Course: Cloud Computing

Week 5: Virtualization

Lecturer: Ikwap Flavia Agatha
MSc. Computer Forensic
PHD in IT (Candidate)
University: Kumi University

Lecture learning out come:

At the end of this lesson, you will be able to:

- Understand the concept of Virtualization
- Understand the components of a virtualized environment
- Understand the different approaches of virtualization
- Understand Virtualization in cloud services
- Understand the advantages of virtualization

Virtualization Concept

- Refers to the act of creating a virtual (rather than actual) version of something, including but not limit to virtual computer hardware platform, operating system, storage device or computer network resources.
- In virtualization a single physical machine (hardware) is configured to use multiple Virtual machines within it
- Virtualization is a widely used technology in the IT industry to increase resource utilization and ROI. It allows the same physical infrastructure to be shared between multiple OSs and applications (Buyya, 2013).

Virtualization Environment-Components

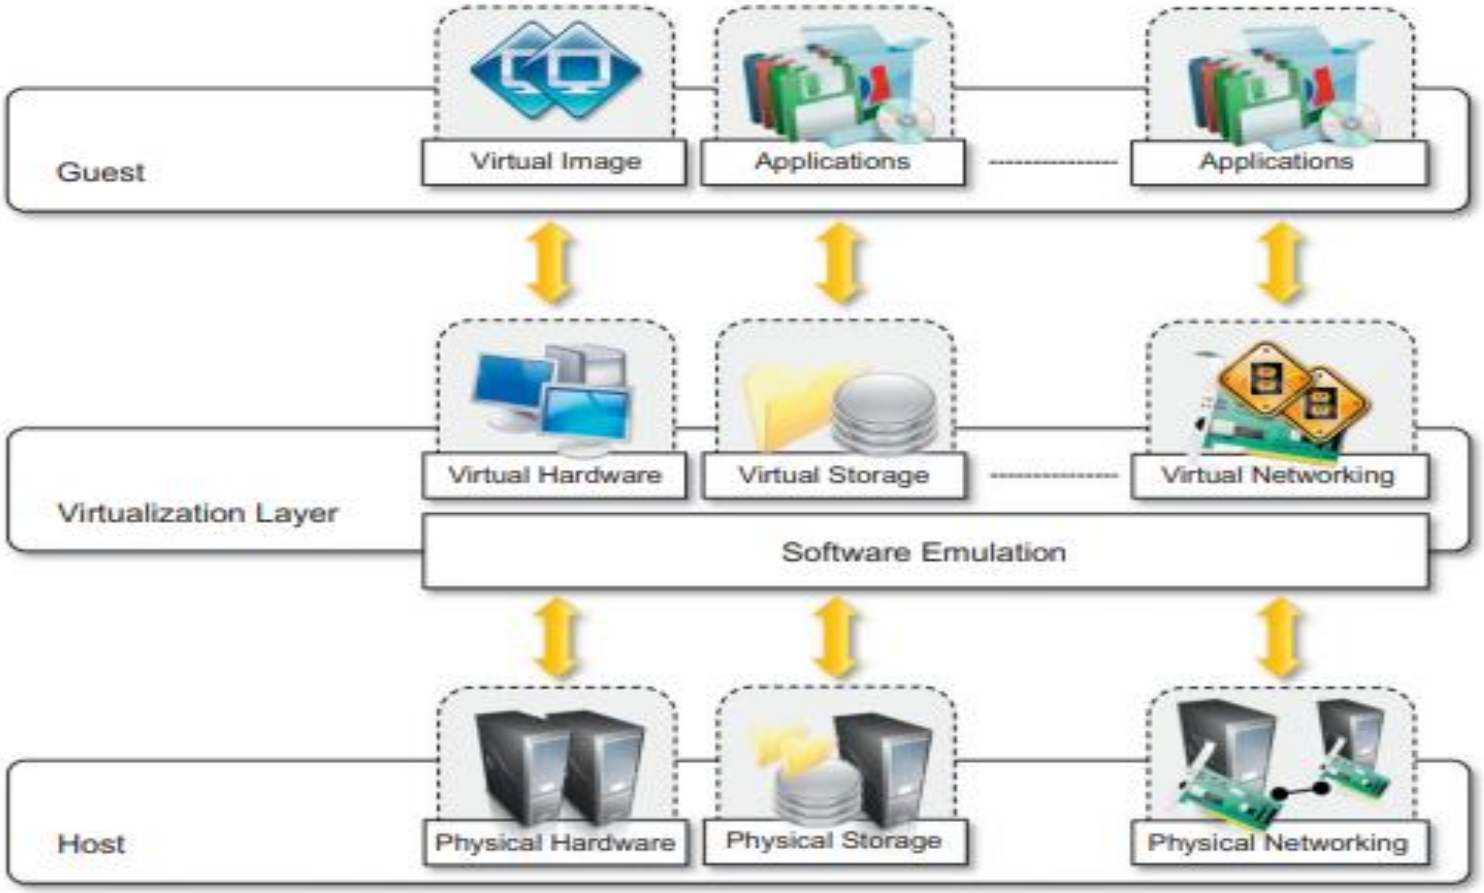

- (Buyya, 2013)

Virtualization Environment-Components

- Guest-. The guest represents the system component that interacts with the virtualization layer rather than with the host.
- Host-The host represents the original environment where the guest is supposed to be managed.
- The virtualization layer: - is responsible for recreating the same or a different environment where the guest will operate.

Virtualization techniques

Execution virtualization

- Execution virtualization includes all techniques that aim to emulate an execution environment that is separate from the one hosting the virtualization layer. All these techniques concentrate their interest on providing support for the execution of programs, whether these are the operating system, a binary specification of a program compiled against an abstract machine model, or an application.

Machine reference model

- At the bottom layer, the model for the hardware is expressed in terms of the Instruction Set Architecture (ISA).
- ISA defines the instruction set for the processor, registers, memory, and interrupts management.

Machine reference model

- The application binary interface (ABI) separates the operating system layer from the applications and libraries, which are managed by the OS. ABI covers details such as low-level data types, alignment, and call conventions and defines a format for executable programs. System calls are defined at this level.
- The highest level of abstraction is represented by the application programming interface (API), which interfaces applications to libraries and/or the underlying operating system.

