



云计算入门

Introduction to Cloud Computing GESC1001

Philippe Fournier-Viger

Professor

School of Humanities and Social Sciences

philfv8@yahoo.com

Fall 2020





CLOUD RESOURCE VIRTUALIZATION

(资源虚拟化)

Based on Chapter 5


Course schedule

Part 1	Introduction and overview
Part 2	Distributed and parallel systems
Part 3	Cloud infrastructure
Part 4	Cloud application paradigm (1)
Part 5	Cloud application paradigm (2)
Part 6	Cloud virtualization and resource management
Part 7	Cloud computing storage systems
Part 8	Cloud computing security
	Final exam

Introduction

A computer has three important resources:

1) A **processor** (处理器), to perform calculations,



2) Memory (计算机存储器)


Memory, to store data

Hard drive (硬盘驱动器)



For long-term storage,
but accessing data is slow

RAM (内存)




For short-term storage,
but accessing data is fast

3) Communication channels (通信通道)

To communicate with other computers

local network

Internet



3) Communication channels (通信通道)

Different types:

« Telephone » cable (RJ11)




RJ45
Ethernet cable




TV (电视) cable



Optical fiber
(光学纤维)



WIFI
(无线网络)



3) Communication channels (通信通道)

They have different characteristics:
bandwidth (带宽), latency (延迟), cost (元)

« Telephone » cable (RJ11)




RJ45
Ethernet cable




TV (电视) cable



Optical fiber
(光学纤维)



WIFI
(无线网络)

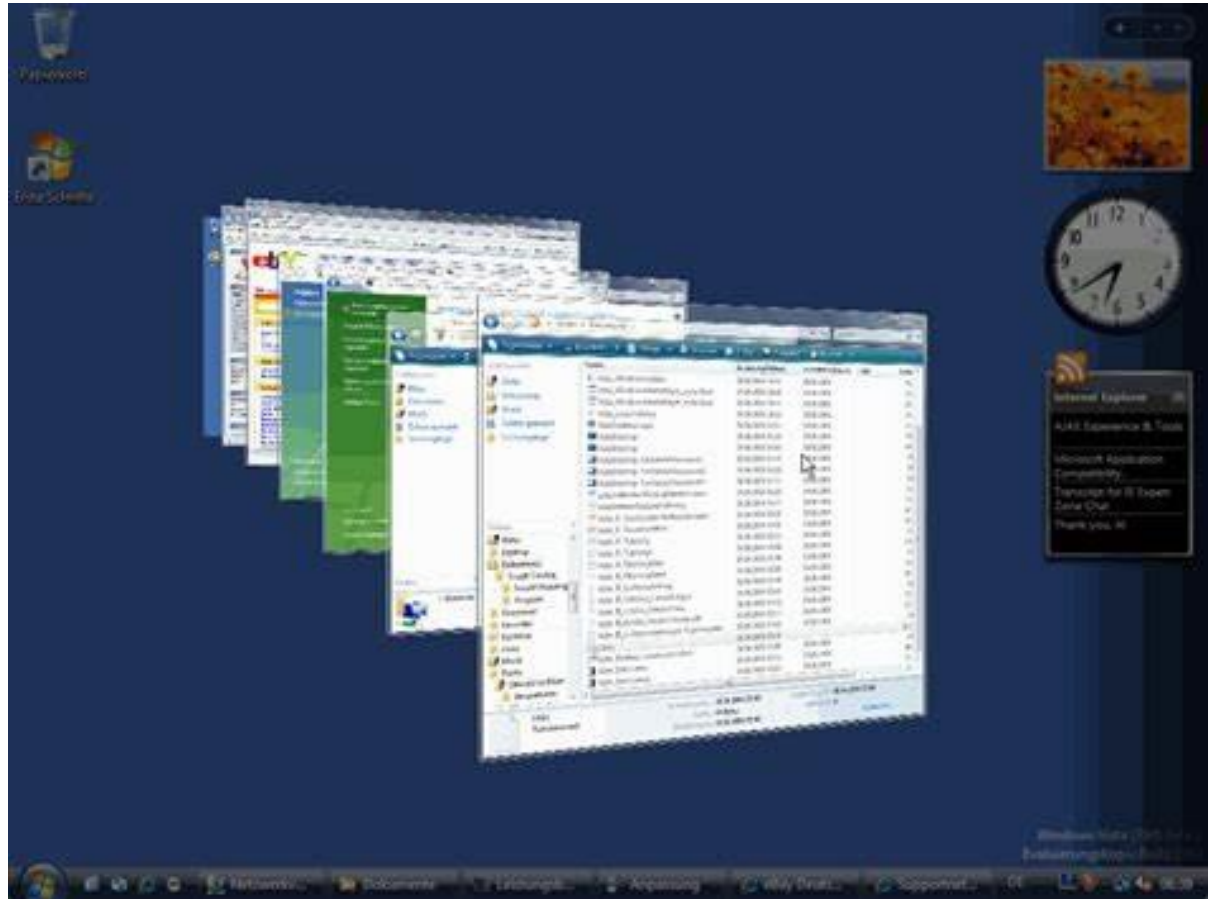


Operating system (操作系统)

- Each computer runs an **operating system** (操作系统)
(e.g. Windows, Android, Linux, MacOS).
- **Why do we need an operating system?**
→

Why do we need an operating system?

1) to provide a **user interface** (用户界面) to the user,



This is a *graphical user interface*

Why do we need an operating system?

1) to provide a **user interface** (用户界面) to the user,

```
Enter today's date (m-d-y): 08-04-81

The IBM Personal Computer DOS
Version 1.00 (C)Copyright IBM Corp 1981

A>dir *.com
IBMBIO      COM           1920  07-23-81
IBMDOS      COM           6400  08-13-81
COMMAND     COM           3231  08-04-81
FORMAT      COM           2560  08-04-81
CHKDSK      COM           1395  08-04-81
SYS         COM            896  08-04-81
DISKCOPY    COM           1216  08-04-81
DISKCOMP    COM           1124  08-04-81
COMP        COM           1620  08-04-81
DATE        COM            252  08-04-81
TIME        COM            250  08-04-81
MODE        COM            860  08-04-81
EDLIN       COM           2392  08-04-81
DEBUG       COM           6049  08-04-81
BASIC       COM          10880  08-04-81
BASICA      COM          16256  08-04-81

A>_
```

This is a *text-based user interface*

Why do we need an operating system?

2) to manage the applications that are running on the computer.

The screenshot displays a Windows 10 desktop with a purple-themed background. In the foreground, the Task Manager application is open, showing the Performance tab. The Performance tab displays system metrics: CPU at 5% (1.63 GHz), Memory at 31% (1.1/3.6 GB), Disk 0 (C:) at 0%, Wi-Fi at 0 Kbps, and Ethernet as Not connected. The Memory section shows a graph of memory usage over 60 seconds and a table of memory composition. The Task Manager window is also open, showing a list of running processes. The process list includes Google Chrome (22), Microsoft PowerPoint, Windows Audio Device Graph Isolation, Windows Resource Manager (2), Antimalware Service Executable, Desktop Window Manager, MBAMessageCenter.exe, and WeChat (32位) (2). The system tray at the bottom right shows the time as 5:25 PM on 4/22/2013.

名称	状态	CPU	内存	磁盘	网络	GPU
Google Chrome (22)		0.1%	861.1 MB	0.1 MB/秒	0 Mbps	
Microsoft PowerPoint		0%	161.6 MB	0 MB/秒	0 Mbps	
Windows 音频设备图形隔离		0.9%	158.8 MB	0 MB/秒	0 Mbps	
Windows 资源管理器 (2)		0.2%	120.4 MB	0 MB/秒	0 Mbps	
Antimalware Service Executa...		0%	108.8 MB	0 MB/秒	0 Mbps	
桌面窗口管理器		0.3%	107.7 MB	0 MB/秒	0 Mbps	C
MBAMessageCenter.exe		0%	106.6 MB	0.1 MB/秒	0 Mbps	
WeChat (32位) (2)		0.1%	36.7 MB	0 MB/秒	0 Mbps	


Why do we need an operating system?

3) to manage the resources of the computer

Memory



Processing resource



Communication channels



And also the devices such as
Keyboard, mouse, screen...



Resource management (资源管理)

Resource management is **complex** for the cloud (云)。


Why?

- many computers,
- many users,
- many applications,
- many operating systems may be used...
- different kind of computers may be used...
- some computers can fail (电脑故障)


Resource management is discussed in detail in **Chapter 6**.

Traditional solution for a data center (数据中心)

- In a traditional **data center** (数据中心), a computer may be accessed by **many users** at the same time.
- An **operating system** (操作系统) is used for each computer in the data center.
 - to manage and share resources with other computers,
 - to protect applications and data from other applications,
 - to make sure that **performance** is good for **all users**.




Could we use an operating system like Windows 10 for the cloud?



It would not be a good idea!

Because, it would be difficult to do

- system administration (系统管理),
- security (系统安全) – how to protect data?,
- resource management (资源管理),
- accounting (who should pay? and how much (元) for using the cloud?)



A better solution is **resource virtualization** (资源虚拟化)

- **Resource virtualization** (资源虚拟化) is the topic of this chapter.
- It is an approach to make it easy to manage the resources of **many** computers in the cloud.
- **Main idea:** each computer runs many **virtual machines** (VM -虚拟机).

