
USN

Internal Assessment Test I – Nov 2024

Sub:	**Cloud Computing**					Sub Code:	**BAD 515C**	Branch:	**AINDS**		
Date:	**06/11/2024**	Duration:	**90 minutes**	Max Marks:	**50**	Sem	**V**		**OBE**		
Answer any FIVE Questions									**MARKS**	**CO**	**RBT**
1		Explain the relationship between distributed systems and cloud computing.							10	CO1	L2
2	a	Briefly define the following: Service-Oriented Architecture							5	CO1	L1
	b	Cyber-physical system							5	CO1	L1

3		What is virtualization. Discuss different levels of virtualization implementation.	10	CO2	L2
4		Explain the IaaS, PaaS, and SaaS cloud service models.	10	CO3	L2
5		Discuss the key features/offerings of any public cloud, Amazon Web Services or Microsoft Azure or Google Cloud Platform.	10	CO3	L3
6		Explain virtualization of CPU, memory, and I/O devices	10	CO2	L3

CI CCI

HOD

1.Explain the relationship between distributed systems and cloud computing.
Ans:

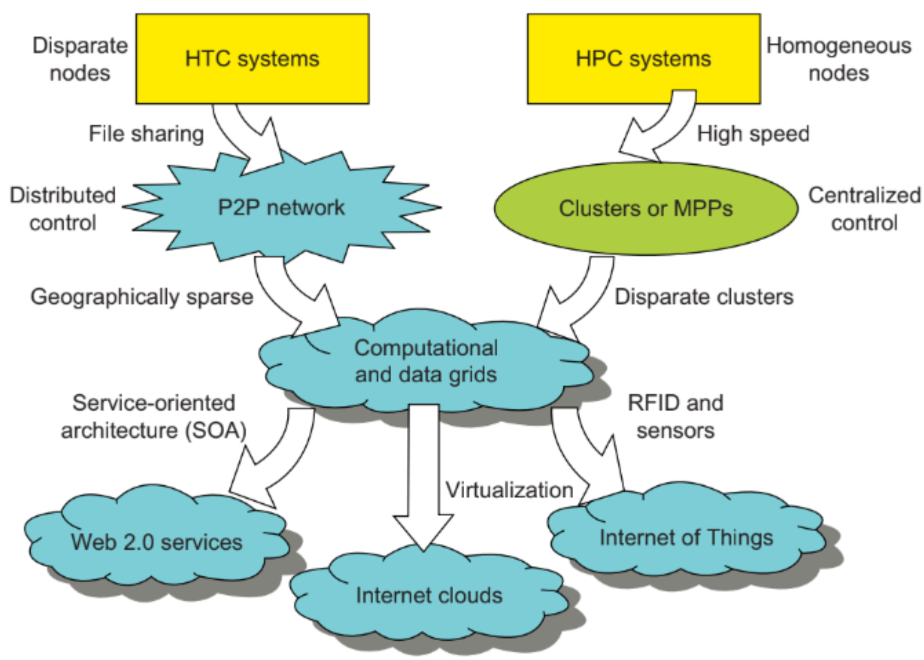

FIGURE 1.1

Evolutionary trend toward parallel, distributed, and cloud computing with clusters, MPPs, P2P networks, grids, clouds, web services, and the Internet of Things.

overlaps with distributed computing to a great extent, and *cloud computing* overlaps with distributed, centralized, and parallel computing. The following list defines these terms more clearly; their architectural and operational differences are discussed further in subsequent chapters.

- **Centralized computing** This is a computing paradigm by which all computer resources are centralized in one physical system. All resources (processors, memory, and storage) are fully shared and tightly coupled within one integrated OS. Many data centers and supercomputers are *centralized systems*, but they are used in parallel, distributed, and cloud computing applications [18,26].
- **Parallel computing** In parallel computing, all processors are either tightly coupled with centralized shared memory or loosely coupled with distributed memory. Some authors refer to this discipline as *parallel processing* [15,27]. Interprocessor communication is accomplished through shared memory or via message passing. A computer system capable of parallel computing is commonly known as a *parallel computer* [28]. Programs running in a parallel computer are called *parallel programs*. The process of writing *parallel programs* is often referred to as *parallel programming* [32].
- **Distributed computing** This is a field of computer science/engineering that studies distributed systems. A *distributed system* [8,13,37,46] consists of multiple autonomous computers, each having its own private memory, communicating through a computer network. Information exchange in a distributed system is accomplished through *message passing*. A computer program that runs in a distributed system is known as a *distributed program*. The process of writing distributed programs is referred to as *distributed programming*.
- **Cloud computing** An *Internet cloud* of resources can be either a centralized or a distributed computing system. The cloud applies parallel or distributed computing, or both. Clouds can be built with physical or virtualized resources over large data centers that are centralized or distributed. Some authors consider cloud computing to be a form of *utility computing* or *service computing* [11,19].