Machine reference model

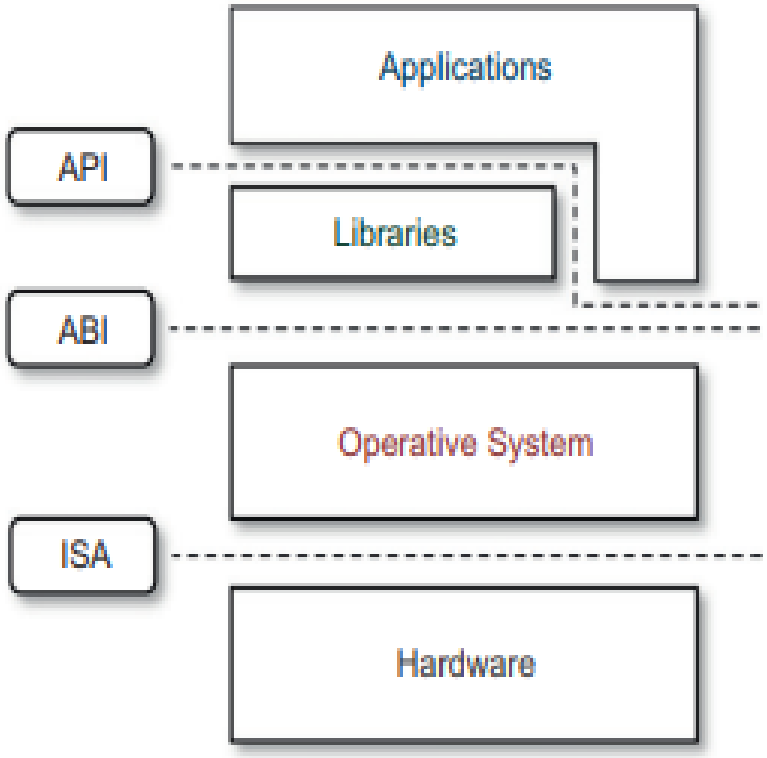

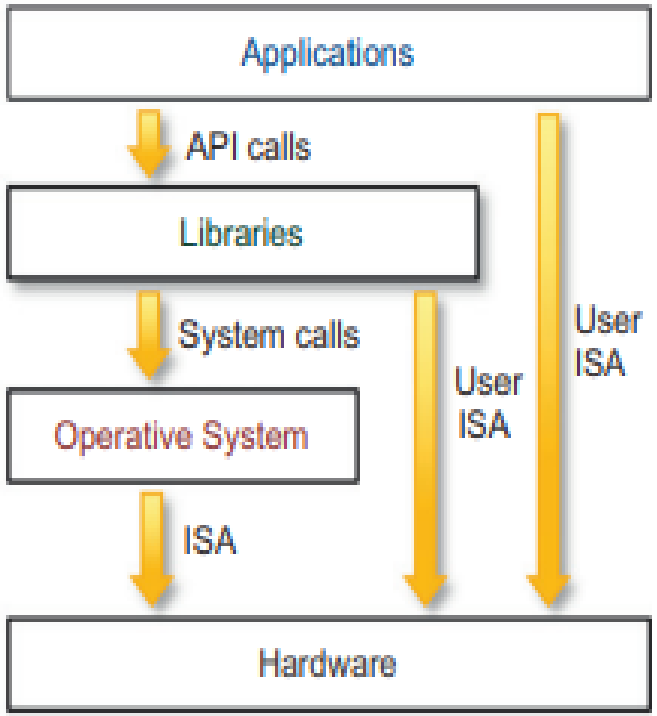

- (Buyya, 2013)

Virtual Machine Manager

- Three main modules, dispatcher, allocator, and interpreter, coordinate their activity in order to emulate the underlying hardware.
- The dispatcher constitutes the entry point of the monitor and reroutes the instructions issued by the virtual machine instance to one of the two other modules.

Virtual Machine Manager

- The allocator is responsible for deciding the system resources to be provided to the VM: whenever a virtual machine tries to execute an instruction that results in changing the machine resources associated with that VM, the allocator is invoked by the dispatcher.
- The interpreter module consists of interpreter routines. These are executed whenever a virtual machine executes a privileged instruction: a trap is triggered and the corresponding routine is executed

Virtual Machine Manager

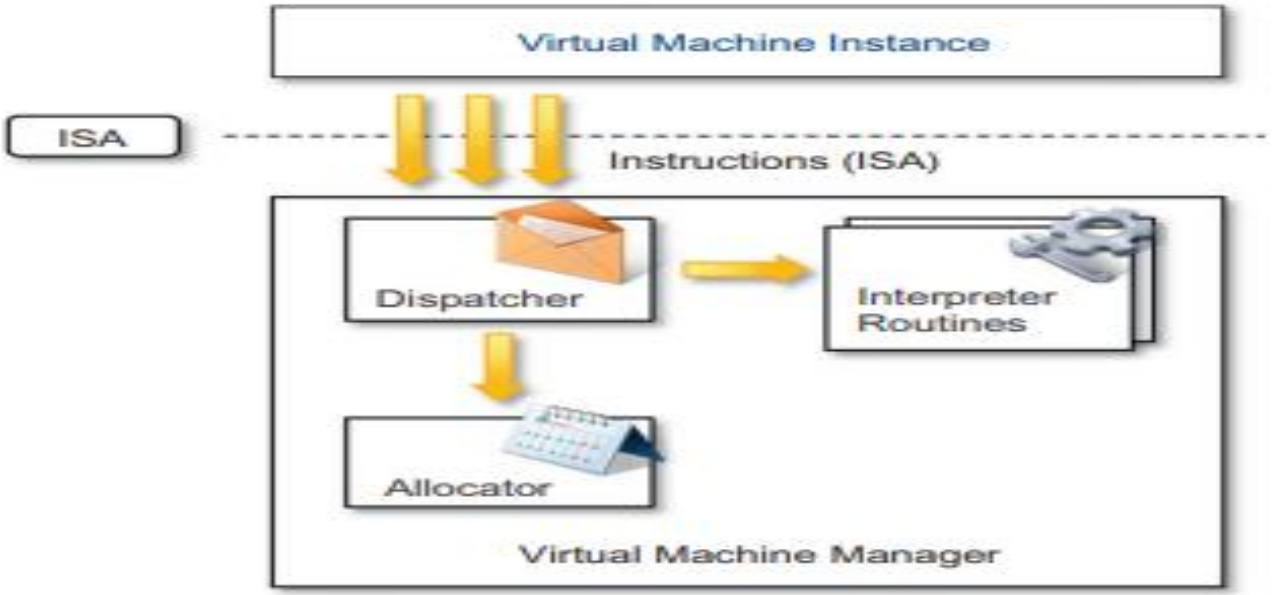

- (Buyya, 2013)

Hardware-level virtualization

- The hypervisor is generally a program or a combination of software and hardware that allows the extraction of the underlying physical hardware. Hardware-level virtualization is also called system virtualization, since it provides ISA to virtual machines, which is the representation of the hardware interface of a system. This is to differentiate it from process virtual machines, which expose ABI to virtual machines

Hardware-level virtualization

- Hypervisors: A fundamental element of hardware virtualization is the hypervisor, or virtual machine manager (VMM). It recreates a hardware environment in which guest operating systems are installed.
- There are two major types of hypervisor:

Type I and Type II

- Type I hypervisors run directly on top of the hardware. Therefore, they take the place of the operating systems and interact directly with the ISA interface exposed by the underlying hardware, and they emulate this interface in order to allow the management of guest operating systems. This type of hypervisor is also called a native virtual machine since it runs natively on hardware.

Type I and Type II

- It can run and access physical resources directly without the help of any host OS.
- This type of hypervisors is used for servers that handle heavy load and require more security. Some examples of type 1 hypervisors include Microsoft Hyper-V, Citrix XenServer, VMWare ESXi, and Oracle VM Server for SPARC.