What is a virtual machine (虚拟机)?

- A **virtual machine** (虚拟机):
an application that simulates a real computer, including its processor, memory, and communication channels.
- A **computer** can run many virtual machines at the same time.
- Each **virtual machine** works like a real computer.
- A **virtual machine** has:
 - its own **operating system** (操作系统)(e.g. **Windows**),
 - its own **applications**.

Note: A more general definition of what is a virtual machine will be presented later.

Example

At home, on a Windows computer, it is possible to use **VirtualBox** or similar applications to create and run virtual machines (虚拟机).



Two
Virtual
machines
(虚拟机)
are running


Virtualization

- A **virtual machine** is also called an **instance** (实例) when we talk about cloud computing.
- The process of running virtual machines on a computer is called **virtualization**.

Virtualization in the cloud →

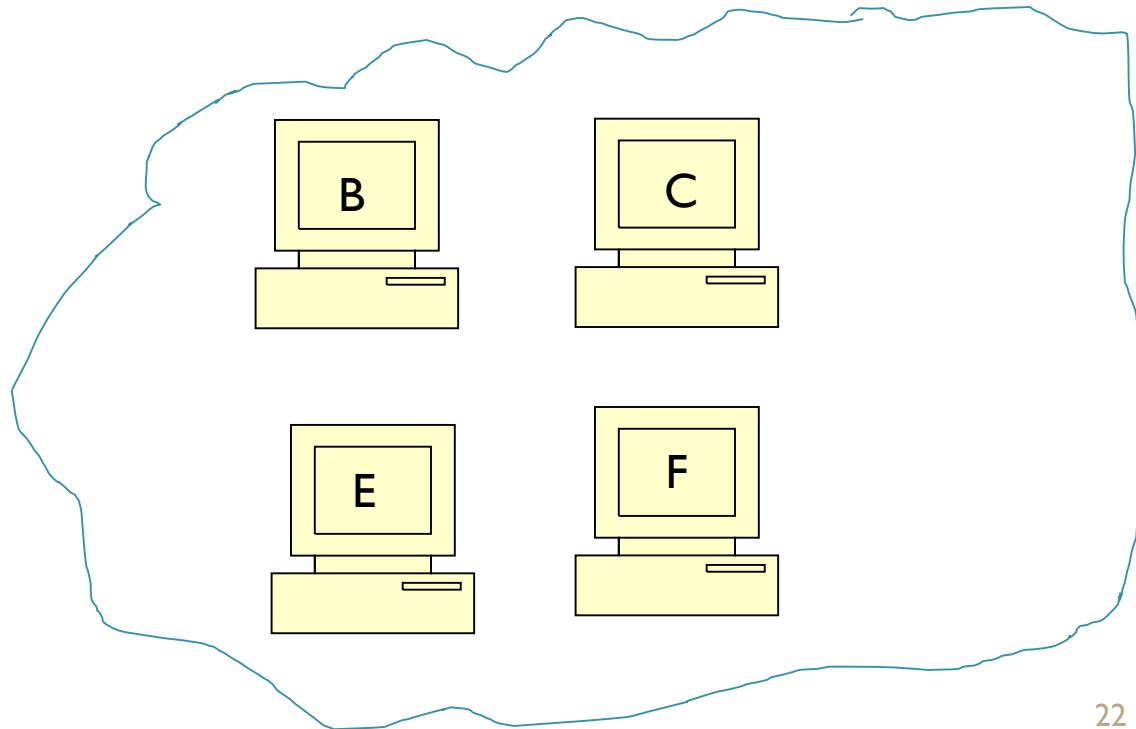
Example

An **image** (虚拟机镜像) is the state of a computer that has been saved into a file



User

Image
AMI




Example



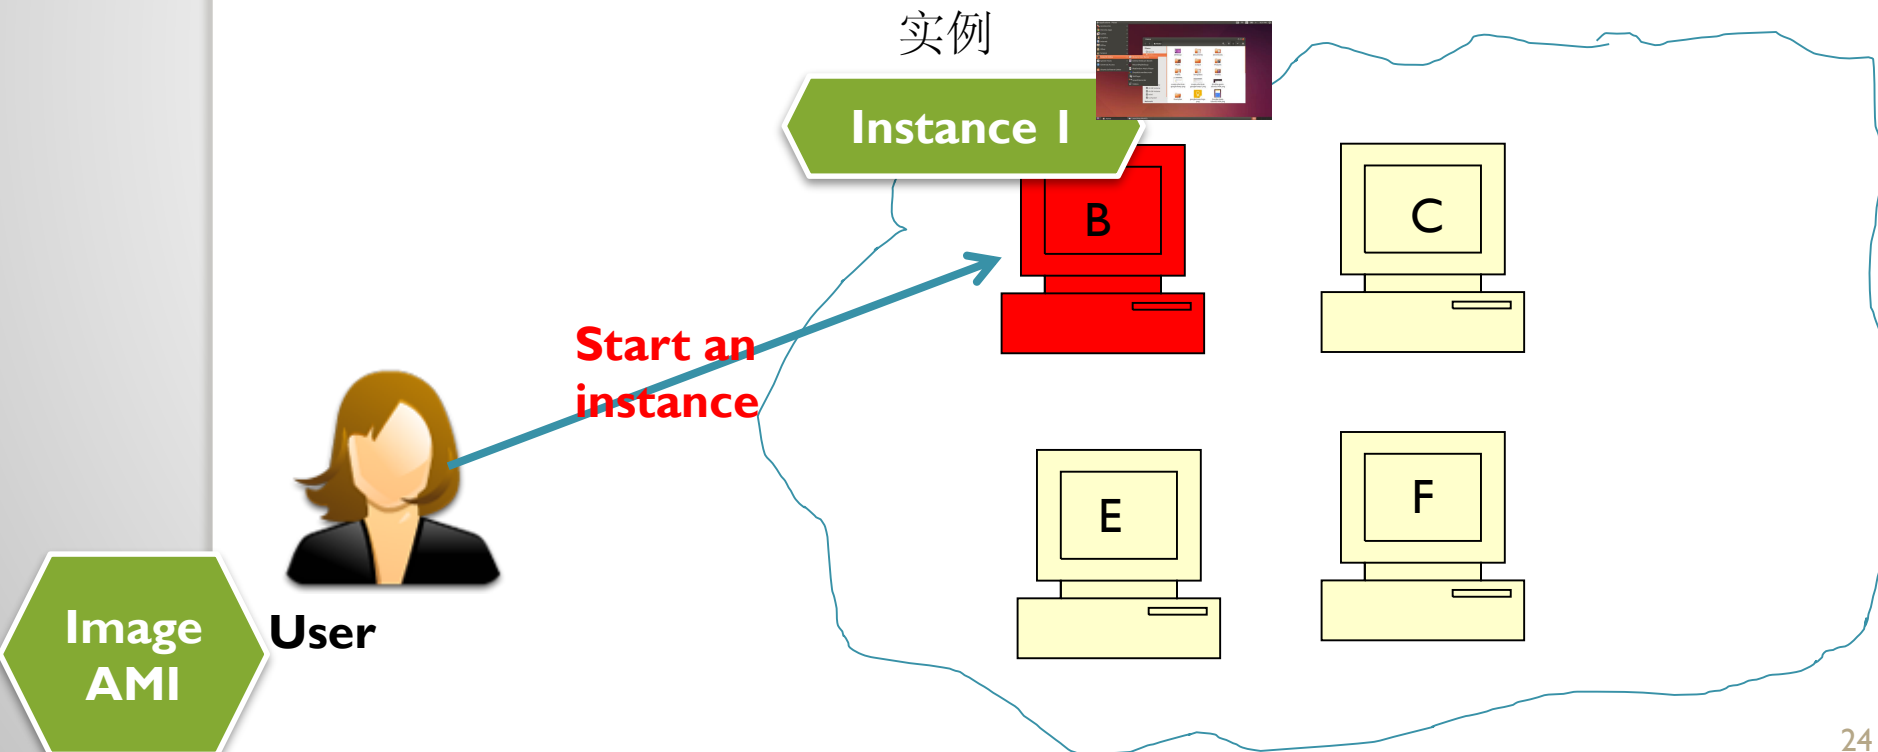
User

**Image
AMI**



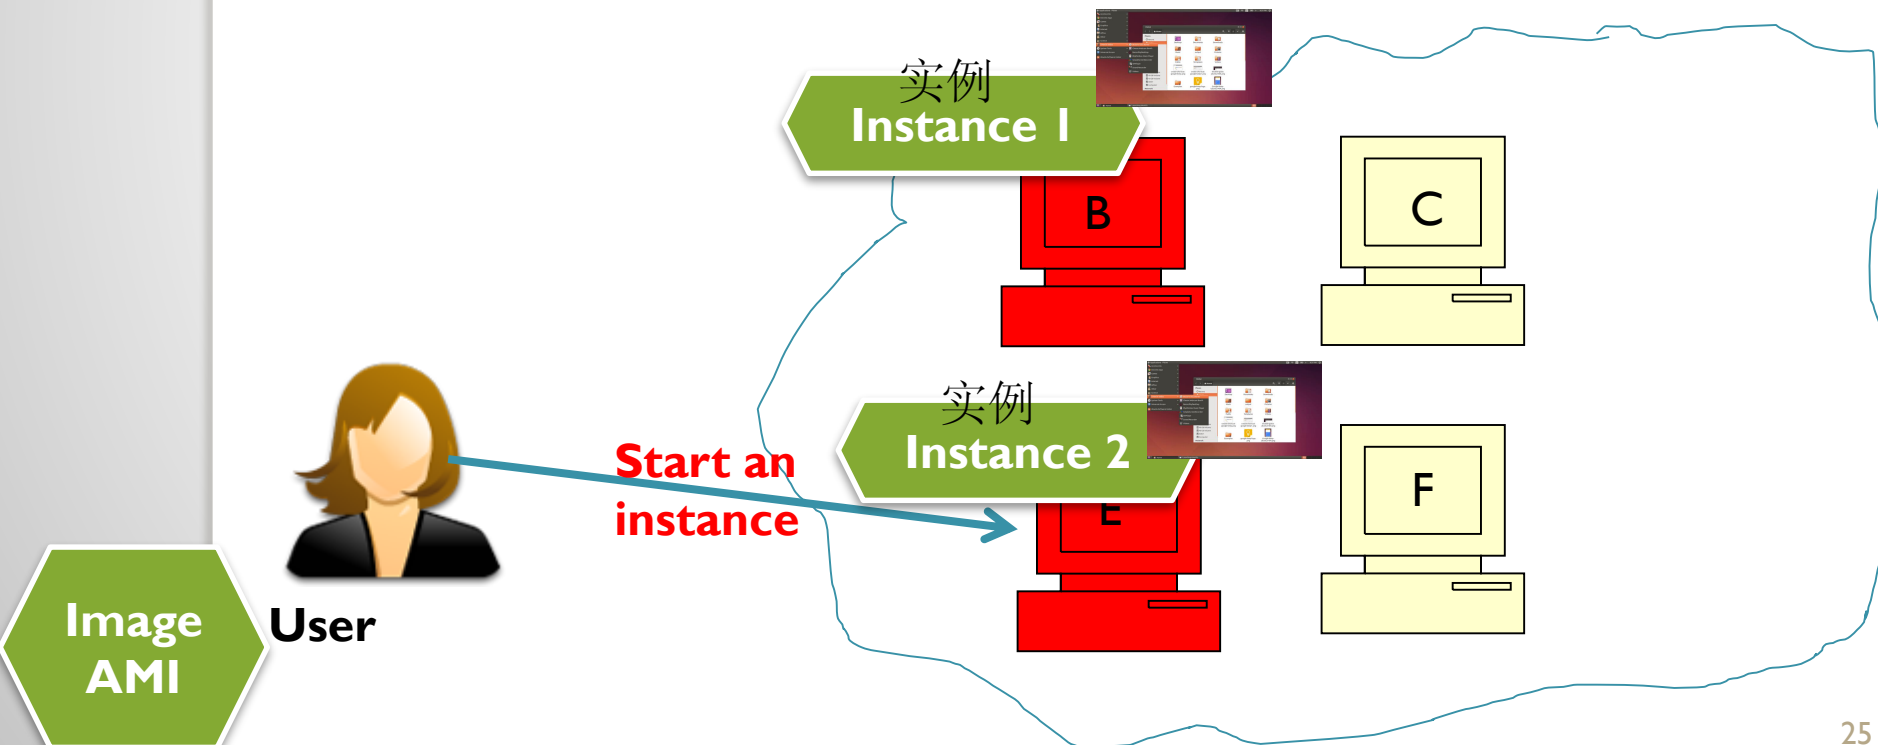
Example

An **image** (虚拟机镜像) can be used to start an **instance** in the cloud (a virtual machine虚拟机)