2a. SOA

https://aws.amazon.com/what-is/service-oriented-architecture/

What is service-oriented architecture?

Service-oriented architecture (SOA) is a method of software development that uses software components called services to create business applications. Each service provides a business capability, and services can also communicate with each other across platforms and languages. Developers use SOA to reuse services in different systems or combine several independent services to perform complex tasks.

For example, multiple business processes in an organization require the user authentication functionality. Instead of rewriting authentication code for all business processes, you can create a single authentication service and reuse it for all applications. Similarly, almost all systems across a healthcare organization, such as patient management systems and electronic health rec (EHR) systems, need to register patients. These systems can call a single, common service to perform the patient registration task.

What are the benefits of service-oriented architecture?

Service-oriented architecture (SOA) has several benefits over the traditional monolithic architectures in which all processes ru a single unit. Some major benefits of SOA include the following:

Faster time to market

Developers reuse services across different business processes to save time and costs. They can assemble applications much fa with SOA than by writing code and performing integrations from scratch.

Efficient maintenance

It's easier to create, update, and debug small services than large code blocks in monolithic applications. Modifying any service SOA does not impact the overall functionality of the business process.

Greater adaptability

SOA is more adaptable to advances in technology. You can modernize your applications efficiently and cost effectively. For example, healthcare organizations can use the functionality of older electronic health record systems in newer cloud-based applications.

2b. Cyber-physical system

A *cyber-physical system* (CPS) is an embedded system which integrates the computing process with the physical world as an interactive and intelligent system. CPSes appear in many computer and TV game systems. The best example is the popular Nintendo Wii interactive gaming system. Cursor-based CPSes can be found in the automotive, aerospace, health care, robotics, manufacturing, battlefield training, and consumer appliance industries, among others. Thus, a full-fledged CPS involves embedded computers, network monitors, and intelligent control of physical processes with humans in the feedback loop. In the physical world, man-machine interactions must be handled in real time by a CPS. Figure 9.23 shows the abstraction architecture of a typical CPS.

A CPS features tight coordination between the system's computational and physical elements. Today, most virtual reality systems tend to focus more on computational elements than on the intense link between the computational and physical elements. To amend this shortcoming, a CPS

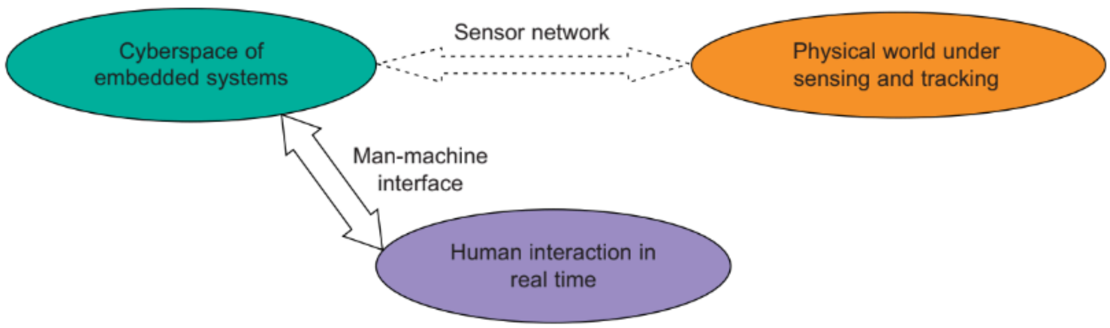

FIGURE 9.23

Three major components working interactively and intelligently in a CPS.

should be designed as a network of interacting elements with physical input and output instead of as a stand-alone device. The notion is extended from the concepts of robotics and sensor networks. The expectation from the CPS community is to improve the link between computational and physical elements. The design objectives include the adaptability, autonomy, efficiency, functionality, reliability, safety, and usability of future CPSes.