Type I

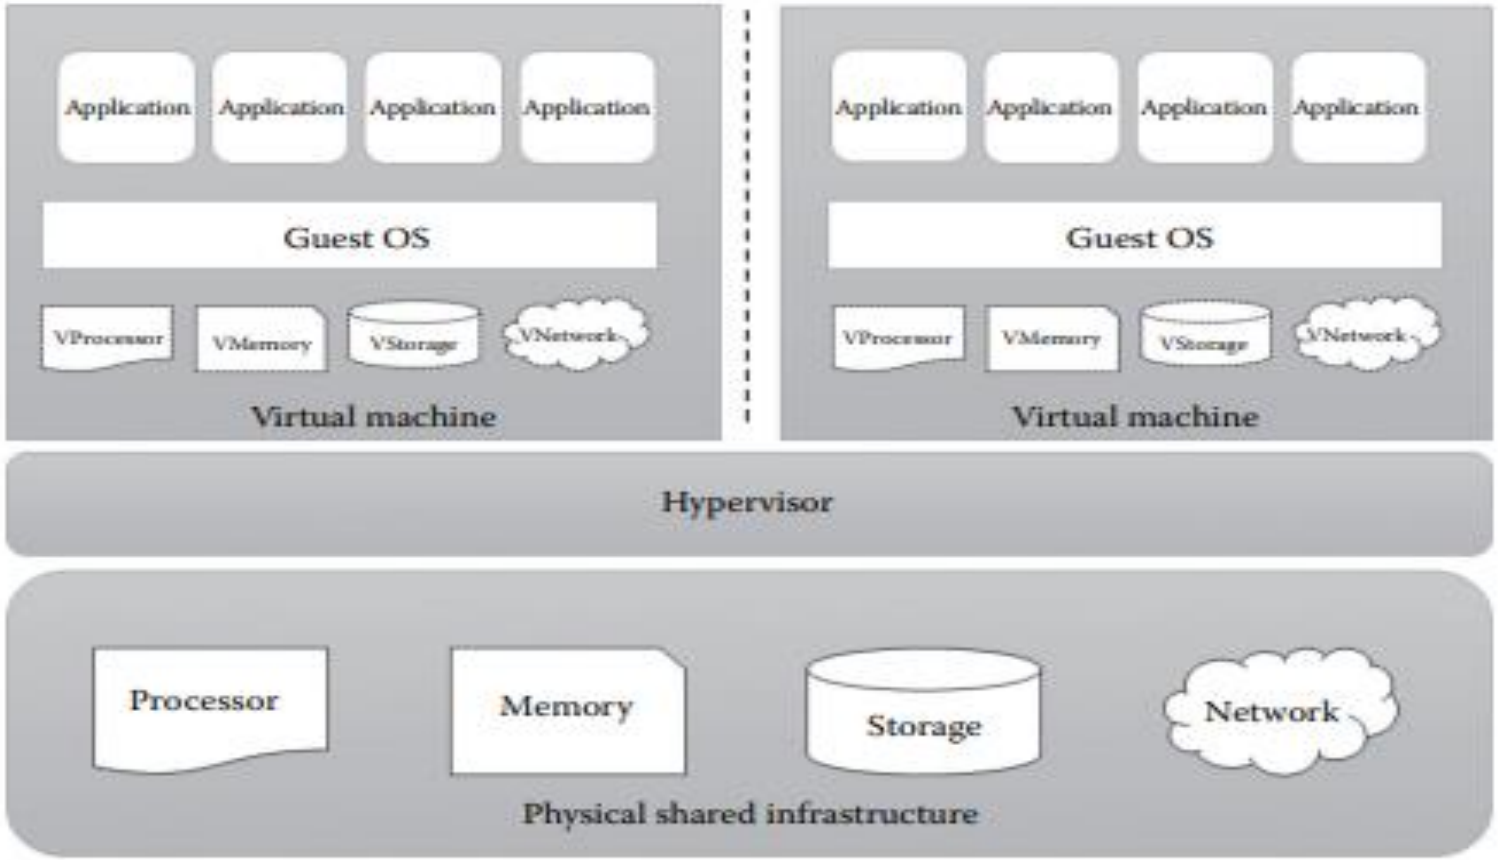

- (K.Chandrasekaran, 2015)

Type I and Type II

- Type II hypervisors require the support of an operating system to provide virtualization services. This means that they are programs managed by the operating system, which interact with it through the ABI and emulate the ISA of virtual hardware for guest operating systems. It requires the host OS and does not have direct access to the physical hardware. These types of hypervisors are installed on the host OS as a software program. This type of hypervisor is also called a hosted virtual machine since it is hosted within an operating system.

Type II

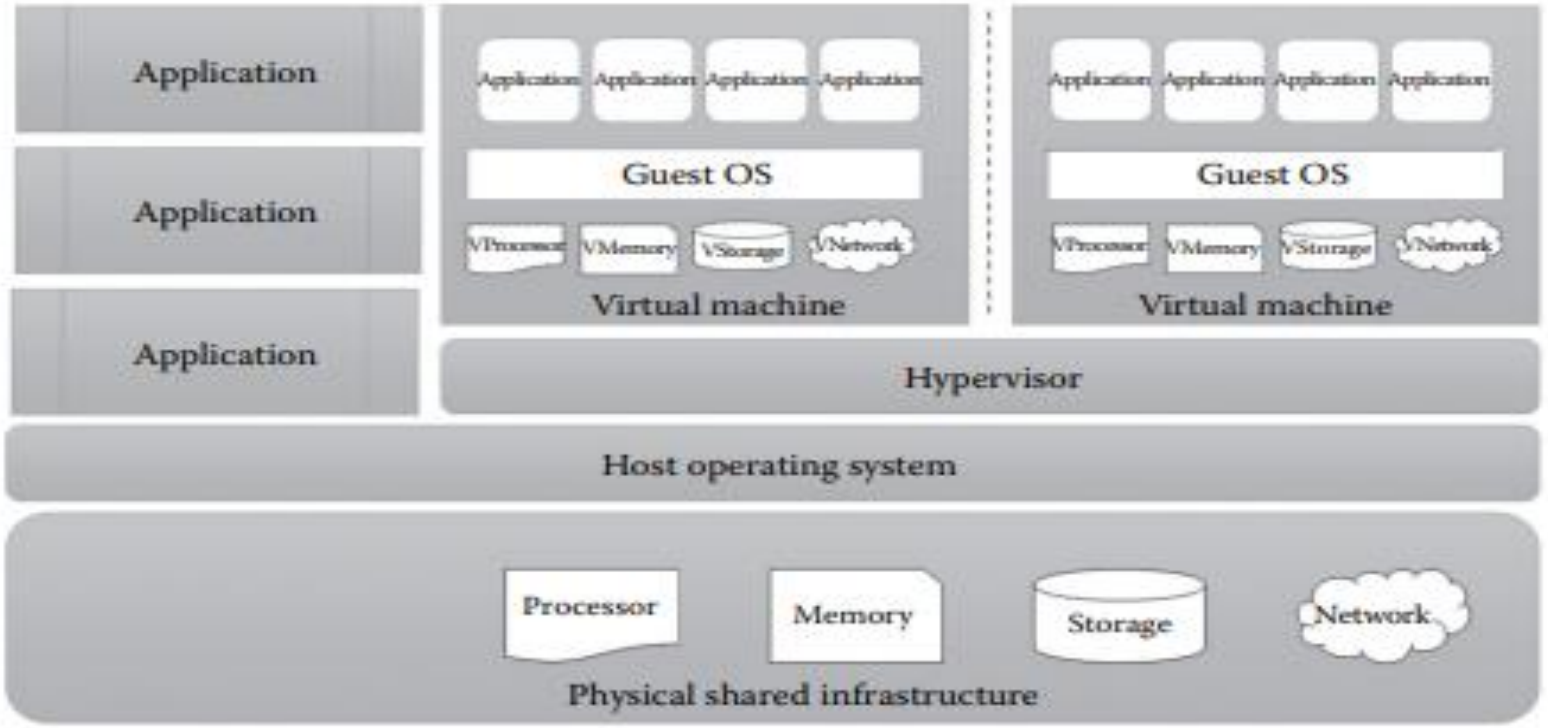

- (K.Chandrasekaran, 2015)

Security Issues in hypervisor

- The hypervisor creates a virtual environment in the data centers and can be attacked through malicious code written by any attacker to disrupt or corrupt the whole server. Hypervisor is the higher authority entity that has the direct access to the hardware hence attracting attackers.
- There are two possibilities of attacking the hypervisor: 1. through the host OS. 2. Through the guest OS
- Attacks from the host OS can be performed by exploiting the vulnerabilities of the host OS

Security Issues in hypervisor

- Once the attacker gets full control over the hypervisor through the compromised OS, the attacker will be able to run all the privileged instructions that can control the actual hardware, he can then carry out malicious activities:
- Denial of service attack
- Stealing the confidential information
- Malicious script.