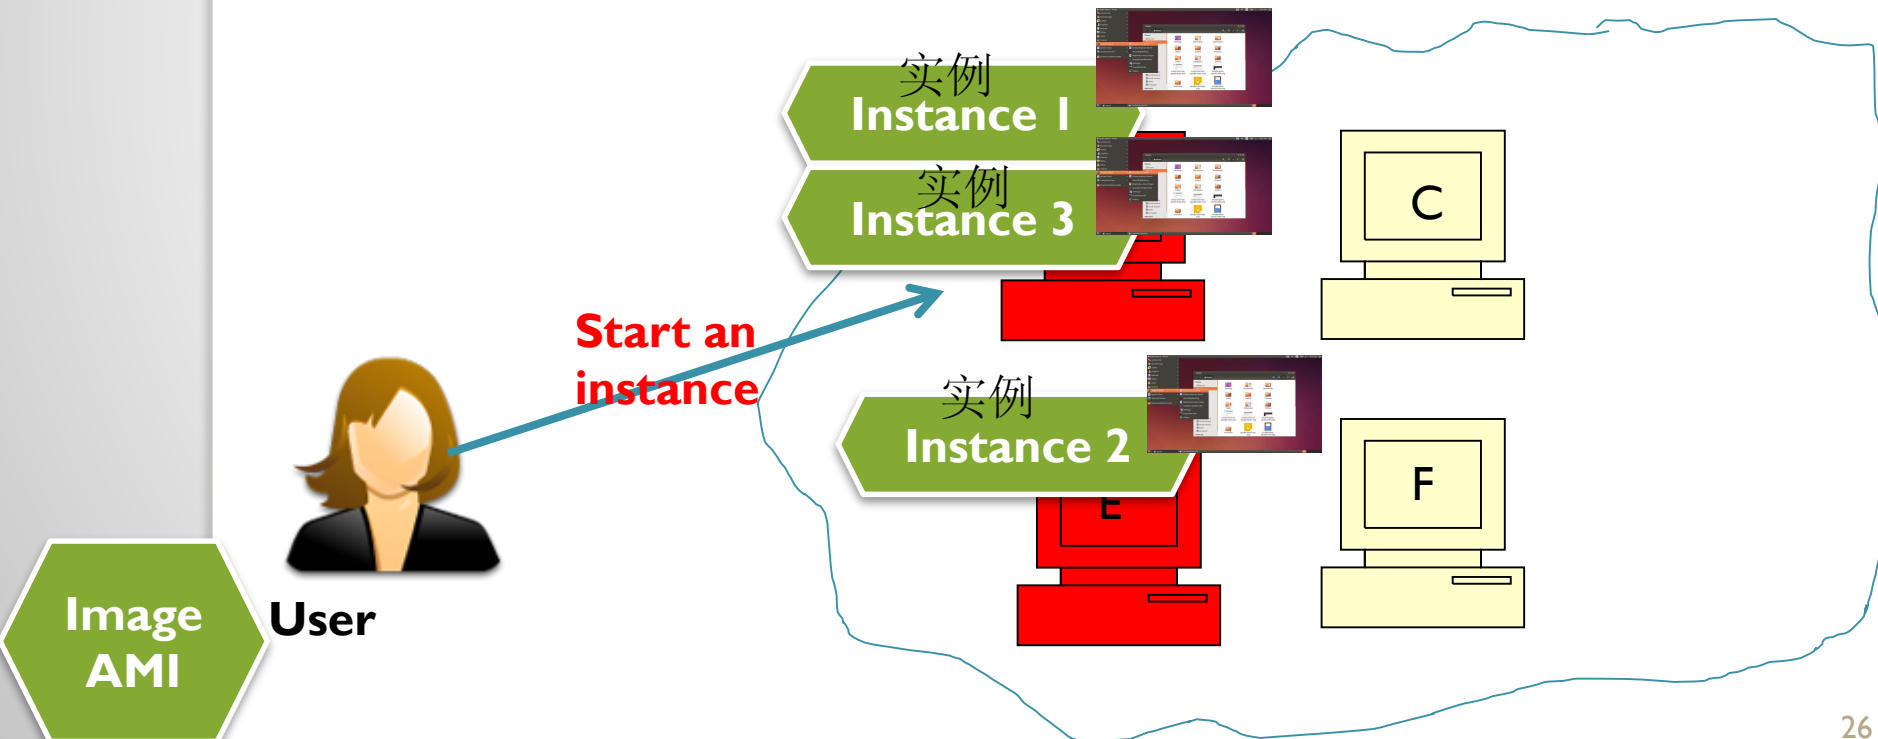
Example

The same image can be used to start many instances



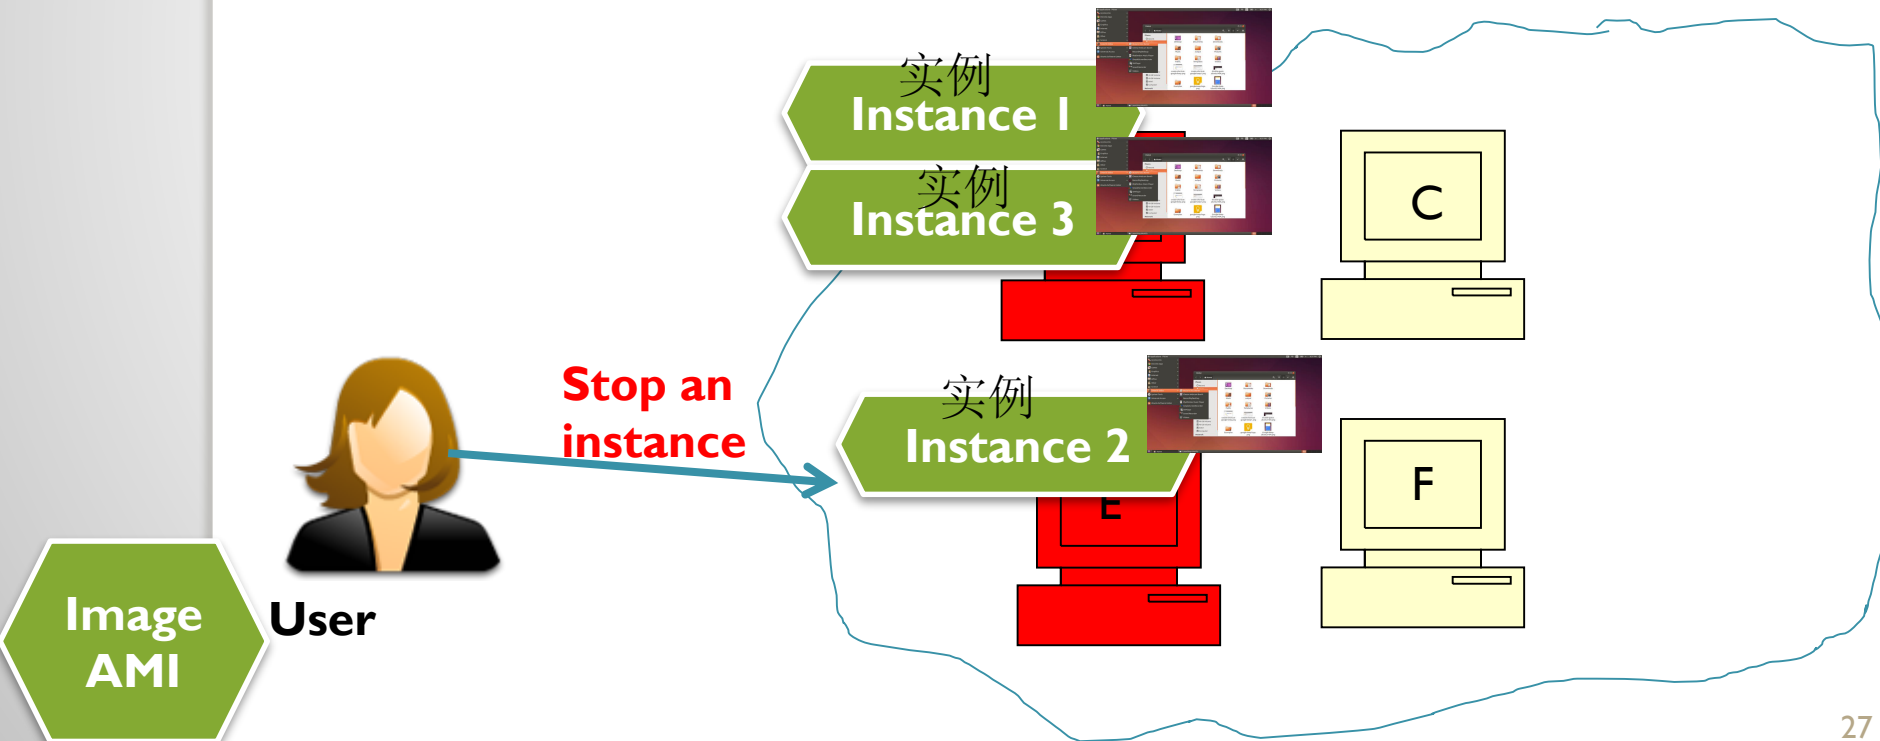
Example

Multiple instances can be run on the same computer



Example

The user can stop an instance.
