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

Introduction

Schedule	
Today	Cloud storage
Next week	Cloud security
	Final exam

Today,

- I will talk about how data is stored in the cloud.

8- CLOUD STORAGE (云存储)

It is not necessary to read the book.
But if you follow the book, we have skipped chapters 6 and 7.

pdf 241-271

Introduction

- A cloud has many computers.
- Many computers =
  - a lot of « **computing power** » (计**算能力**)
  - a lot of **storage space** (**存储空间**)

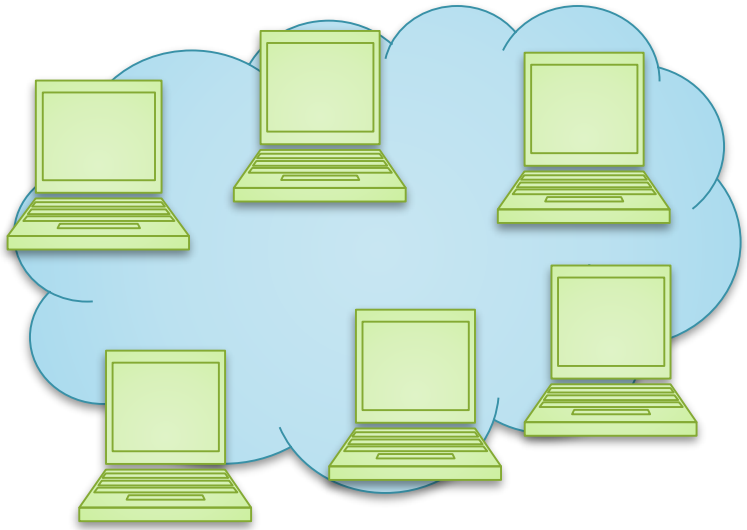

Introduction

- Users of the cloud can **store** and/or **modify data** in the cloud
  - Use the *Baidu Cloud* (百度云) to store files from your computer.
- Users of the cloud can **read data** (get data) from the cloud
  - Watching videos from *Youku* (优酷) or *TikTok* (抖音)

- For some clouds, data can be accessed using the Internet.

Introduction

- **Cloud data** can be accessed using **computers** or **mobile devices** (移动设备).


- **Mobile devices often have limited resources.**
- **Thus, mobiles devices** can use the **cloud** to store audio and video files.
  (**e.g.** *Xiaomi* cloud)
- The cloud is suitable for **video delivery**
  (**e.g.** *Youku, Toudou, TikTok*)

Another example - *AMAP* (高德地图)

