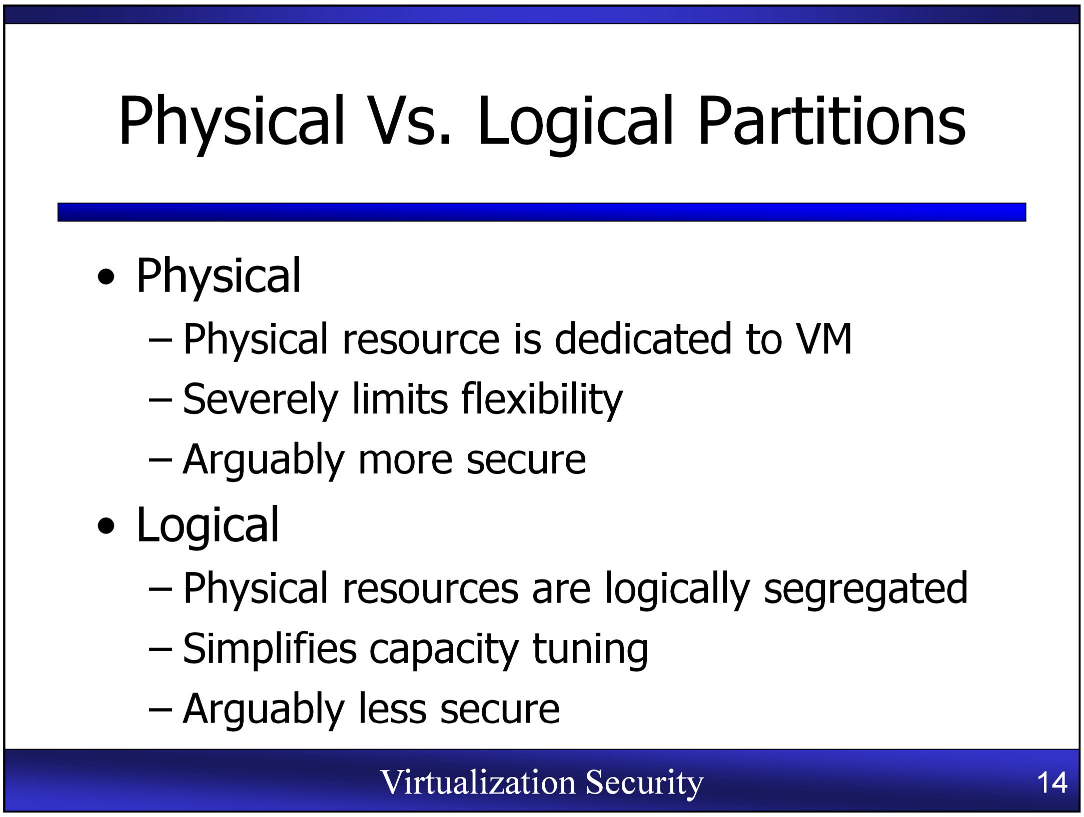

Physical Vs. Logical Partitions

- Physical
  - Physical resource is dedicated to VM
  - Severely limits flexibility
  - Arguably more secure
- Logical
  - Physical resources are logically segregated
  - Simplifies capacity tuning
  - Arguably less secure

When working with virtualized resources, pay attention to whether you are working with physical or logical resources. A physical resource is an actual physical device, such as a network card or storage drive. A logical resource is a virtualized resource configured to appear as a physical resource. For example a virtual machine may “see” a 50 GB hard drive. In reality this may simply be a logical partition of a 300 GB physical drive, which is being shared across multiple virtual machines (VMs).

While resources in a virtualized environment are typically shared between VMs, it is possible to allocate a physical resource to a specific VM instance. For example a specific storage array could be dedicated for use by a specific VM. While this reduces flexibility and increases cost, it does has the benefit of ensuring data storage cannot be inadvertently accessed across VMs.