Through the host OS

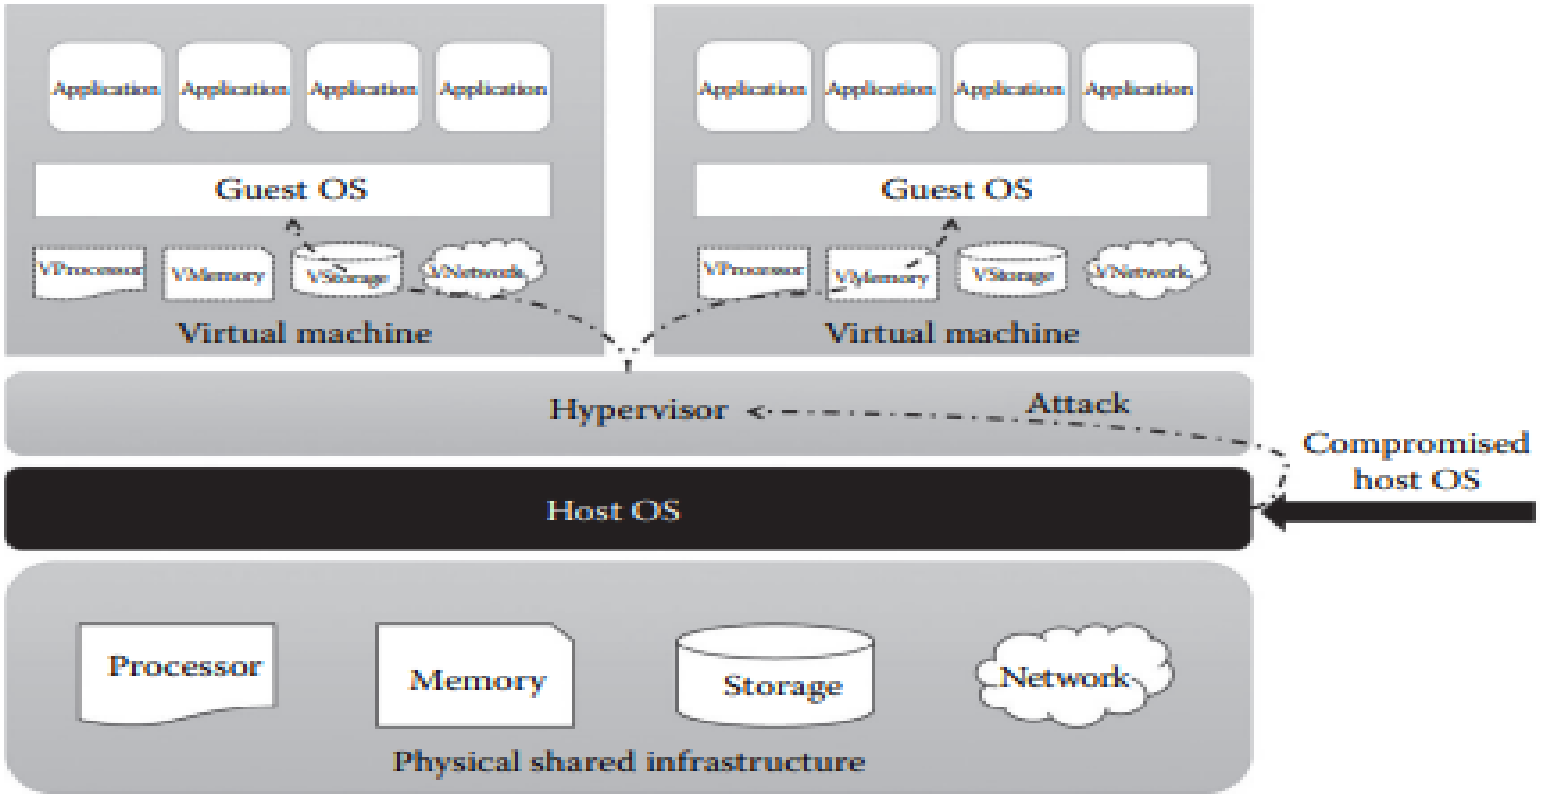

- (K.Chandrasekaran, 2015)

Through the guest OS

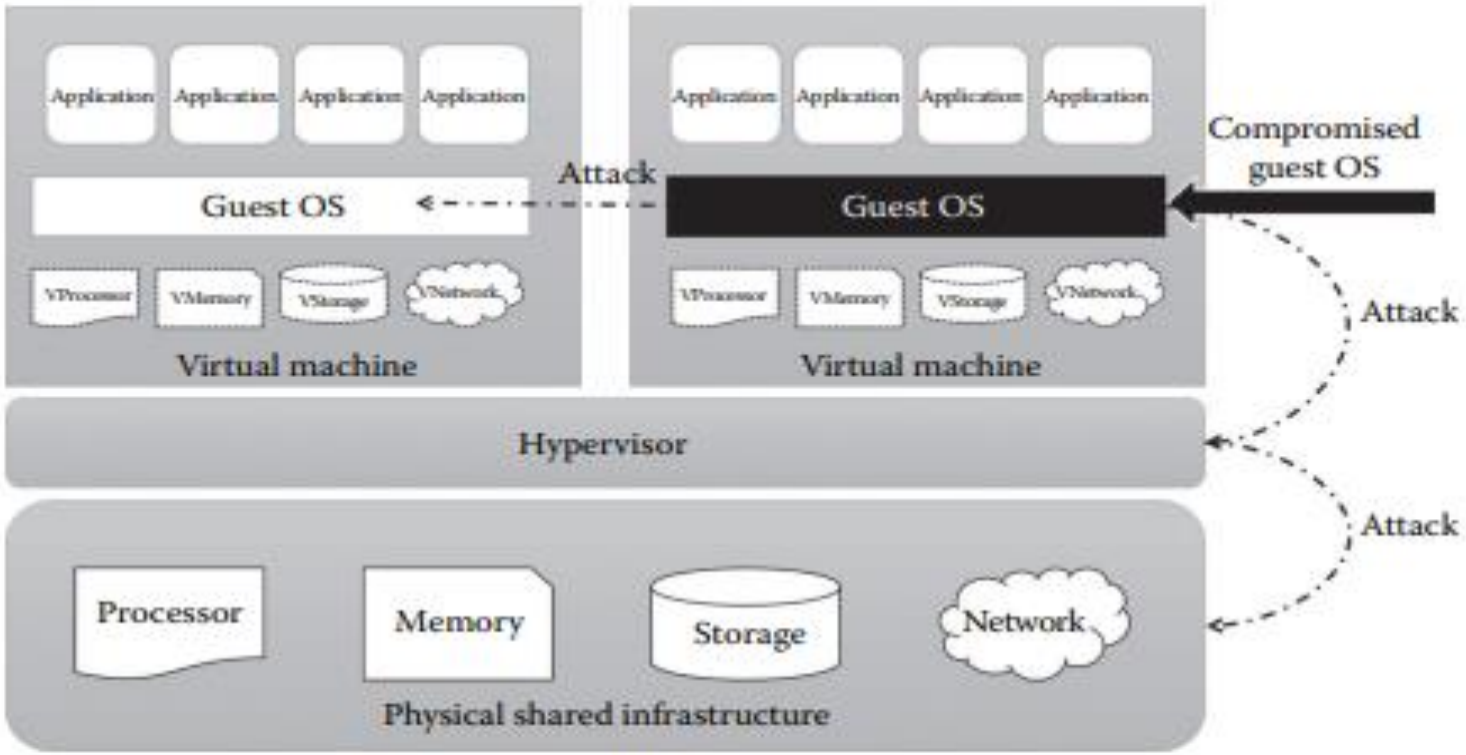

- (K.Chandrasekaran, 2015)

Through the guest OS

- Since the guest OS is communicating with the hypervisor to get virtual resources, any malicious code from the guest OS or VMs can compromise the hypervisor. Normally, the attacks from the guest OS will try to abuse the underlying resources.

Approaches of virtualization

- Full virtualization
- Para-virtualization
- Partial Virtualization.

Full-virtualization.

- Full virtualization refers to the ability to run a program, like an operating system, directly on top of a virtual machine and without any modification, as though it were run on the raw hardware. To make this possible, virtual machine managers must provide a complete emulation of the entire underlying hardware.
- The advantage of full virtualization is complete isolation, which leads to enhanced security, ease emulation of different architectures, and coexistence of different systems on the same platform.

Full-virtualization.

- The challenge is the interception of privileged instructions such as I/O instructions: Since they change the state of the resources exposed by the host, they have to be contained within the virtual machine manager.
- A successful and efficient implementation of full virtualization is obtained with a combination of hardware and software, not allowing potentially harmful instructions to be executed directly on the host.

Para-virtualization

- Para-virtualization techniques expose a software interface to the virtual machine that is slightly modified from the host and, as a consequence, guests need to be modified.
- The aim of para-virtualization is to provide the capability to demand the execution of performance-critical operations directly on the host.

Para-virtualization

- This allows a simpler implementation of virtual machine managers that have to simply transfer the execution of these operations, which were hard to virtualize, directly to the host.
- This technique has been successfully used by Xen for providing virtualization solutions for Linux-based operating systems specifically ported to run on Xen hypervisors.