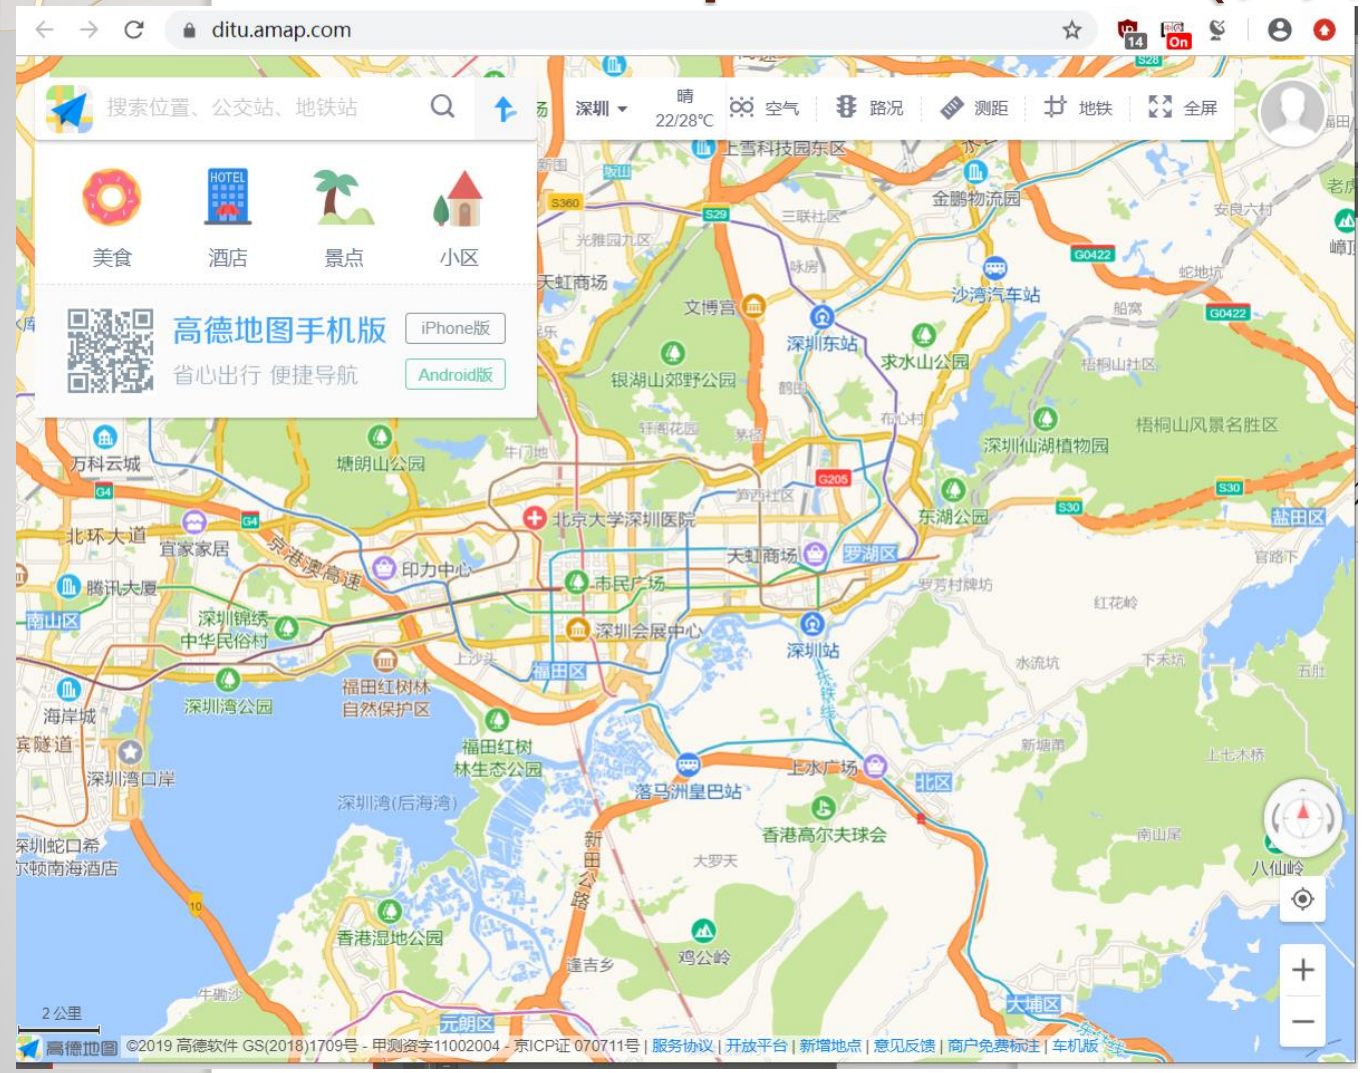

- A very big map of China is stored in the cloud.
- When using *AMAP* on your cellphone, *AMAP* will download **parts of this** map.

- To use *AMAP* offline (without internet), you can download the maps to your phone (**e.g.** map of Shenzhen).

Introduction

- **Cloud applications** may **constantly receive new data**.
- **e.g.** users are constantly uploading new videos on a video sharing website such as *Youku* or *Douyin*.
- **Cloud applications** often **collect data** about:
  - their **services (e.g.** performance**),**
  - the **users of these services**

    (e.g. how users utilize the cloud application?, how often? what they like / don't like?**)**.

Introduction

- The owner of a cloud application can **analyze the collected data** to obtain useful information. (**e.g.** to improve the application, for marketing, to launch a new product, pay-for-what-you-use, etc.)
- **Collecting data about users** is one of the reasons why some big companies offer **free** cloud services.
  "*If you are not paying for it, then you are the product*" → but not always true!
- **e.g.** some big companies may use data about users to show advertisements, or even **sell the data.**

元

The business of data…

- In the West, companies like *Google* and *Facebook* offer free services to collect data about users, and then use this data to sell targeted advertisements.
- Data may be sold to other companies, and even be resold.
- Some companies also get data from multiple sources and combine it to know more about people.
- **Data brokers**: companies or individuals that collect and then sell or resell data.