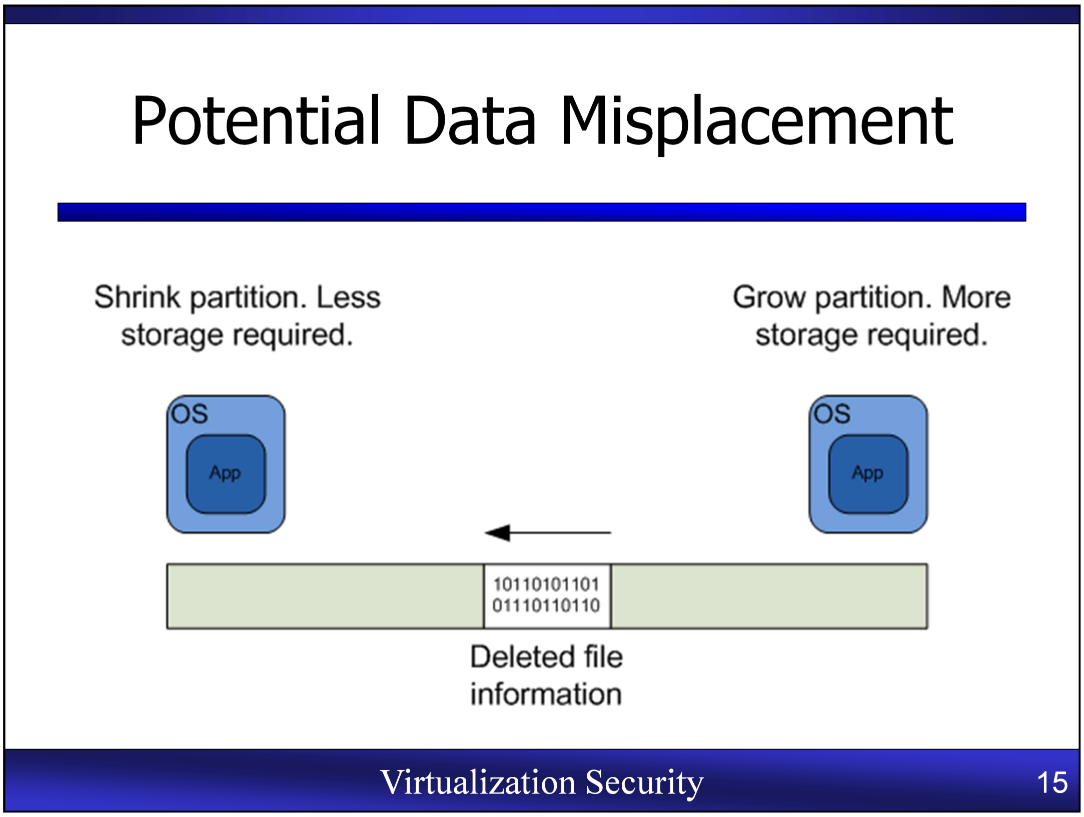

Potential Data Misplacement

Shrink partition. Less storage required.

Grow partition. More storage required.

Deleted file information

Virtualization Security 15

The above slide shows one of the potential security issues that can occur when storage resources are shared. Remember that in a IaaS environment each VM is typically stored as a single file. As storage requirements change, those files may be resized. Reducing the size of one partition and increasing the size of another creates the possibility that sectors containing deleted file information will effectively move from one VM to another. This could permit the owner of the second VM to recover file information stored as part of the first VM.

Again, dedicating physical storage ensure that this issue does not surface. Another possible solution is to encrypt all file information stored to disk. If encrypted, moved sectors would be unreadable without the appropriate key(s).

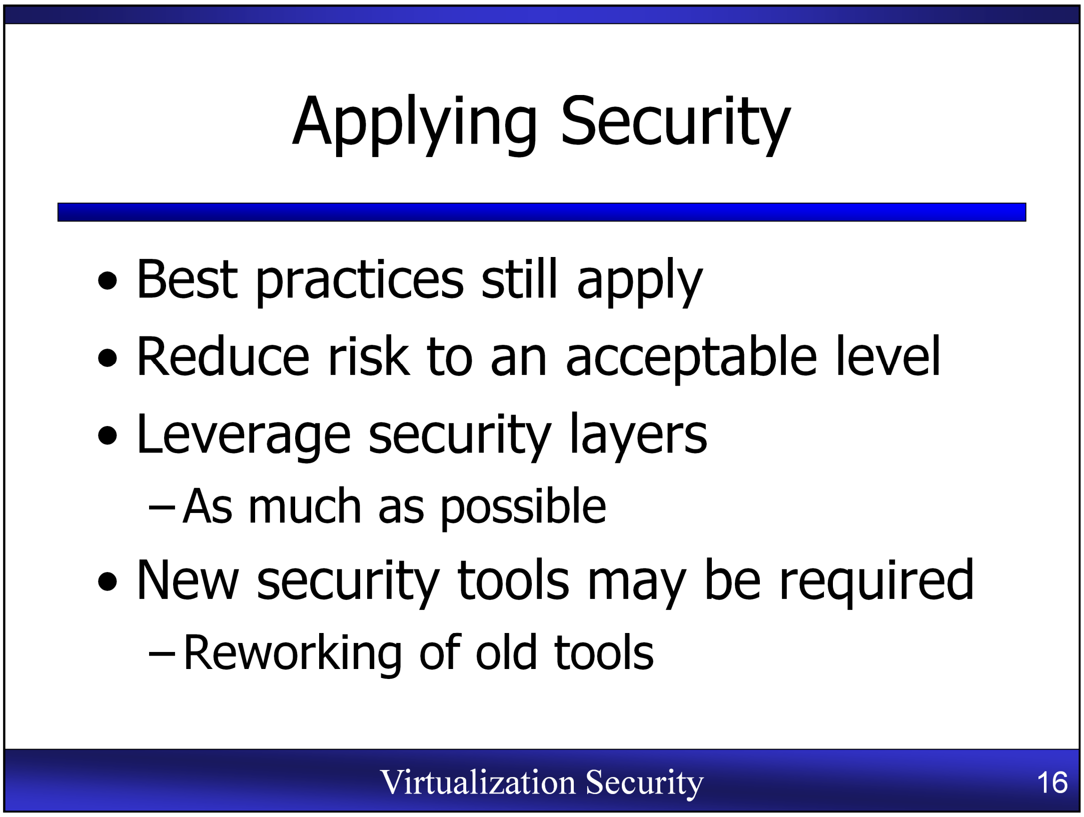

Applying Security

- Best practices still apply
- Reduce risk to an acceptable level
- Leverage security layers
  - As much as possible
- New security tools may be required
  - Reworking of old tools