These three components interact with one another heavily. Sensor networks and man-machine interfaces need to be established to glue the three sectors together. In a CPS, the cyberspace and physical worlds are intertwined. The embedded computing, sensed real-world data, and real-time responses are all equally important. New understandings and techniques are needed to deal with this new fusion. New advances are desired to broaden the potential of CPSes in several dimensions. One is human intervention at the right moment, such as collision avoidance in driving a vehicle. Doctors need precision in robotic surgery and workers need nano-level manufacturing.

It is desired to design CPSes to replace humans in dangerous or inaccessible environments such as battlefields, rescue from earthquake rubble, and deep-sea exploration. Users demand close coordination in air traffic control and war fighting and enhanced efficiency and augmentation of human capabilities in health care monitoring and life-saving operations. Since 2006, the U.S. National Science Foundation (NSF) has identified CPS as a key area of research. The NSF and other federal agencies have sponsored several workshops on CPS in recent years.

3a. What is virtualization. Discuss different levels of virtualization implementation.

Virtualization is technology that you can use to create virtual representations of ser machines. Virtual software mimics the functions of physical hardware to run multipl physical machine. Businesses use virtualization to use their hardware resources effici investment. It also powers cloud computing services that help organizations manag

What is virtualization?

To properly understand Kernel-based Virtual Machine (KVM), you first need to understand some basic concepts in *virtualization*. Virtualization is a process that allows a computer to share its hardware resources with multiple digitally separated environments. Each virtualized environment runs within its allocated resources, such as memory, processing power, and storage. With virtualization, organizations can switch between different operating systems on the same server without rebooting.

Virtual machines and hypervisors are two important concepts in virtualization.

Virtual machine

A *virtual machine* is a software-defined computer that runs on a physical computer with a separate operating system and computing resources. The physical computer is called the *host machine* and virtual machines are *guest machines*. Multiple virtual machines can run on a single physical machine. Virtual machines are abstracted from the computer hardware by a hypervisor.

Hypervisor

The *hypervisor* is a software component that manages multiple virtual machines in a computer. It ensures that each virtual machine gets the allocated resources and does not interfere with the operation of other virtual machines. There are two types of hypervisors.

Type 1 hypervisor

A type 1 hypervisor, or bare-metal hypervisor, is a hypervisor program installed directly on the computer's hardware instead of the operating system. Therefore, type 1 hypervisors have better performance and are commonly used by enterprise applications. KVM uses the type 1 hypervisor to host multiple virtual machines on the Linux operating system.

Type 2 hypervisor

Also known as a hosted hypervisor, the type 2 hypervisor is installed on an operating system. Type 2 hypervisors are suitable for end-user computing.

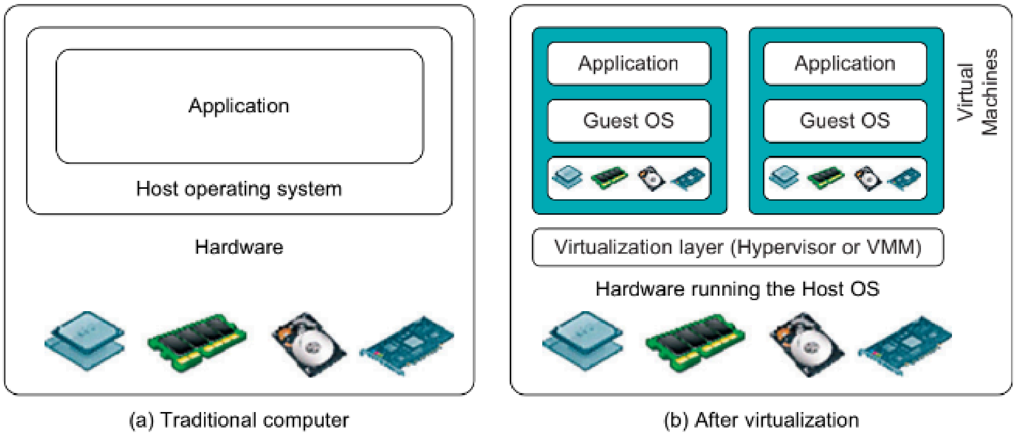

FIGURE 3.1

The architecture of a computer system before and after virtualization, where VMM stands for virtual machine monitor.