Illegal data brokers

- **Example from 2018**
  - The **Wuxi police** has arrested113 persons that were illegally collecting/selling information about people.
  - The criminal network was international. It involved people from China, Myanmar, etc.
  - **Data includes**: location, bank balances, credit information, property details, and familial information.
  - **Data was used for:** network fraud, taking out microloans, and violent debt collection.

The news →

https://tech.qq.com/a/20181011/003346.htm

热门精选

较真丨给小学生戴“脑电波走神监测环”，这事真有点不科学

华为Mate30系列5G版首销：4999元起 配麒麟990 5G芯片

昨夜，黄浦江的游船上来了一群世界上最聪明的人

广州10个月破获电信诈骗案4400宗 拦截止付逾5亿元

保险流量恶战：获客暴涨至上千元，重回线下开门店

监测头环被责令停用：系哈佛博士研发，公司刚成立就送到小学实验

中腾微网获千万级A轮融资，“资本”推动“科创”升级

揭秘个人信息交易黑市：内部分工明确 日交易额百万

互联网 人民日报 2018-10-11 08:31

[摘要]境外中间商是全国侵犯公民个人信息黑色产业的关键环节。目前，这批境外中间商已垄断国内绝大部分侵犯公民个人信息的地下产业链，甚至还有专门的价目表，日交易量达数十万条。

循线追踪、摧毁链条、打击源头，江苏无锡警方破获一起重大案件

除掉个人信息交易“地下黑市”(聚焦个人信息保护)

“代查各类信息，只有你想不到，没有我查不到……”今年2月，江苏无锡警方网安部门在网络巡查中发现，某网络软件聊天平台上，有人打出了这样一则广告。

夸张的是，这条广告竟然“所言不虚”，发布广告者只是犯罪团伙中的冰山一角。

Types of data about persons (1)

- **Identifying data**
  - Name, address, e-mail, phone number
- **Sensitive identifying data**
  - ID card number, birth date, passport number
- **Public data**
  - Mariage record, divorce record,…
- **Demographic data**
  - Age, Race, Religion…
- **Vehicle data**
  - Vehicle identification, insurance data

https://onlinemba.unc.edu/blog/data-brokers-infographic/ (2015)

Types of data about persons (2)

- **General interest data**
  - Music, Movies, Sports, …
- **Technology data**
  - Mobile phone usage, location, social network activity
- **Financial data**
  - Salary, money in the bank, etc.
- **Travel data**
  - Train and airplane tickets purchases…
  - Frequent flyer number

https://onlinemba.unc.edu/blog/data-brokers-infographic/ (2015)

Types of data about persons (3)

- **Health data**
  - Purchases of medicines,
  - Results of medical tests, etc.

- …

https://onlinemba.unc.edu/blog/data-brokers-infographic/ (2015)

Where data is obtained from?