While both cloud and virtualization bring new security challenges, many of the same logic and processes we have used to reduce risk in the past can still be applied. What changes is the security tools we can leverage as well as how security should best be applied.

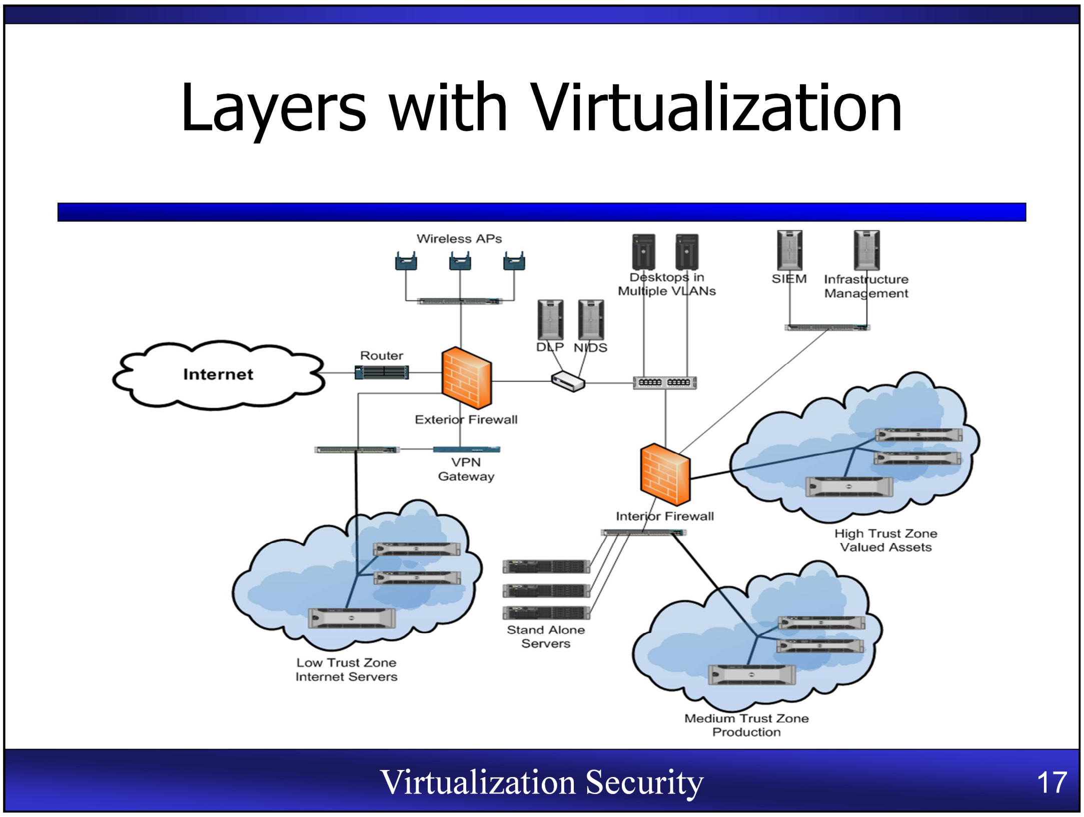

Layers with Virtualization

Wireless APs
Desktops in Multiple VLANs
SIEM
Infrastructure Management
DLP
NIDS
Router
Internet
Exterior Firewall
VPN Gateway
Interior Firewall
High Trust Zone Valued Assets
Low Trust Zone Internet Servers
Stand Alone Servers
Medium Trust Zone Production

Virtualization Security 17

Typically when we determine which servers to virtualize, we look at performance metrics such as average server utilization. The lower the utilization level, the more likely the server will make a good candidate for virtualization.

Security also needs to be part of this equation. When we virtualize a server with no additional security controls (such as hypervisor malware control), we can potentially increase the risk to that server. This may be acceptable for low value data, or it may be completely unacceptable for extremely sensitive information. A good risk analysis will guide us either way. This is where the risk zones shown above come into our design. For example all of the virtualized servers in the “medium trust zone” will most likely require only minor security enhancements to mitigate risk to the proper level. The “high trust” zone, however, will contain servers that will most likely require additional security precautions. So by grouping our servers by risk level we not only enhance manageability but make more efficient use of our security resources.