The user can restart an instance.

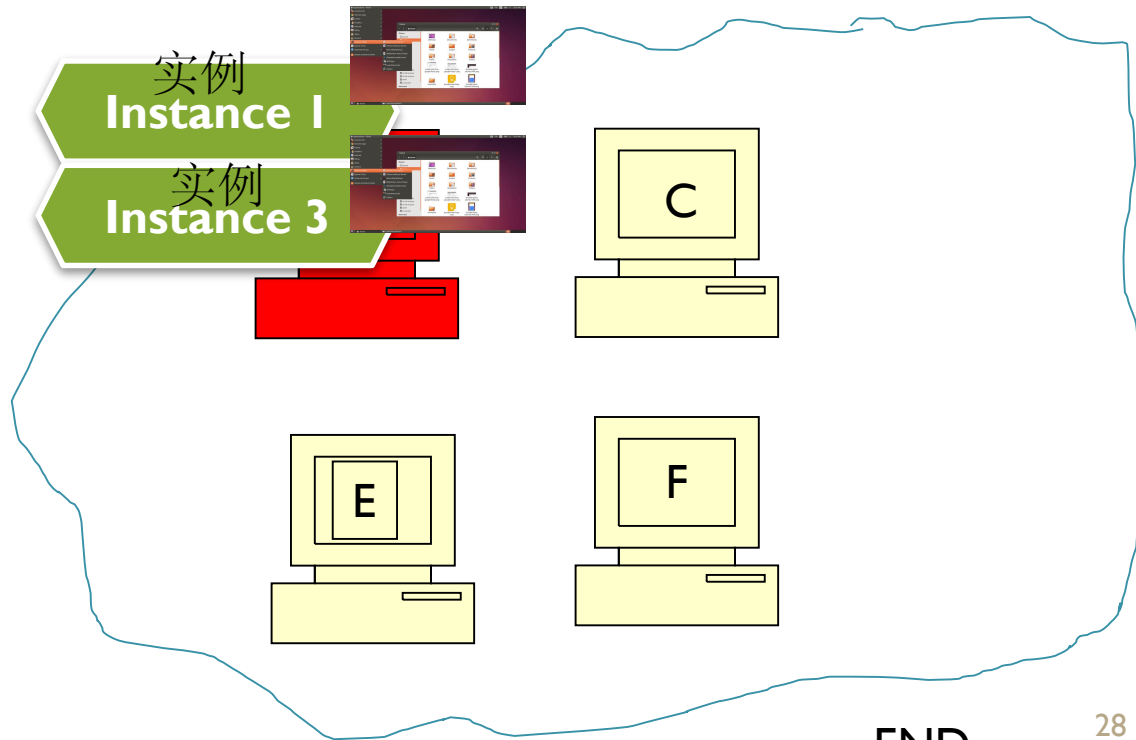


Illustration

Image
AMI



User




END

Why virtualization is useful for the cloud?

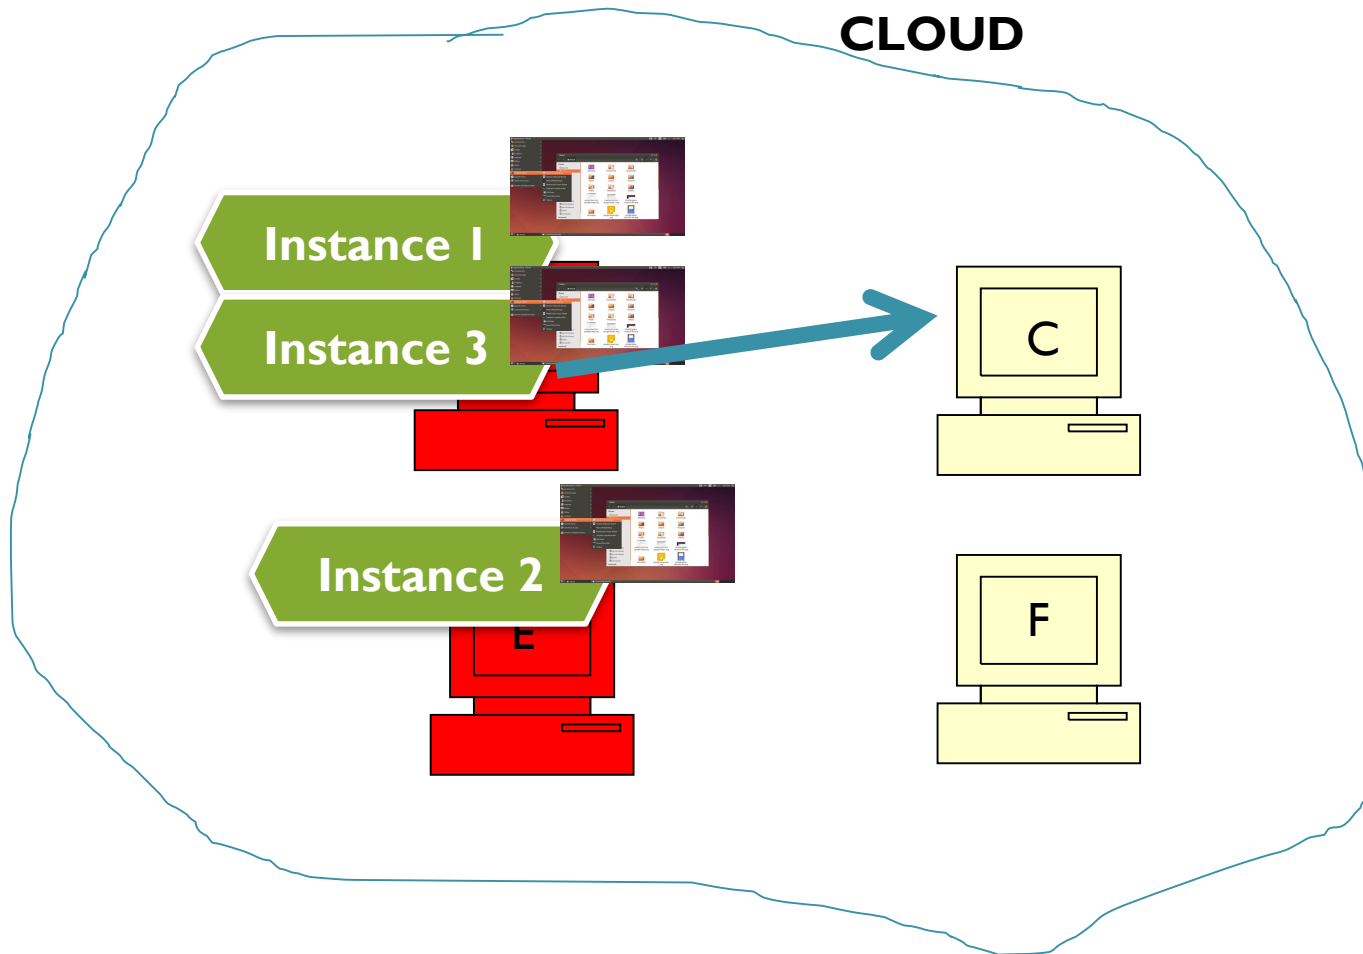
1) Virtualization makes it easy to do **load balancing** (负载均衡):