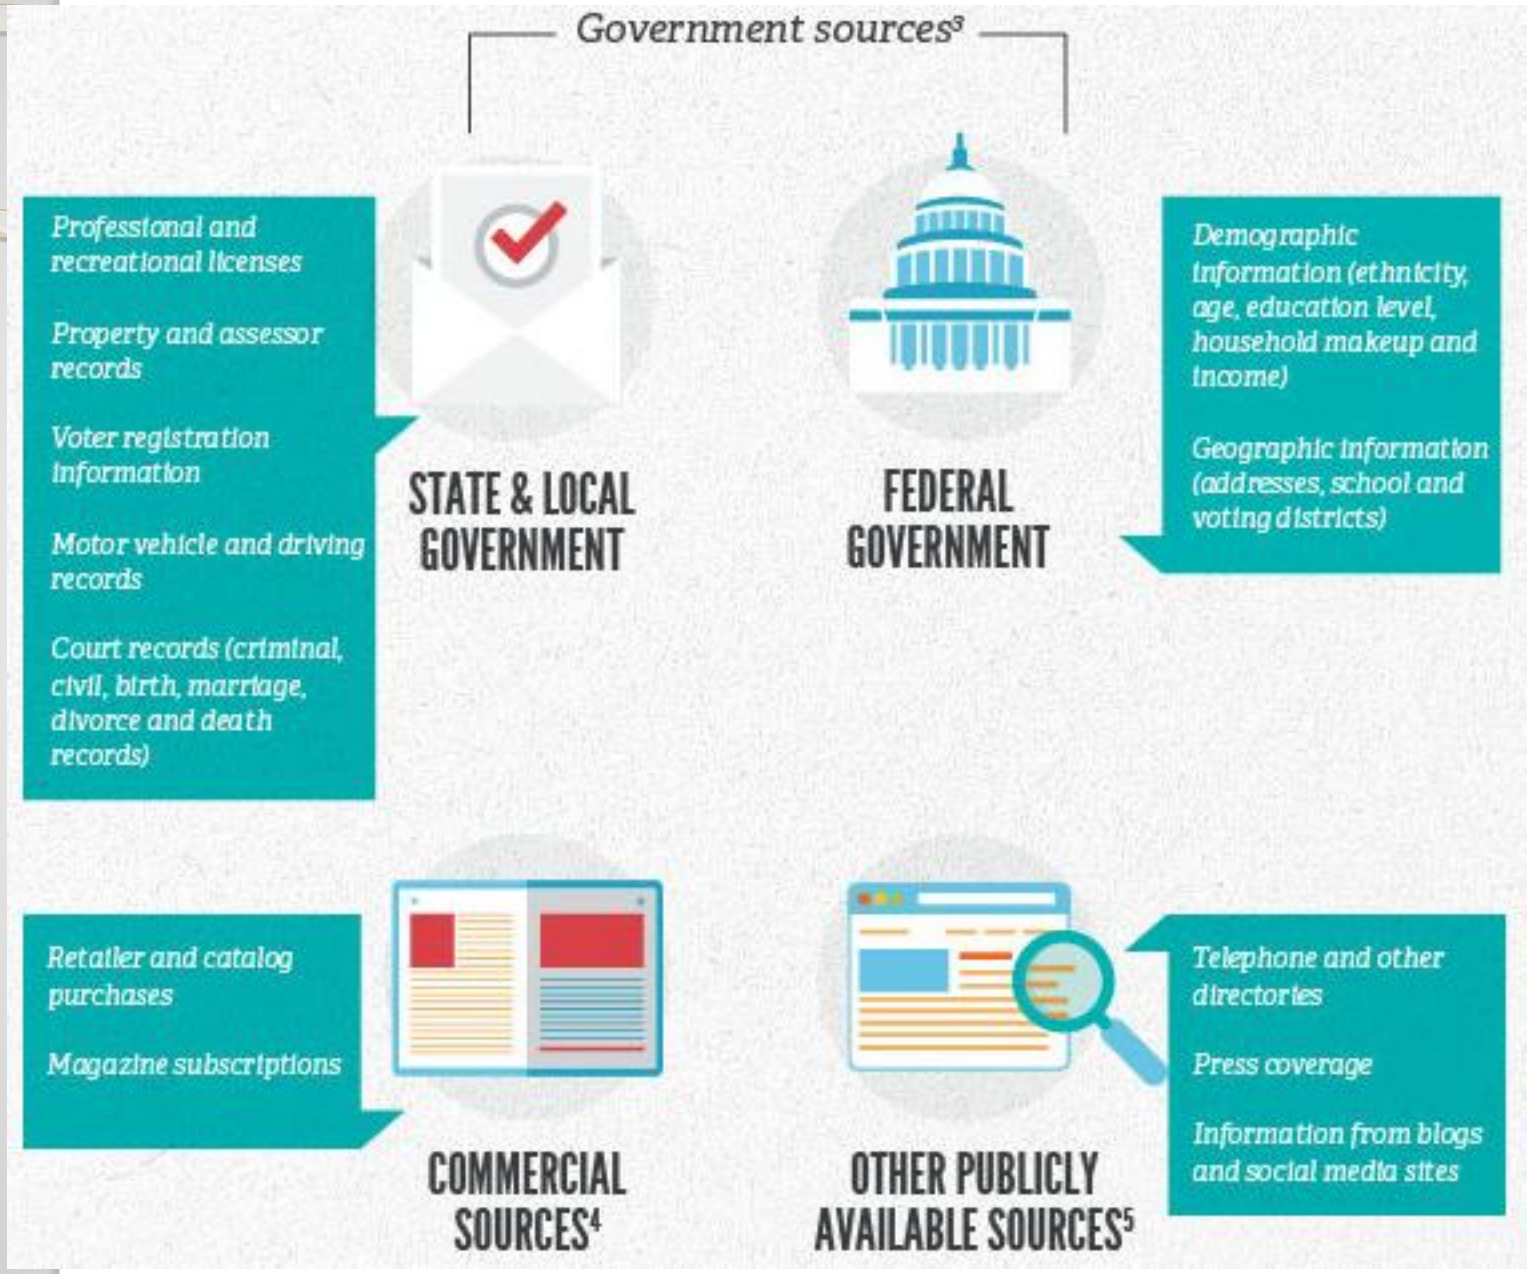

https://onlinemba.unc.edu/blog/data-brokers-infographic/ (2015)

What is data worth to you?

https://onlinemba.unc.edu/blog/data-brokers-infographic/ (2015)

Keeping data private?