Partial-virtualization

- Provides a partial emulation of the underlying hardware, thus not allowing the complete execution of the guest operating system in complete isolation. Partial virtualization allows many applications to run transparently, but not all the features of the operating system can be supported.

Virtualization Opportunities

- Virtualization is the process of abstracting the physical resources to the pool of virtual resources that can be given to any virtual machines (VMs). The different resources like memory, processors, storage, and network can be virtualized using proper virtualization technologies.

Processor Virtualization

- Processor virtualization allows the VMs to share the virtual processors that are abstracted from the physical processors available at the underlying infrastructure.
- The virtualization layer extracts the physical processor to the pool of virtual processors that is shared by the VMs. The virtualization layer will be normally any hypervisors. But processor virtualization can also be achieved from distributed servers

Processor Virtualization

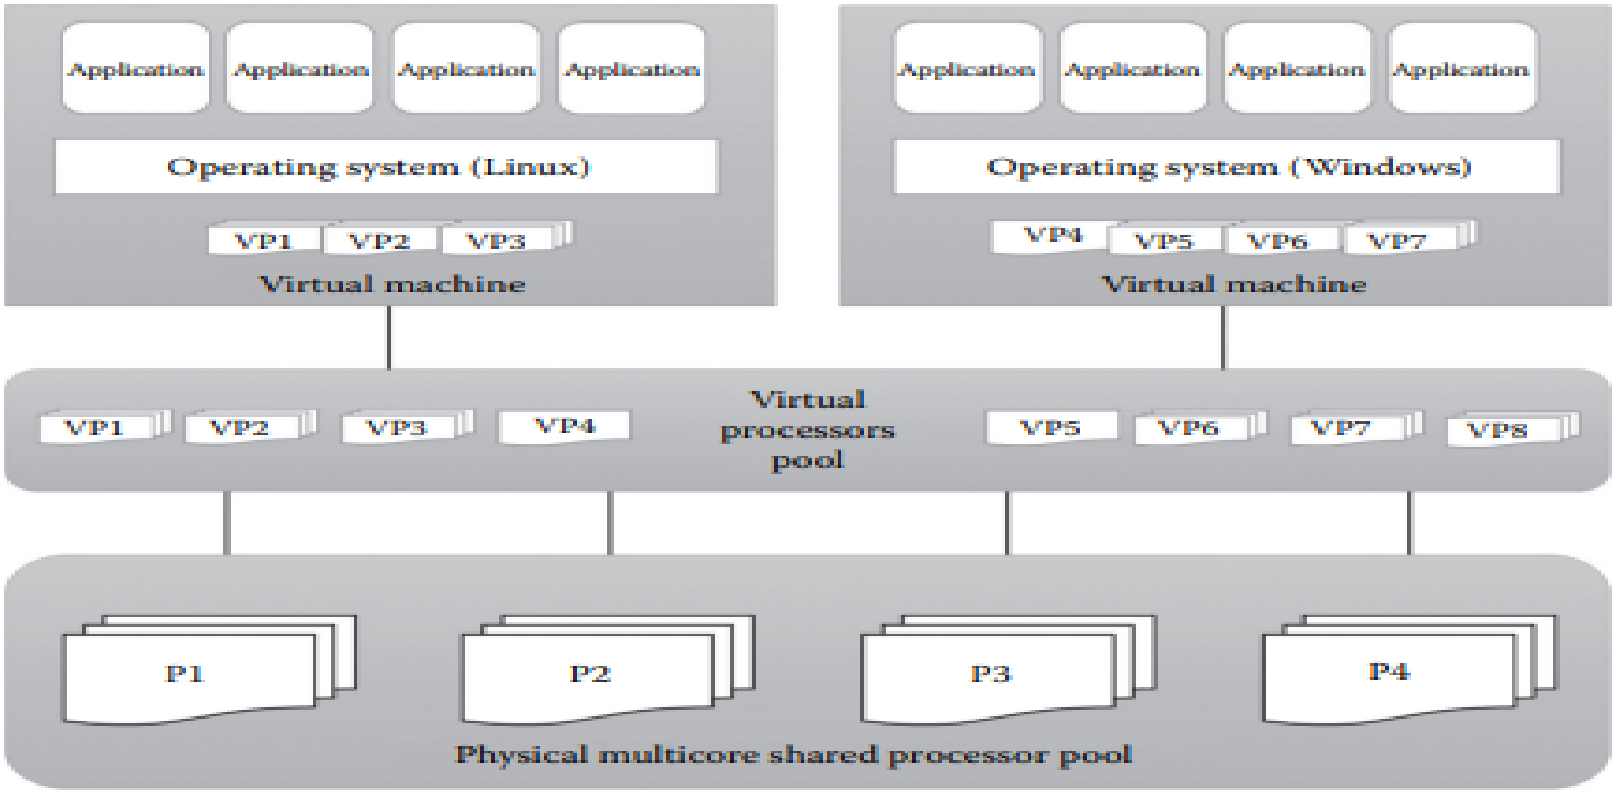

- (K.Chandrasekaran, 2015)

Memory Virtualization

- The process of providing a virtual main memory to the VMs is known as memory virtualization or main memory virtualization. In main memory virtualization, the physical main memory is mapped to the virtual main memory as in the virtual memory concepts in most of the OSs. All the modern x86 processors are supporting main memory virtualization.

Memory Virtualization

- Main memory virtualization can also be achieved by using the hypervisor software. In the virtualized data centers, the unused main memory of the different servers will consolidate as a virtual main memory pool and can be given to the VMs.