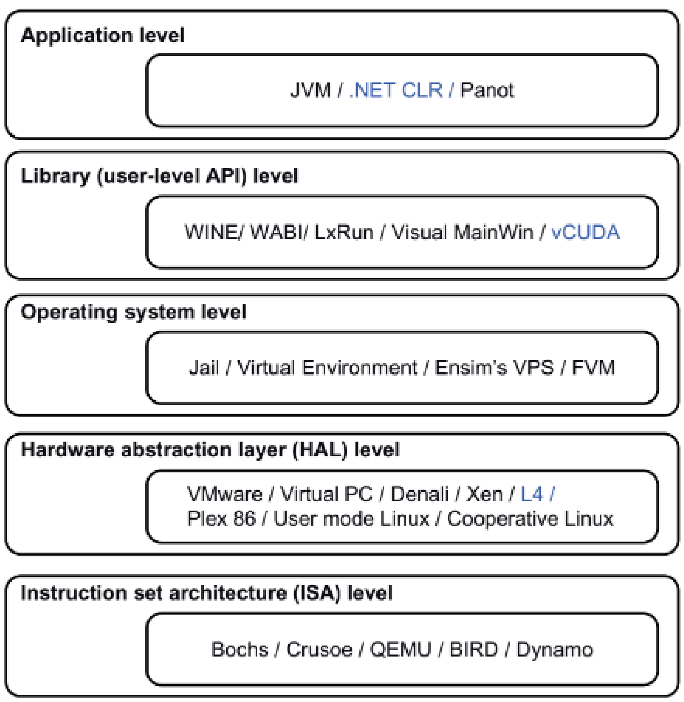

FIGURE 3.2

Virtualization ranging from hardware to applications in five abstraction levels.

4 Explain the IaaS, PaaS, and SaaS cloud service models.

- Y cells indicate what you manage
- C cells indicate what the cloud service provider manages

	Own IT	IaaS	PaaS	SaaS
Data	Y	Y	Y	Y
Application	Y	Y	Y	C
Runtime or software that runs the application	Y	Y	C	C
Middleware or software that monitors the application	Y	Y	C	C
Operating systems on which the application runs	Y	C	C	C
Virtualization technology	Y	C	C	C
Server machines	Y	C	C	C
Storage devices	Y	C	C	C
Network appliances	Y	C	C	C

IaaS pros

- Highest level of control over infrastructure
- On-demand scalability
- No single point of failure for higher reliability
- Reduced upfront capital expenditures (for example, pay-as-you-go pricing)
- Fewer provisioning delays and wasted resources
- Accelerated development and time to market

IaaS cons

- Responsible for your own data security
- Requires hands-on configuration and m
- Difficulties securing legacy application

PaaS pros	PaaS cons
• Instant access to a complete, easy-to-use development platform • Cloud service provider is responsible for maintenance and securing infrastructure • Available over any internet connection on any device • On-demand scalability	• Application stack can be limited to the mos components • Vendor lock-in may be an issue depending service provider • Less control over operations and the overa • More limited customizations
SaaS pros	**SaaS cons**
• Easy to set up and start using • The provider manages and maintains everything, from hardware to software • Software is accessible over any internet connection on any device	• No control over any of the infrastructure o • Integration issues with your existing tools a • Vendor lock-in may be an issue depending service provider • Little to no customization