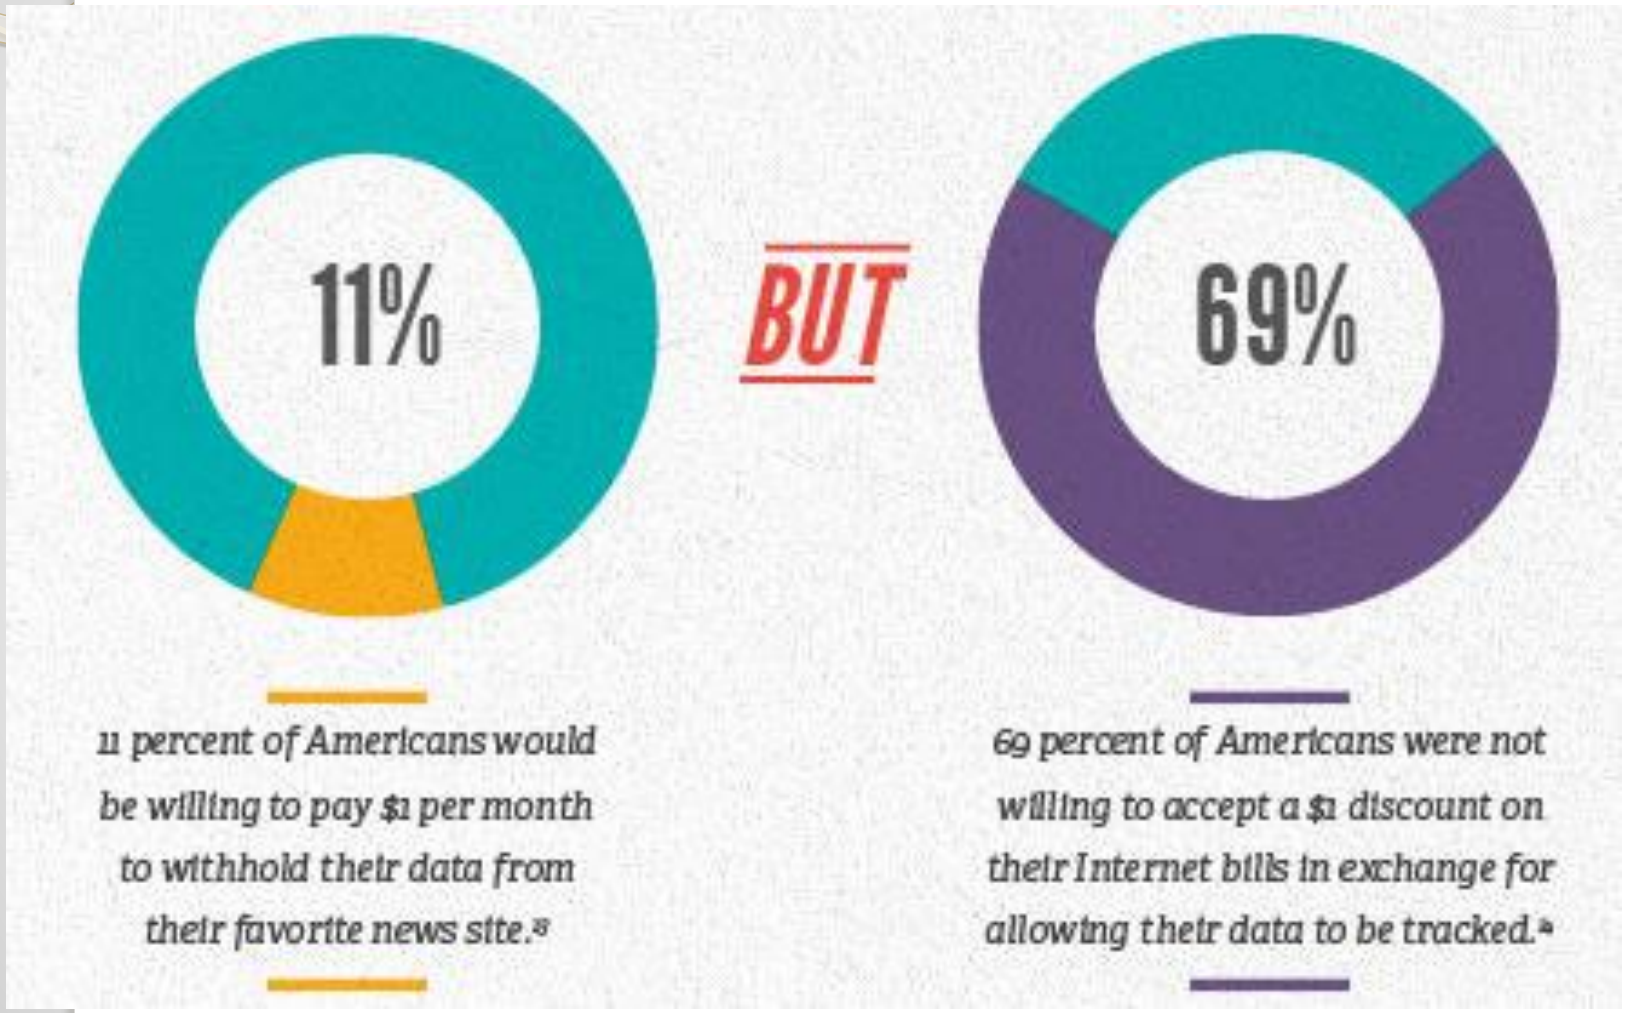

https://onlinemba.unc.edu/blog/data-brokers-infographic/ (2015)

Or having personnalized services?