Memory Virtualization

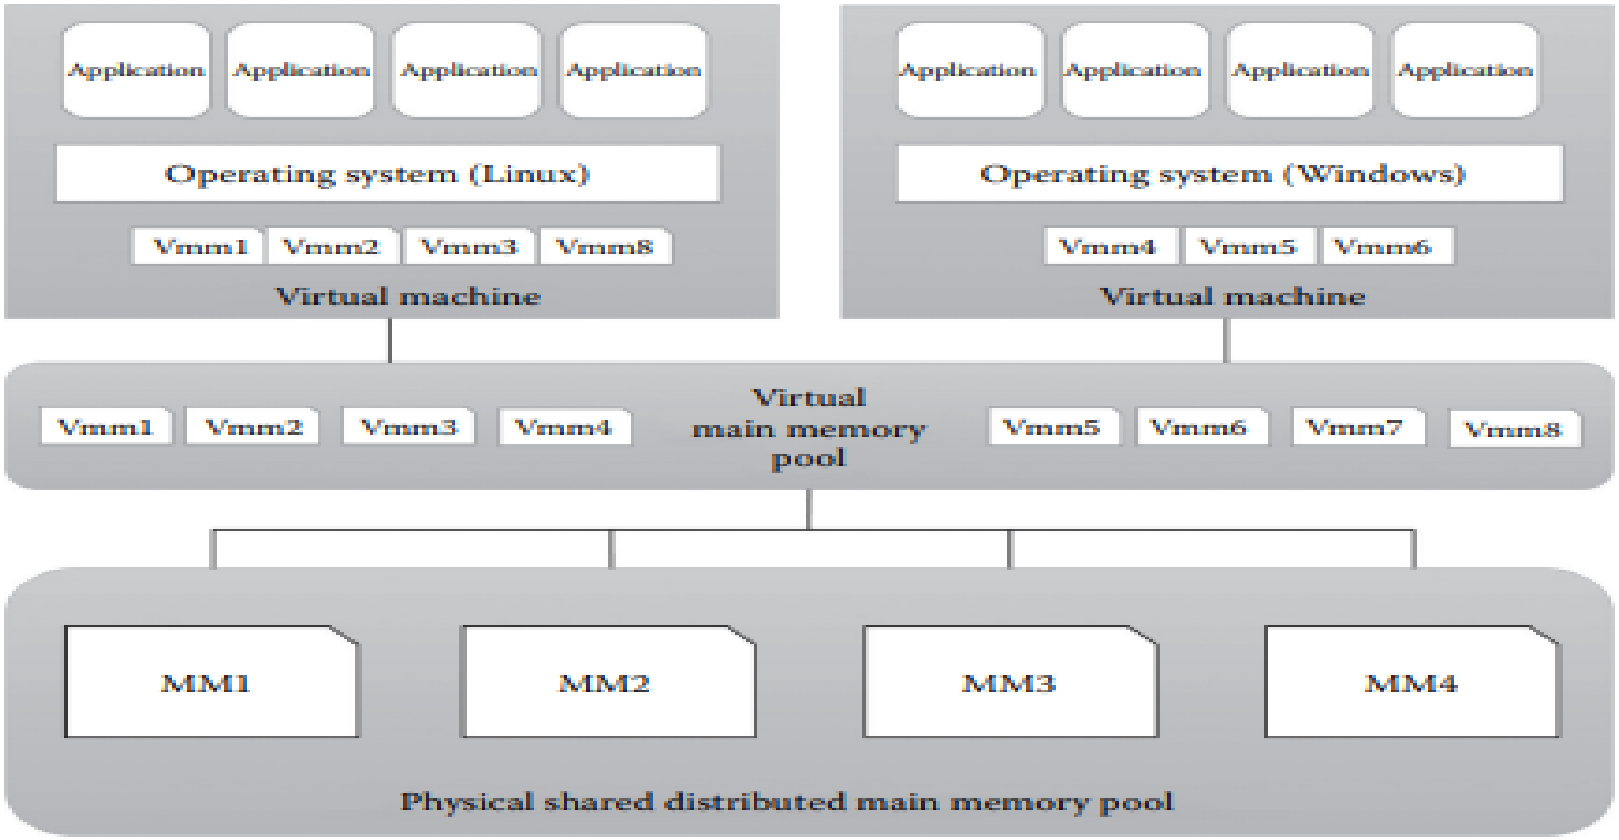

- (K.Chandrasekaran, 2015)

Storage Virtualization

- Storage virtualization is a form of resource virtualization where multiple physical storage disks are extracted as a pool of virtual storage disks to the VMs. The virtualized storage will be called a logical storage. Storage virtualization is mainly used for maintaining a backup or replica of the data that are stored on the VMs. It can also be achieved through the hypervisors. The other advanced storage virtualization techniques are storage area networks (SAN) and network-attached storage (NAS).

Storage Virtualization

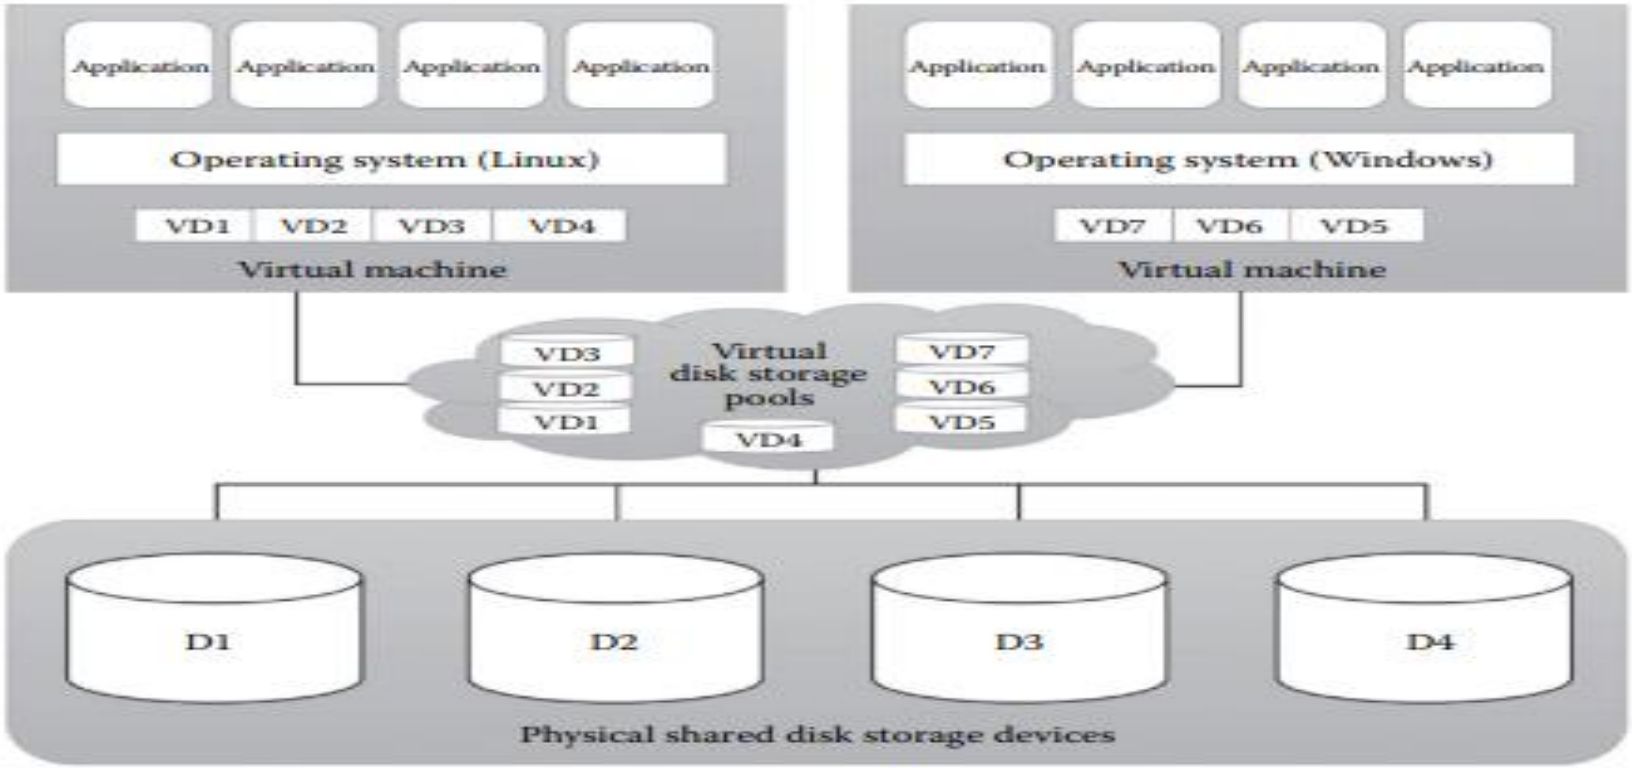

- (K.Chandrasekaran, 2015)

Operating system-level virtualization

- Helps to create different and separated execution environments for applications that are managed concurrently. There is no virtual machine manager or hypervisor, and the virtualization is done within a single operating system, where the OS kernel allows for multiple isolated user space instances. The kernel is also responsible for sharing the system resources among instances and for limiting the impact of instances on each other.

Operating system-level virtualization

- A user space instance in general contains a proper view of the file system, which is completely isolated, and separate IP addresses, software configurations, and access to devices. Operating systems supporting this type of virtualization are general-purpose, timeshared operating systems with the capability to provide stronger namespace and resource isolation.
- Examples of operating system-level virtualizations are FreeBSD Jails, IBM Logical Partition (LPAR), Solaris Zones and Containers, Parallels Virtuozzo Containers, OpenVZ, iCore Virtual Accounts, Free Virtual Private Server (FreeVPS), and others.

Data Virtualization

- Data virtualization is the ability to retrieve the data without knowing its type and the physical location where it is stored. It collects data from the different sources to a single logical/virtual volume of data. This logical data can be accessed from any applications such as web services, E-commerce applications, web portals, Software as a Service (SaaS) applications, and mobile application. Data virtualization hides the type of the data and the location of the data for the application that access it. It is mainly used in data integration, business intelligence, and cloud computing.