- Each computer in the cloud can run several virtual machines.
- If a computer is too busy, the **state of some of its virtual machines can be saved**.
- Then, these virtual machines can be **sent to other computers** that are less busy.
- The other computers can continue running these virtual machines.

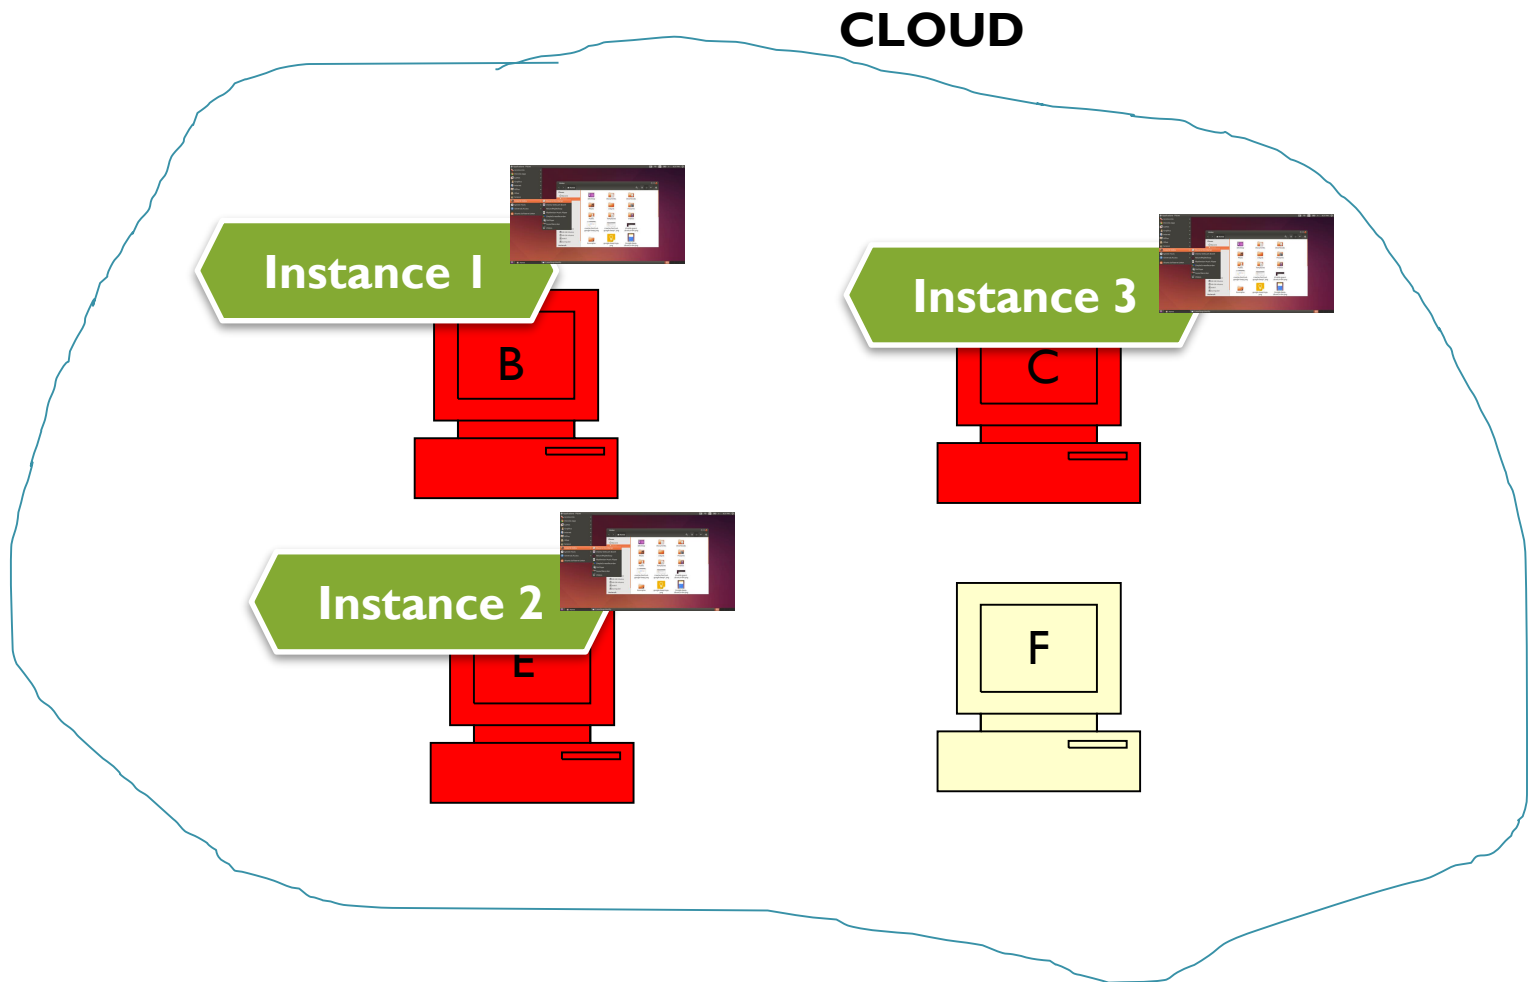
l) makes it easy to do **load balancing** (负载平衡):



l) makes it easy to do **load balancing** (负载平衡):




l) makes it easy to do **load balancing** (负载平衡):



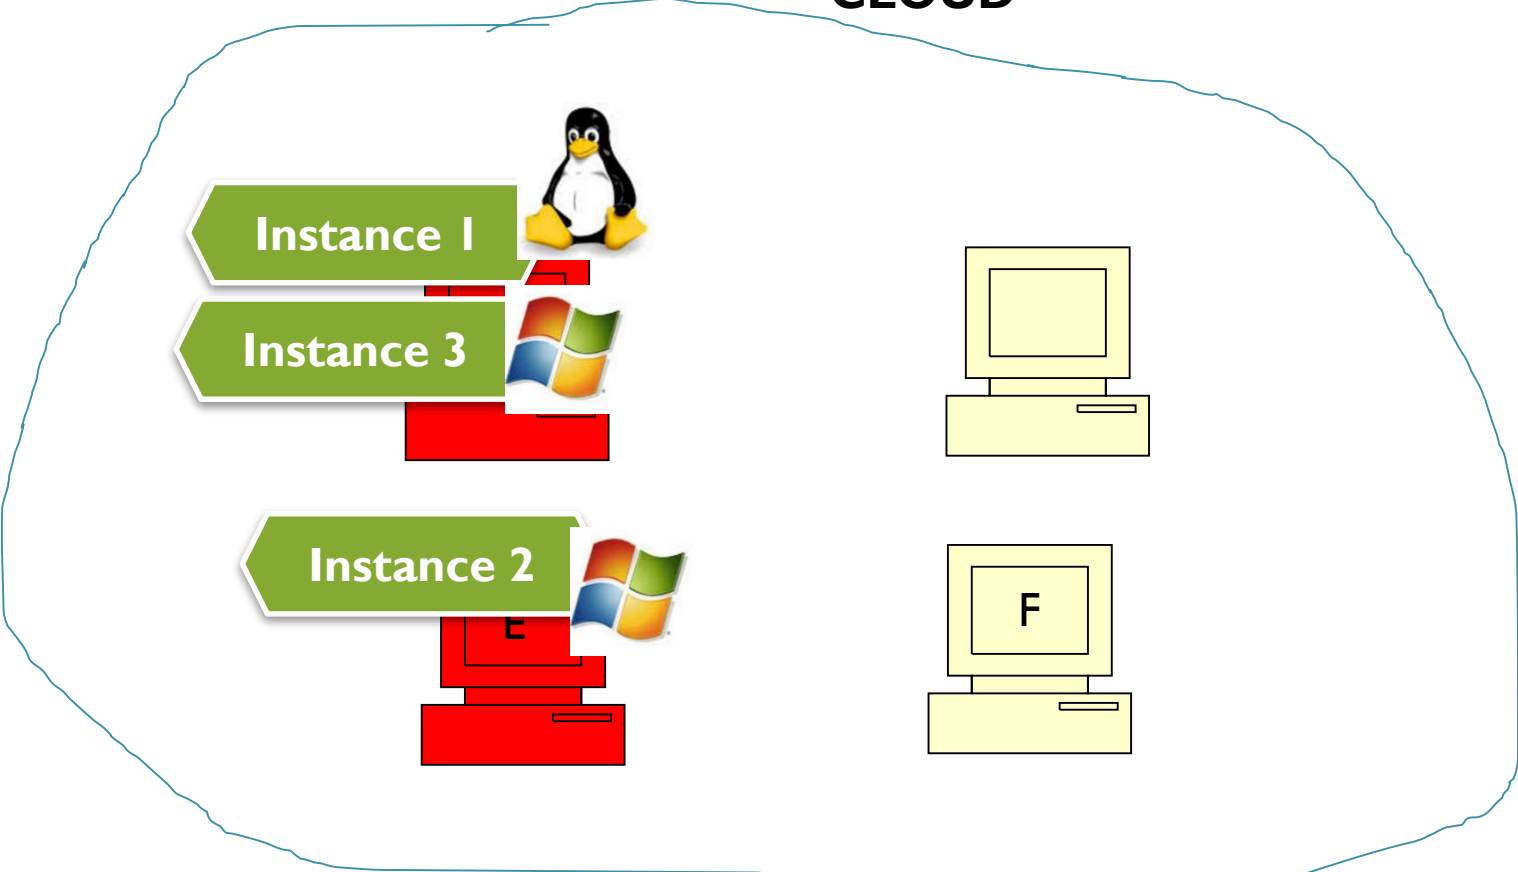
Why virtualization is useful for the cloud?