5. Discuss the key features/offerings of any public cloud, Amazon Web Services

https://www.geeksforgeeks.org/what-are-the-important-aws-cloud-services/

6. Explain virtualization of CPU, memory, and I/O devices

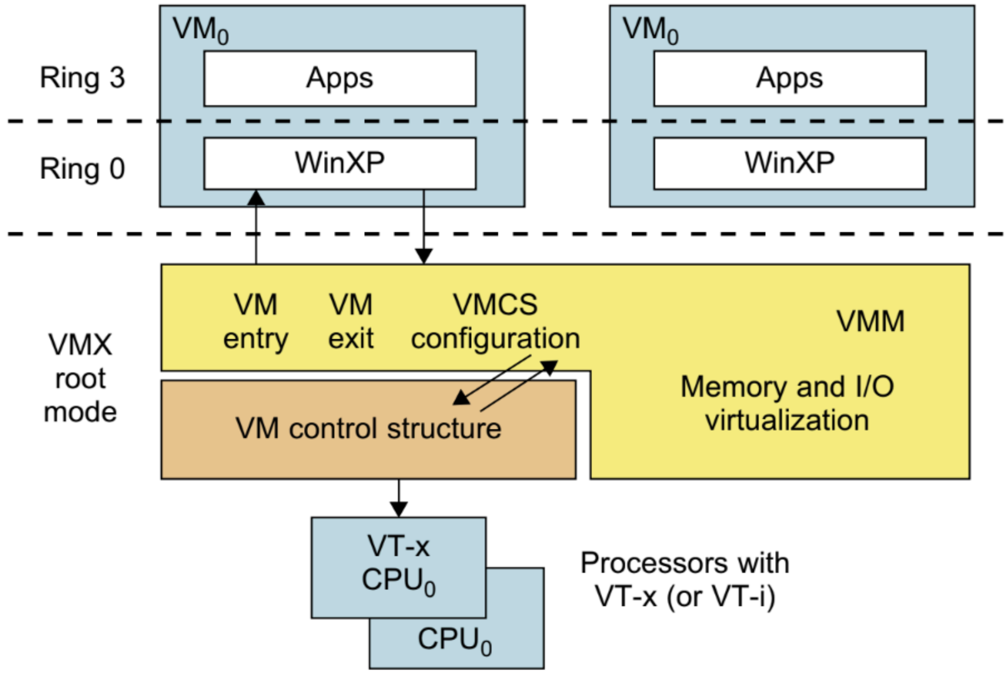

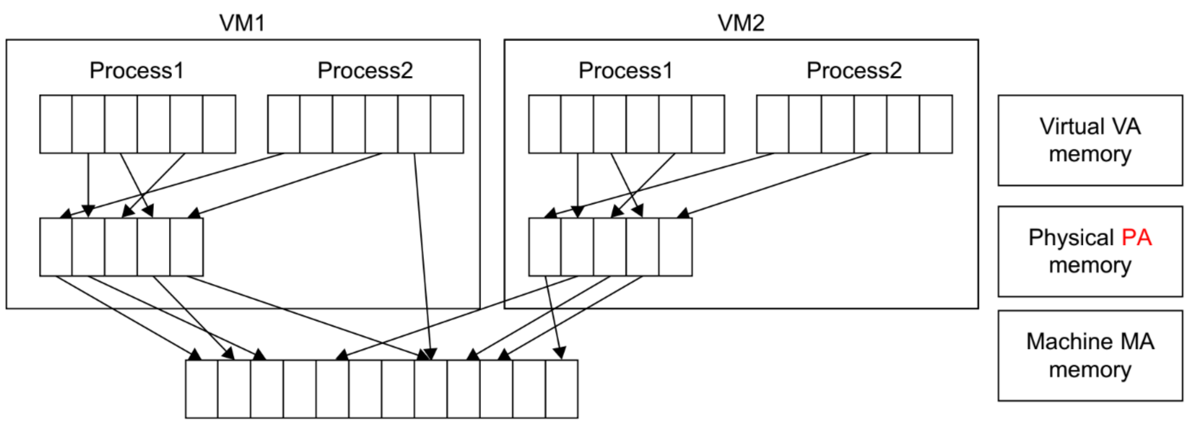

FIGURE 3.12

Two-level memory mapping procedure.

(*Courtesy of R. Rblig, et al. [68]*)

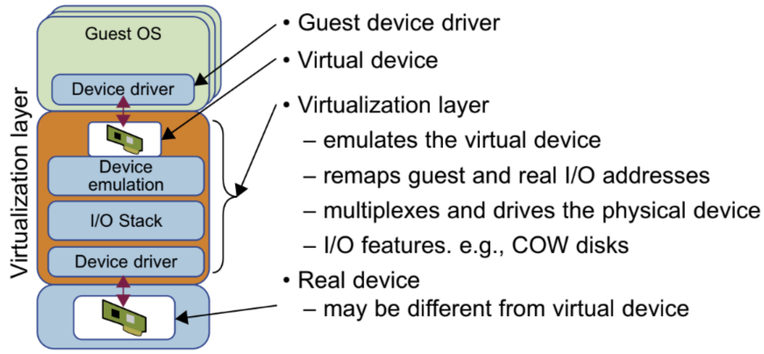

FIGURE 3.14

Device emulation for I/O virtualization implemented inside the middle layer that maps real I/O devices into the virtual devices for the guest device driver to use.

(*Courtesy of V. Chadha, et al. [10] and Y. Dong, et al. [15]*)