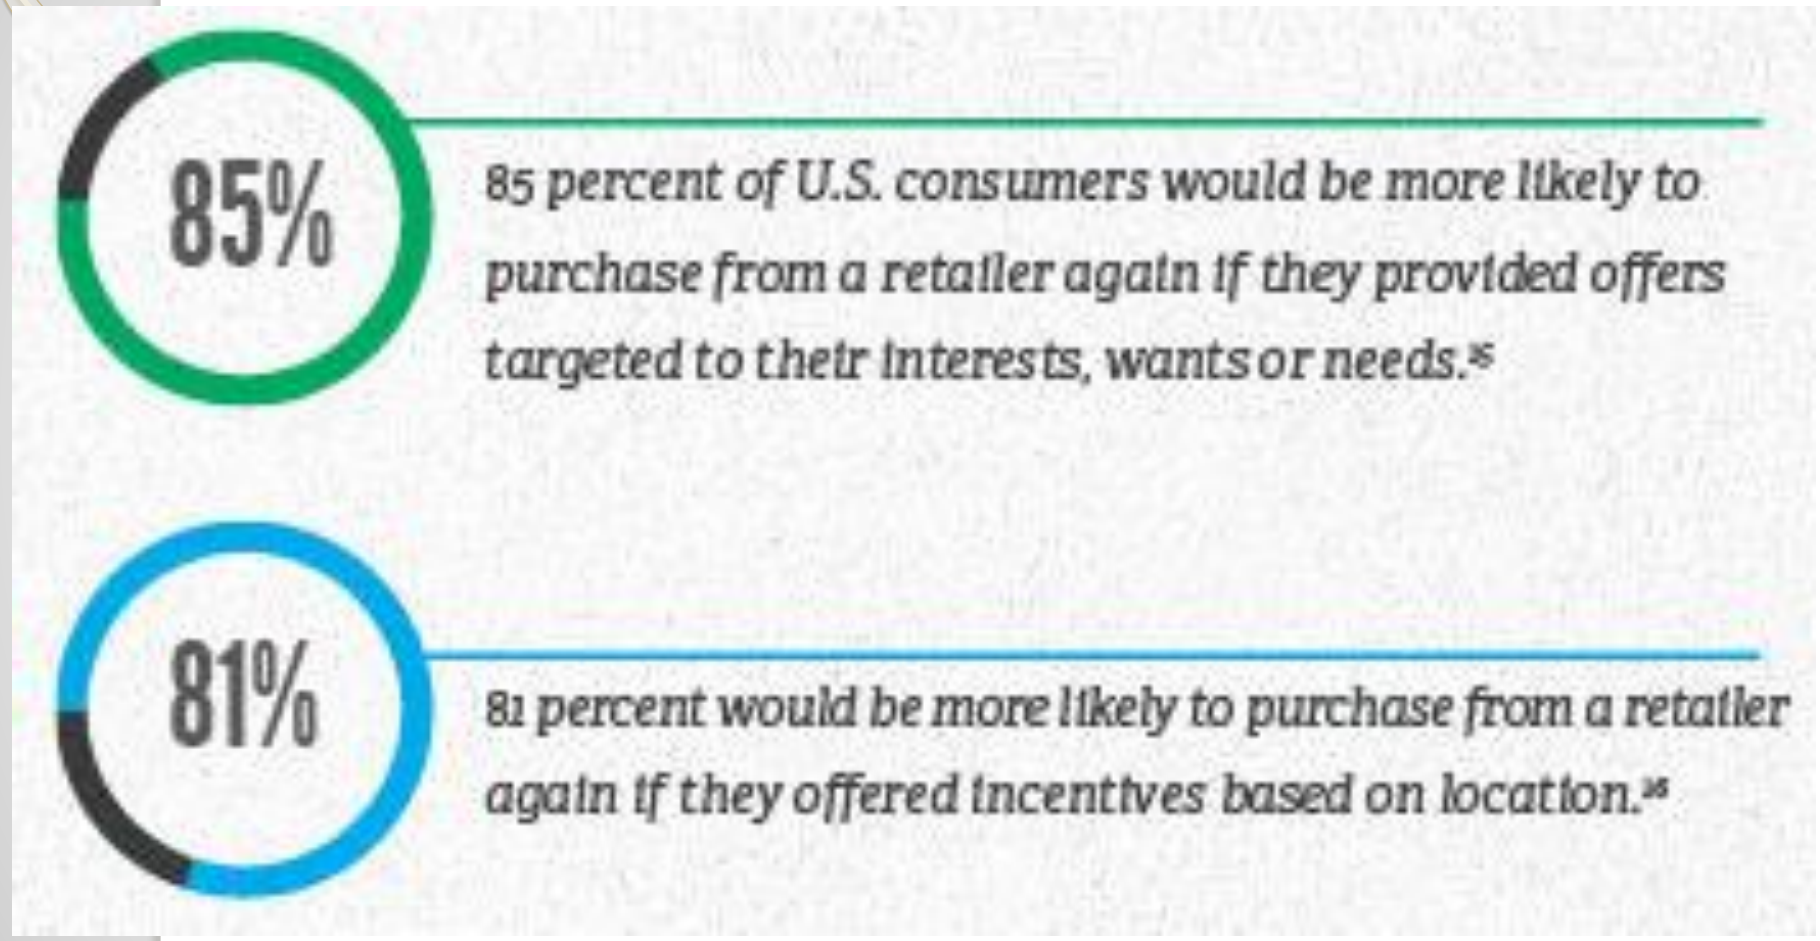

https://onlinemba.unc.edu/blog/data-brokers-infographic/ (2015)

Laws and regulations for data

- Different countries have different laws and regulations about how data must be stored, shared, can be used, etc.
- In **China**, we have the **Cybersecurity and Data Protection Law (2017)**
  - Network operators must tell the **purpose**, **method**, and **scope of data** collection and use, and obtain consent from users.
  - Network operators shall **not disclose, alter, or destroy personal data** without the consent of users

… continue on next slide →

Laws and regulations for data

- In the event of a **data breach,** network operators must take **actions to fix the problem, inform users**, and report to the appropriate government agencies

(this is just a rough unofficial translation of some parts of the law – for more details, see the original law in Chinese language)

Data storage and processing

- **Data storage** and **data processing** in the cloud are closely related.

- If it takes more time to **access the data**, it will take more time to **process the data** (处理数据).
- **Access time**: the time needed to obtain some amount of data.
- **If we have more data**, then **we need more computing power** to be able to process the data at the same speed.

Data storage and processing

- **Cloud applications** often **process large amount of data**.
- **How to process large amount of data quickly?**
  - use **data replication** (数据复制)
    - stores multiple copies of data,
    - store data in multiple locations.
  - use appropriate **storage management strategies**
  - **(存储管理策略)**.
  - write cloud applications that use the data in an **efficient** (有效率的) **way**.

Using the data in an efficient way?

Example

- I have data about 1 million employees.

- I want to calculate the average salary.

- **How?**

Using the data in an efficient way?

Example

- I have data about 1 million employees.

- I want to calculate the average salary.
- **A <u>simple</u> approach:**
  1. Read the data to calculate the sum of salaries.
  2. Read the data *again* to count the number of employees
  3. Divide the sum by the number of employees to obtain the average salary.