2) It allows users to use an **operating system** that they like or need:

- e.g. some people prefer using Linux instead of Windows
- e.g. some people want to use some applications that run only on Windows or on Linux
- ...




CLOUD



Why virtualization is useful for the cloud?


3) Using virtual machines increases **security**:

- **All virtual machines** running on a same computer **are isolated from each other** (彼此隔离).
- A user running applications in a virtual machine will **not** be able to **steal data** (窃取数据) in another virtual machine or cause other security problems (**viruses, etc.**).



Why virtualization is useful for the cloud?


- If a **hacker takes control of a virtual machine** on a computer, he will not have access to the other virtual machines.



Why virtualization is useful for the cloud?

4) It is easy to manage computer resources:

- Modern **virtualization software** (applications 应用) allows to define **how much resources** can be used by each virtual machine.
 - processor time,
 - memory,
 - storage space.
- Thus, **resources can be reserved and guaranteed** for each user.



Mary



10 % of the
CPU

Lea



50 % of the
CPU

Paul

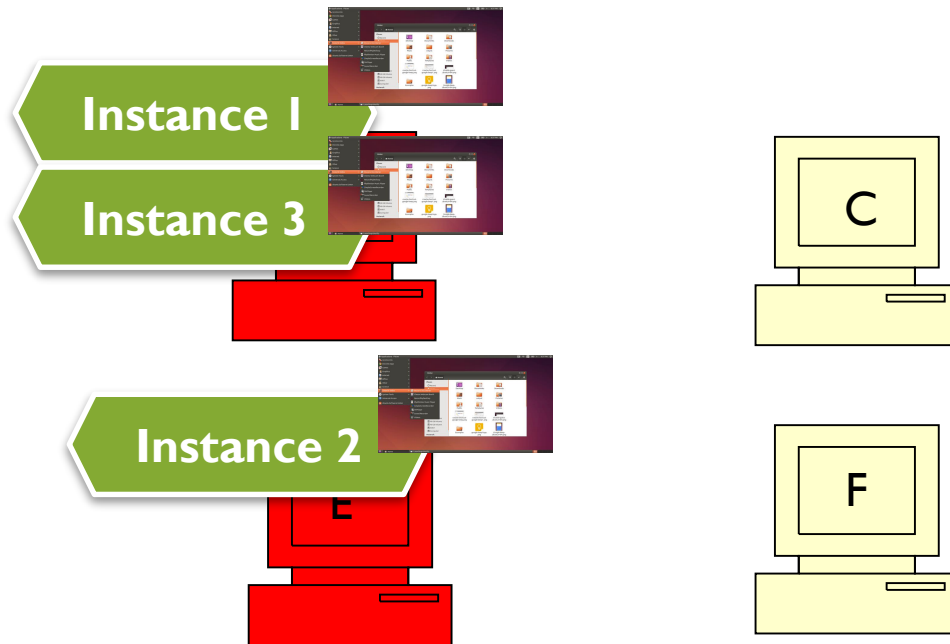


40 % of the
CPU

Why virtualization is useful for the cloud?

5) It allows **elasticity** (弹性计算)

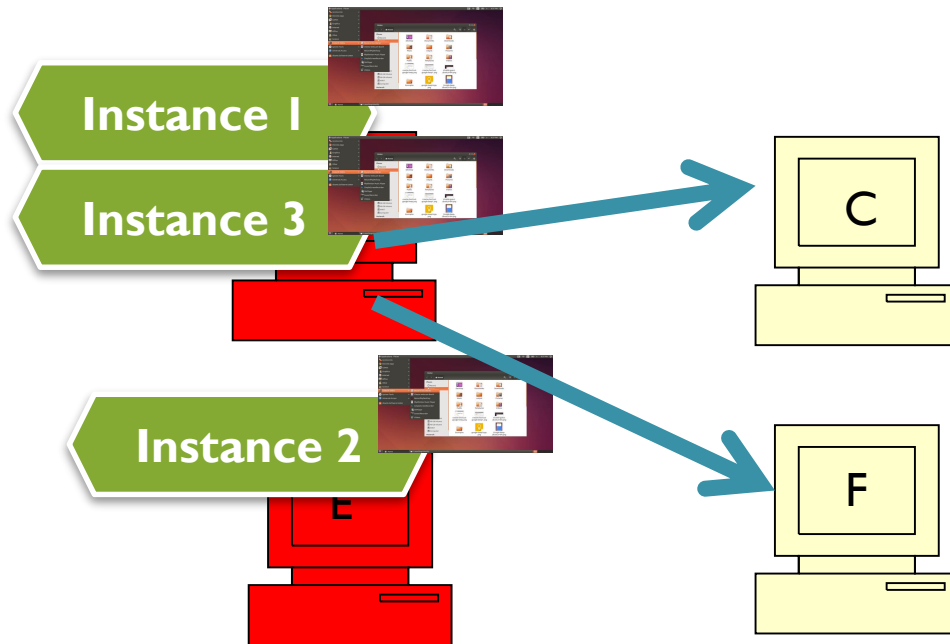
- If virtual machines are too busy or do not have enough resources, the state of a virtual machine can be saved and copied to other computers.
- These computers can run copies of the virtual machine.



Why virtualization is useful for the cloud?

5) It allows **elasticity** (弹性计算)

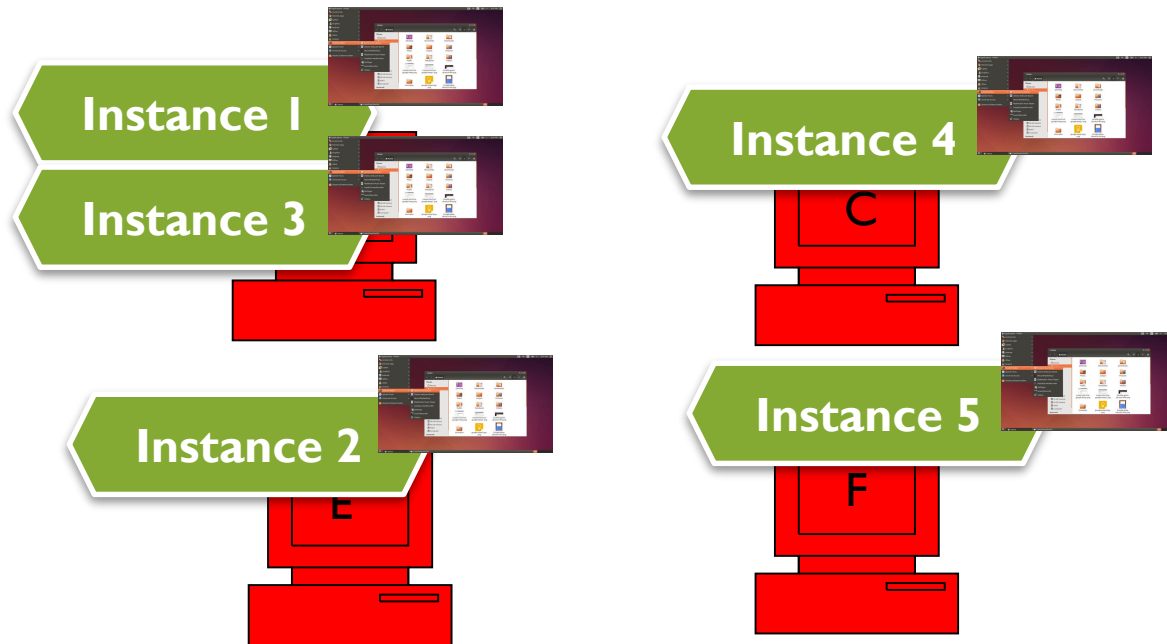
- If virtual machines are too busy or do not have enough resources, the state of a virtual machine can be saved and copied to other computers.
- These computers can run copies of the virtual machine.