Data Virtualization

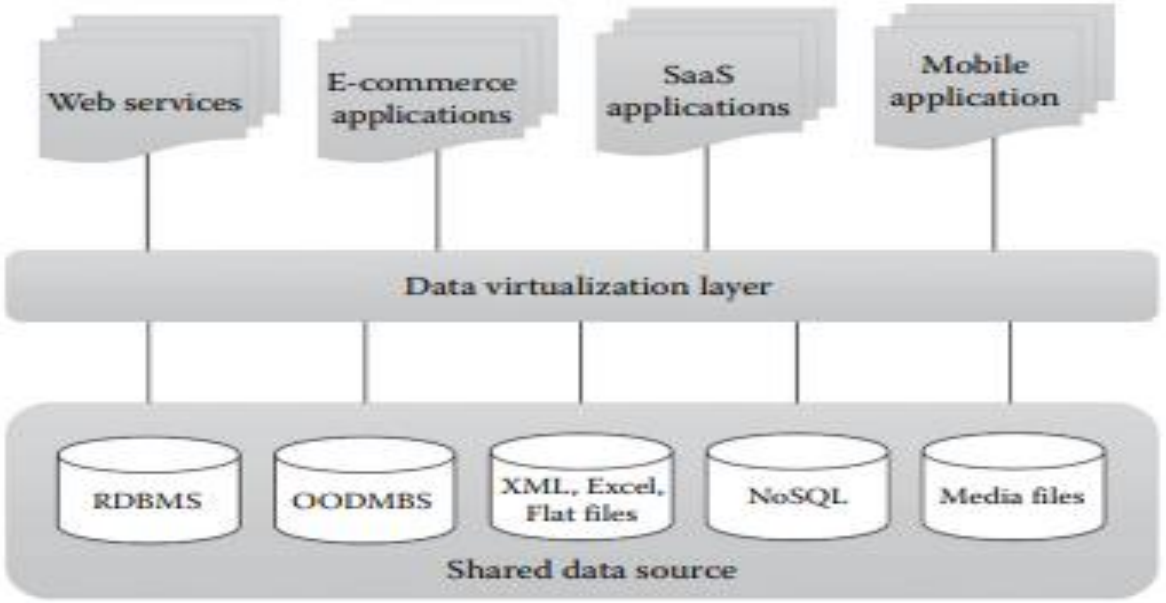

- (K.Chandrasekaran, 2015)

Application Virtualization

- Application virtualization is the enabling technology for SaaS of cloud computing. The application virtualization offers the ability to the user to use the application without the need to install any software or tools in the machine. The applications will be developed and hosted in the central server. The hosted application will be virtualized, and the users will be given the separated/isolated virtual copy to access.

Application virtualization

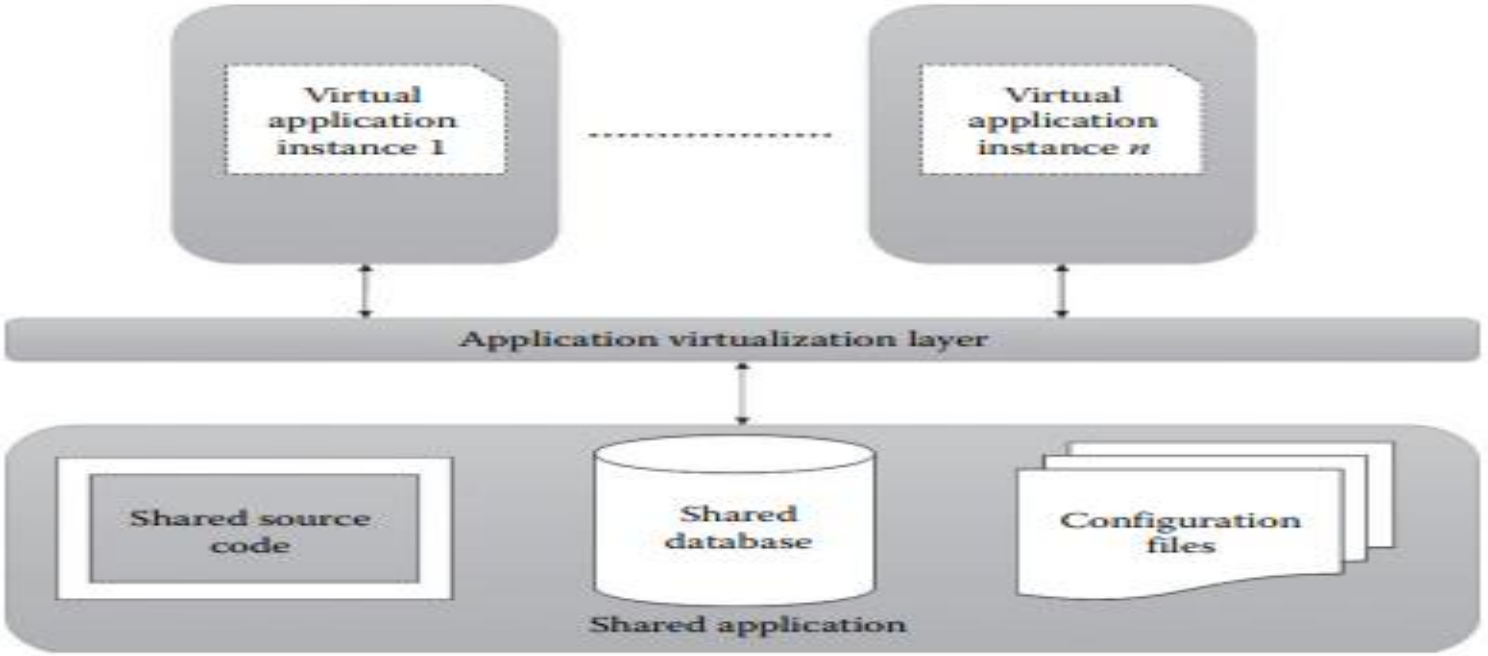

- (K.Chandrasekaran, 2015)

Network Virtualization

- Network virtualization is a type of resource virtualization in which the physical network can be extracted to create a virtual network. Physical network components like router, switch, and Network Interface Card (NIC) will be controlled by the virtualization software to provide virtual network components. The virtual network is a single software-based entity that contains the network hardware and software resources. Network virtualization can be achieved from internal network or by combining many external networks.