Why virtualization is useful for the cloud?

5) It allows **elasticity** (弹性计算)

- If virtual machines are too busy or do not have enough resources, the state of a virtual machine can be saved and copied to other computers.
- These computers can run copies of the virtual machine.



Why virtualization is useful for the cloud?

5) It is easy to **monitor performance** (监视性能)

- The **performance of each virtual machine can be monitored** easily in term of:
 - processor time,
 - network usage,
 - storage space, etc.
- This is **important** for the “**pay for what you use**” model of cloud computing.



www.Vecto.rs · 24920

What is needed to do resource virtualization?

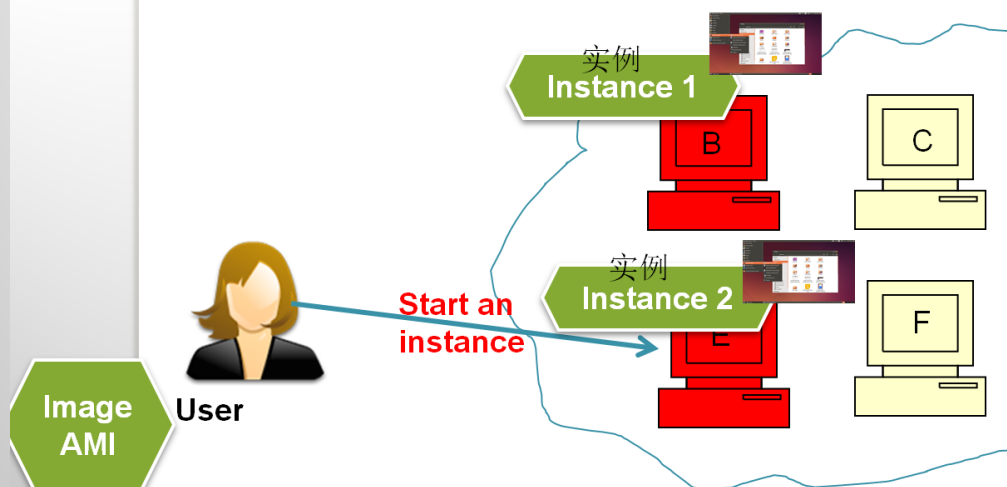
It is useful to have:

- **powerful processors** (强大的处理器),
- **enough memory and storage space,**
- **hardware support** (硬件支持)
(new computers often have specific hardware features to support virtualization),
- ...

Virtualization is convenient for the user

For example, consider the **Amazon EC2 Cloud**.

- A **user** can create an **Amazon Machine Image (AMI)** containing an operating system.
- He can then use the image to create as many virtual machines as he wants in the cloud.
- He use **monitoring** and **management tools** provided by Amazon.



Two approaches for virtualization

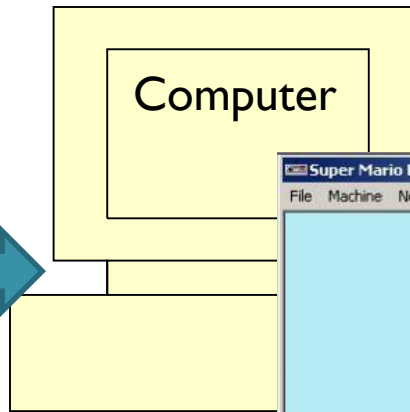
I) Full virtualization (全虚拟化):

- the virtual machine is an exact copy of the **physical hardware** (物理硬件) of a computer.
a typical desktop computer
- any **operating system** that can run on the real hardware can also run in the virtual machine.
Windows, Linux...
- any **applications** for the operating system can be used.
Wechat, QQ,...

Another example

Some video games were once only played on a video game console and a TV

Nowadays, using a virtual machine, we can play these games on a desktop computer or smartphone



Two approaches for virtualization

2) Paravirtualization (半虚拟化):

- the virtual machine does **not** support all functions of the physical hardware.
- thus, the operating system and applications must be **modified** to run on the virtual machine.

Virtual Machine Manager (hypervisor - 虚拟机管理程序)

- **To run multiple machines** at the same time on a computer, a software called a **virtual machine manager (VMM)** is often used.
- **Some are run from an operating system** such as Windows or Linux
 - e.g. **Red Hat Virtual Machine Manager**
- **Others are run directly on the hardware** (硬件)
 - e.g. **Oracle VM server**

Virtual Machine Manager (hypervisor - 虚拟机管理程序)

What can it do?

- *create, start and stop* virtual machines,
- *save or load* a virtual machine,
- *monitor* the performance of virtual machines (in real-time, or over a period of time),

5.1

Virtualization, an old concept...


- It has been first used in the 1950s.
- **1959**: a computer named *Atlas* was using a virtual memory (United Kingdom – 英国)
- **Virtualization** is a general concept that is not only applied to virtual machines.

5.1

Virtualization

Virtualization simulates access to a **physical object** by four ways:

1. **Multiplexing (多路复用)**: Create multiple virtual objects using one physical object.
(e.g. a processor is shared by multiple applications)



Virtualization

2. **Aggregation (聚集)**: Create a virtual object from multiple physical objects.
(e.g. several hard drives are viewed as a single hard drive by a computer)
3. **Emulation (仿真)**: Construct a virtual object from a different type of physical object.
(e.g. the hard drive of a computer is used to simulate the RAM memory of a computer)
4. **A combination of 1 and 3.**

Guest operating system (客户操作系统)

- An operating system running in a virtual machine is called a “**guest operating system**” (客户操作系统).
- A **guest operating system** interacts with virtual hardware in the same way that it would interact with real hardware.
- The **virtual machine manager (VMM)** traps all interactions between the **guest operating system** and the real hardware.



Disadvantages of virtualization

Virtualization has **side effects** (副作用).

Performance may not be so good:

- All operations performed in a virtual machine are intercepted by the VMM before being executed.
- Because of this, a virtual machine may be slow.
- In general, running an application in a virtual machine is always slower.

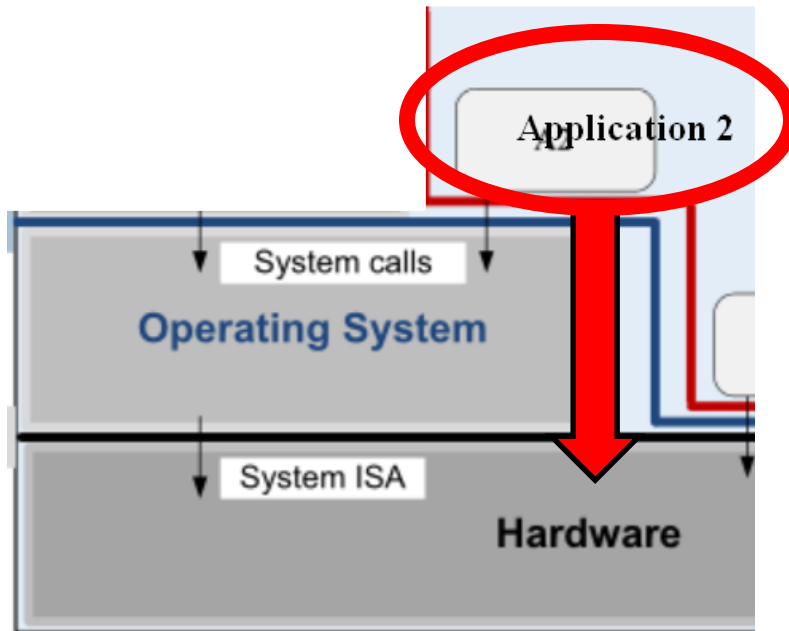
Disadvantages of virtualization

Higher hardware costs:

- More **expensive computers** may be bought for running virtual machines (VMs) compared to computers running a traditional operating system
- **Why?**
Virtualization needs **faster and/or multicore processors**, more memory, larger disks, and a faster network access.

5.2 Viewing computer systems as layers

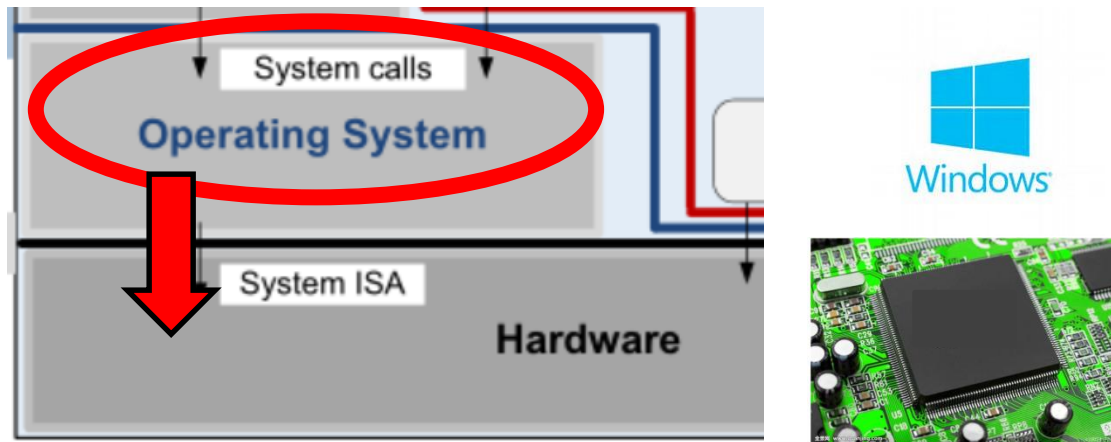
A computer system can be viewed as a set of layers:



An **application** typically interacts with the hardware (the processor, memory and storage) by first communicating with the operating system. Then the operating system tells the hardware what to do.

5.2 Viewing computer systems as layers

A computer system can be viewed as a set of layers:



Or an application can run directly on the hardware.
This is the case of the **operating system**.

Virtual machine (虚拟机)

Computer (电脑)

Operating System (操作系统) e.g. **Windows**



Virtual machine (虚拟机)

Computer (电脑)

Operating System (操作系统) e.g. **Windows**

Applications (应用)



Virtual machine (虚拟机)

Computer (电脑)

Operating System (操作系统) e.g. **Windows**

Applications (应用)



Virtual Machine Manager (虚拟机管理器)




Virtual machine (虚拟机)

Computer (电脑)

Operating System (操作系统) e.g. **Windows**

Applications (应用)



Virtual Machine Manager (虚拟机管理器)

Virtual Machines (虚拟机)



An alternative

Computer (电脑)

Virtual Machine Manager (虚拟机管理器)



An alternative

Computer (电脑)

Virtual Machine Manager (虚拟机管理器)

Virtual Machines (虚拟机)



5.3

Virtual machine manager (VMM) (虚拟机管理程序)

- It **manages how resources** of a computer **are shared** between one or more virtual machines.
- It **monitors performance** and takes actions when performance is not good.
 - **e.g. If not enough RAM memory is available, the content of a virtual machine may be saved to disk to free some memory.**



Virtual machine manager (VMM) (虚拟机管理程序)

- **It checks what the virtual machines or applications are doing to ensure safety.**


From the hardware point of view

- The **processor** of a computer (CPU) can perform a set of operations.
- A modern processor often offer **three modes**:
 1. **Kernel mode** (内核模式): any operations can be performed on the processor. (e.g. this mode is used by the operating system such as Windows and the VMM)



From the hardware point of view

- 2. User mode:** some operations considered dangerous (for security) are not allowed to be performed on the processor (e.g. this mode is used by applications to protect the operating system)
- 3. Virtual machine mode:** this mode restricts what a virtual machine can do on the processor (e.g. this mode is often used for virtual machines to protect other virtual machines and the operating system)

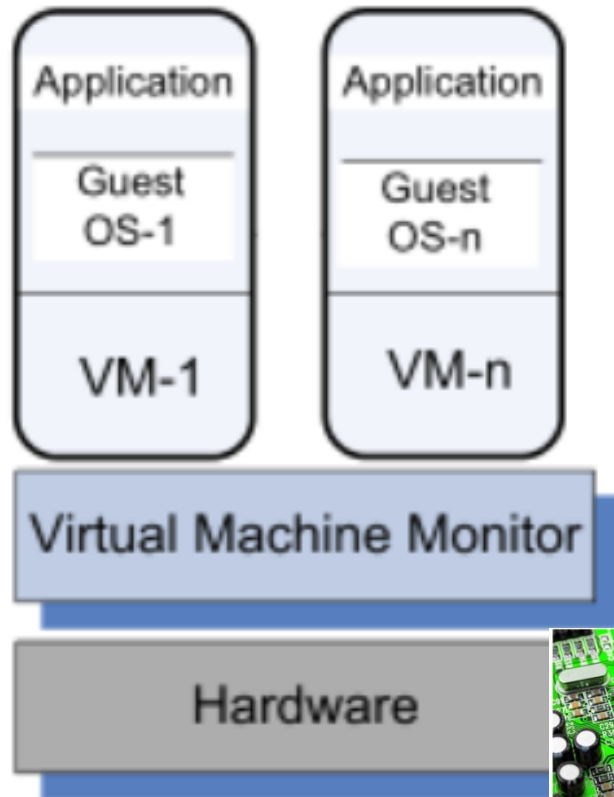


Two main types of virtual machines (VM)

- **System virtual machine:**
 - This type of virtual machine **can run an operating system** such as windows and **several applications.**
 - This is what we discussed mainly until now.
- **Application virtual machine:**
 - This type of virtual machine is designed to **run a single application.**
 - **e.g. Java** programs that are run directly in a virtual machine without using an operating system.

Traditional virtual machines

The VMM supports multiple virtual machines and runs directly on the hardware

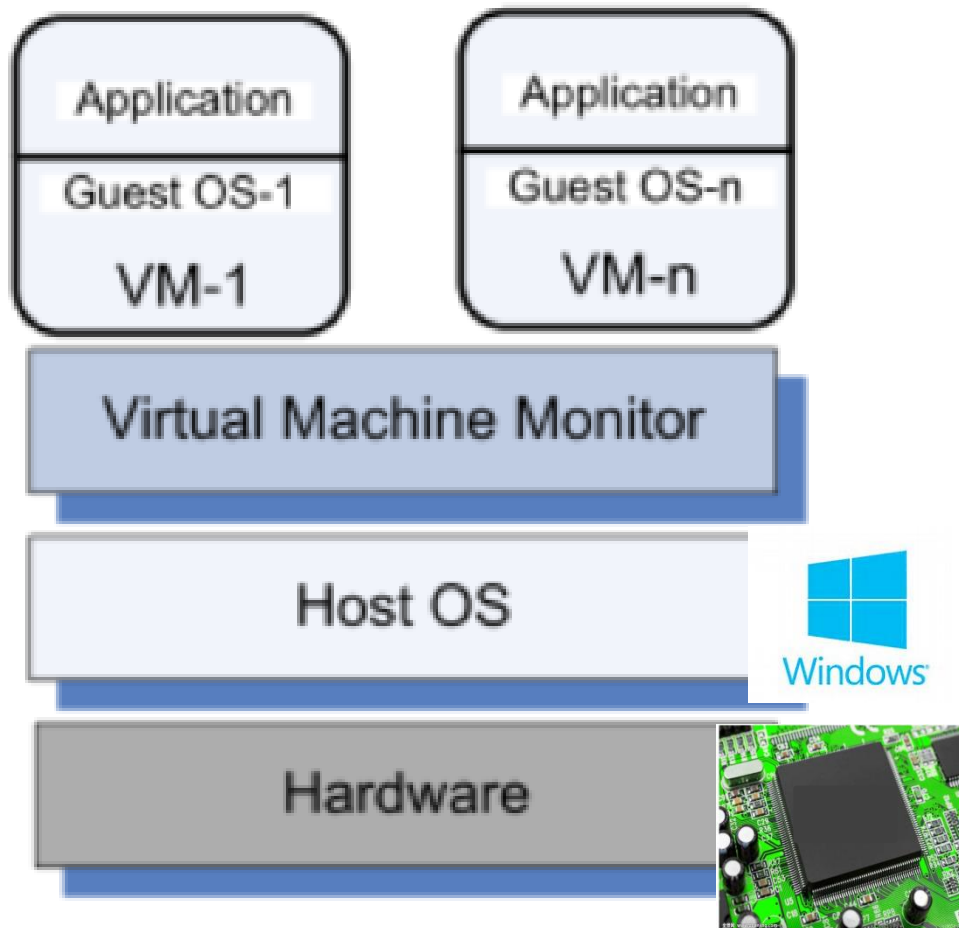


Advantage: performance

Example: [VMWare ESX](#)

Hosted virtual machines

The VMM runs under an operating system



Advantages:

The VM is easier to install
The VM can use some components of the operating system

Disadvantage:

Slower!


Example: [VirtualPC](#),
[VirtualBox...](#)

Performance penalty (性能代价)

- Using a virtual machine is slower than not using a virtual machine.
- How much slower **depends on the type of virtual machines.**
- If a virtual machine is running **directly on the hardware the difference is small.**
- **If** a virtual machine is running inside another operating system such as **Windows** the **difference will be greater.**

Performance penalty


- Some virtual machines will communicate directly with the processor to do calculations.
- This **can be done** when a virtual machine **is executed on the same** physical hardware that it simulates.
- This is **fast**.
- But they will still communicate with the VMM to use the network or read/write data (**for security**).



Performance penalty

- If a virtual machine simulates a kind of computer that is different from the hardware, then all operations performed by the VM must be translated by the VMM to be executed on the hardware.


This is **slower**.



5.5

Performance security isolation

- In theory, using virtual machines is very **safe**.
- **Even if an attacker gain access to a virtual machine, he cannot access other virtual machines** running on the same computer.
- This is true, **except if the VMM has some security problems**.



Performance security isolation

- Another reason why virtual machines are secure is that they are often restricted to only perform some operations.
- **e.g.** the Xen VMM has only 28 special operating system operations, while a standard Linux operating system may have more than 280.

Conclusion

- In this part, we have discussed **resource virtualization** for the cloud.

Chapter 5. D. C. Marinescu. Cloud Computing Theory and Practice, Morgan Kaufmann, 2013.