Network Virtualization

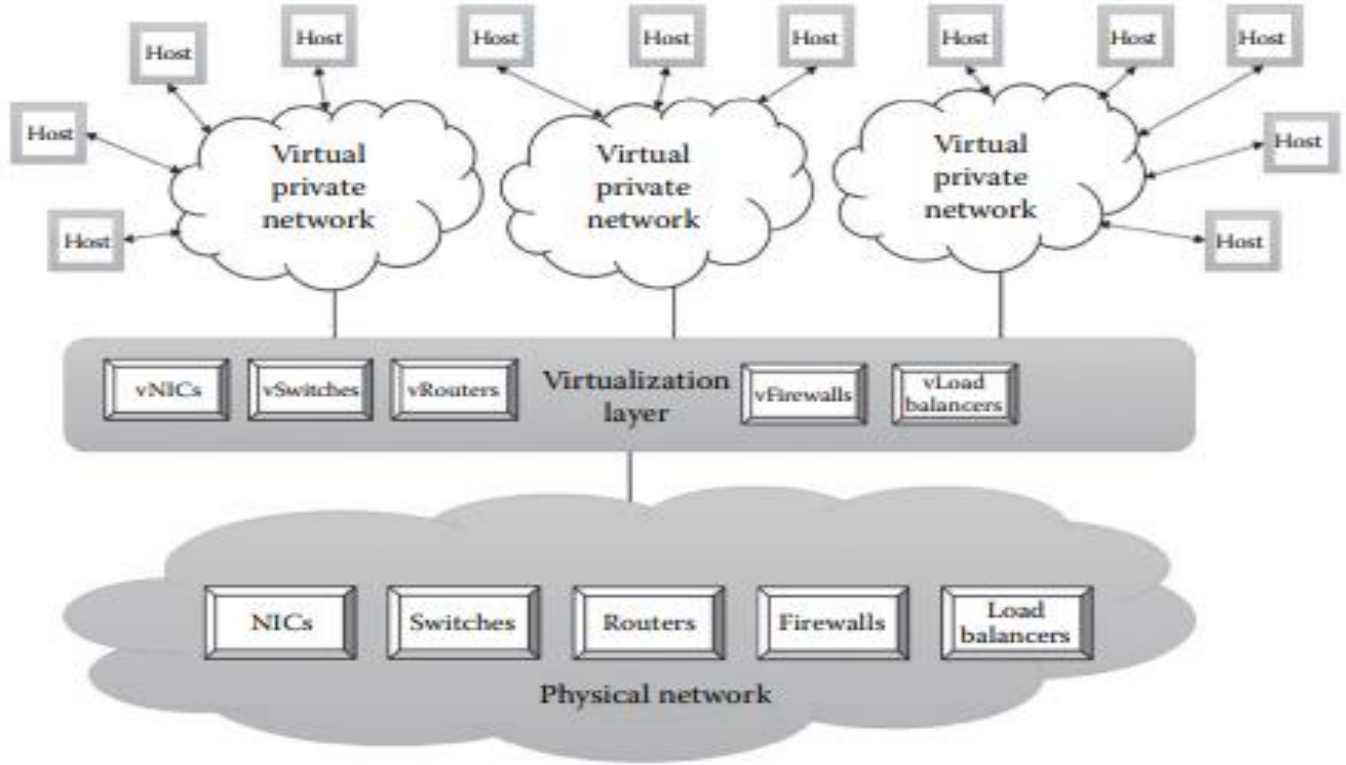

Desktop /client-server Virtualization

- Desktop virtualization enables to store the users' operating system on a server in a data center (this basically gives someone an entire computing platform without the hardware). Through this type of virtualization, employees can work conveniently from their homes. The data transfer is secured, and any risk of data theft is minimized. Examples of desktop virtualization are VMware ThinApps Citrix XenApps, VMware View and Microsoft Remote Desktop Services.

Server Virtualization:

- Instead of assigning one task to one server, in server virtualization multiple tasks run from one server. This causes an increase in performance and the operating cost is reduced. Few examples of server virtualization are FreeVPS, LinuxVserver and OpenVZ

Server Virtualization

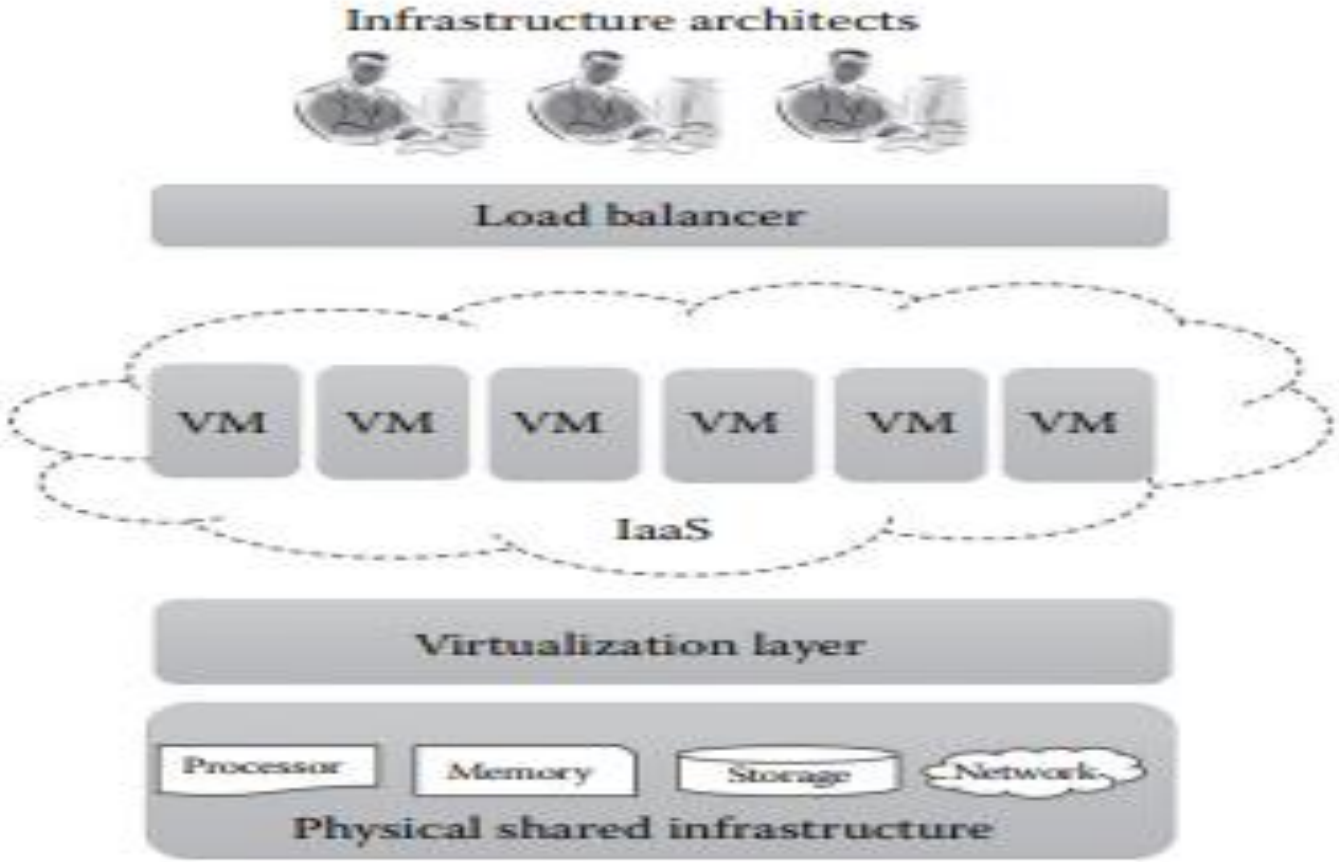

- (K.Chandrasekaran, 2015)

Virtualization in Cloud services

- Platform as a service
- Utilizes the OS-level, database-level, programming language–level virtualization to provide the virtual development platform to the end users. PaaS providers will provide a variety of client tools such as WebCLI, REST APIs, and Web UI to the developers for accessing the virtual platform.

Software as a service

- Utilizes application-level virtualization to deploy the application. The SaaS application allows multiple customers to share the same instance of an application. This technology is popularly known as multitenancy. The scalability of the application will be increased by the software load balancer, which will transfer the additional load to the new application/database server.
- Virtualization is used as an enabling technology to provide multitenant infrastructure in platform, and SaaS.

Virtualization in Cloud services

- Other cloud services that use virtualization
- Network as a Service using network virtualization; Storage as a Service using storage virtualization; and Database as a Service using database virtualization.

Advantages of virtualization

- Managed execution and isolation allow building secure and controllable computing environments. A virtual execution environment can be configured as a sandbox, thus preventing any harmful operation to cross the borders of the virtual host.
- Allocation of resources and their partitioning among different guests is simplified

Advantages of virtualization

- Portability -Virtual machine instances are normally represented by one or more files that can be easily transported with respect to physical systems.
- Through virtualization it is possible to achieve a more efficient use of resources. Multiple systems can securely coexist and share the resources of the underlying host, without interfering with each other

Disadvantages

Performance degradation: The causes of performance degradation can be traced back to the overhead introduced by the following activities:

- Maintaining the status of virtual processors
- Support of privileged instructions (trap and simulate privileged instructions)
- Support of paging within VM
- Console functions

Disadvantages

- Inefficiency and degraded user experience Virtualization can sometime lead to an inefficient use of the host. In particular, some of the specific features of the host cannot be exposed by the abstraction layer and then become inaccessible.
- Security holes and new threats: Virtualization opens the door to a new and unexpected form of phishing. In the case of hardware virtualization, malicious programs can preload themselves before the operating system and act as a thin virtual machine manager.

Next Lesson

- Cloud infrastructure and Architecture

Referencing

- Chandrasekaran, K. (2015). *Essentials of Cloud computing.* USA: Taylor and Francis Group.
- Gowda, T. (2021). OVERVIEW OF VIRTUALIZATION IN CLOUD COMPUTING. *International Research Journal of Modernization in Engineering Technology and Science* , 1-6.
- Rajkumar, B. (2013). *Mastering Cloud Computing.* USA: Morgan Kaufmann.
- Sodhi, B. S. (2017). *Virtualization and cloud computing.* US: CC BY SA.