This can be very slow because we need to read the data twice (two times)!

Using the data in an efficient way?

Example

- I have data about 1 million employees.

- I want to calculate their average salary.
- **A better approach:**
  1. Read the data *only one time* to calculate the sum of salaries **AND** the number of employees.
  2. Divide the sum by the number of employees to obtain the average

  This is faster because we read the data only once (one time)!

Using the data in an efficient way?

Example

- I have data about 1 million employees.

- I want to calculate their average salary.

Data location is also important!

(e.g. all the data is in Shenzhen, or some data is in Beijing and some in Shenzhen**).**

The two previous approaches are what we call ***algorithms*** (算法), that is a *way* of calculating something.

A good programmer (程序员) will design *good algorithms*.

- Give the expected result,
- Has good performance (speed, memory…)

Big Data (大数据)

- The term "**Big Data**" (大数据) has become very popular.
- **What is *big data*?**
- **A simple definition:** data that is so **large** that it cannot be stored and processed on a single computer.
- Nowadays, the term **Big Data** is used with many different meanings, that are often far from the original meaning of Big Data.
- One of the most **popular definition of Big Data** →

Big Data

Data that has three main characteristics (the three "Vs") :

- **Volume (数据的数量)**: a large amount of data
- **Velocity (速度)** : the data is arriving at a very high speed (**e.g.** new messages published on Weibo every second)
- **Variety (各种各样)**: the data is of different types such as text, images, audio, video, graphs,….

Big Data

Some persons have added two more Vs:

- **Veracity** (真实性): the data may be of poor quality (inaccurate) or not trustworthy (不值得信赖).
- **Value** (价值): it is important to try to use big data to obtain some "business value", that is to extract useful knowledge from data.

Thus, Big Data = 5 Vs

Some examples

- **A laptop computer (笔记本电脑) may store 80 GB to 4000 GB**
- **In 2010,** the four main detectors at the **Large Hadron Collider (LHC),** a particle accelerator, produced **13 petabytes (PB)** of data (**13,000, 000 GB**).
- The **Sloan Digital Sky Survey (SDSS)** is a telescope that collects about **200 GB** of data per night.

File system (文件系统) of a computer

- A **file system** (文件系统) is **a way** of **organizing** and **storing data** in a computer.
- **In a desktop computer** (台式电脑) a file system is used such as **NTFS, FAT, HFS, ext3**…

Hard drive
硬盘驱动器

Table of content (**内容目录**)
File « Movie.avi » is located at …
File « file.doc » is located at …

File system (文件系统) for big data

How about big data?

- **For the cloud**, the data is stored on multiple computers.
- For this reason, we need **file systems** (文件系统) that are designed for organizing data on multiple computers.
- **Some file systems used for the cloud are: GFS** and **HDFS** …

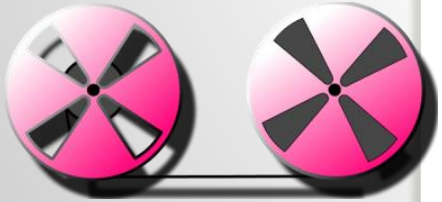

THE EVOLUTION OF STORAGE TECHNOLOGY

Storage capacity (存储容量)

- The **storage capacity** of a storage device is **measured** in **number of bytes**.
  1 bit = 0 or 1
- 1 byte = 8 bits = approximately 1 letter.
- Storage needs vary.

Storage Term	Approximate Number of Bytes	Exact Number of Bytes
Kilobyte (KB)	1 thousand	2^{10} or 1,024
Megabyte (MB)	1 million	2^{20} or 1,048,576
Gigabyte (GB)	1 billion	2^{30} or 1,073,741,824
Terabyte (TB)	1 trillion	2^{40} or 1,099,511,627,776
Petabyte (PB)	1 quadrillion	2^{50} or 1,125,899,906,842,624
Exabyte (EB)	1 quintillion	2^{60} or 1,152,921,504,606,846,976
Zettabyte (ZB)	1 sextillion	2^{70} or 1,180,591,620,717,411,303,424
Yottabyte (YB)	1 septillion	2^{80} or 1,208,925,819,614,629,174,706,176

Cost of storing data

Every year, it becomes **cheaper** to **store data**, thanks to technological advances.

Storage cost US $ / gigabyte (GB) 1980-2009

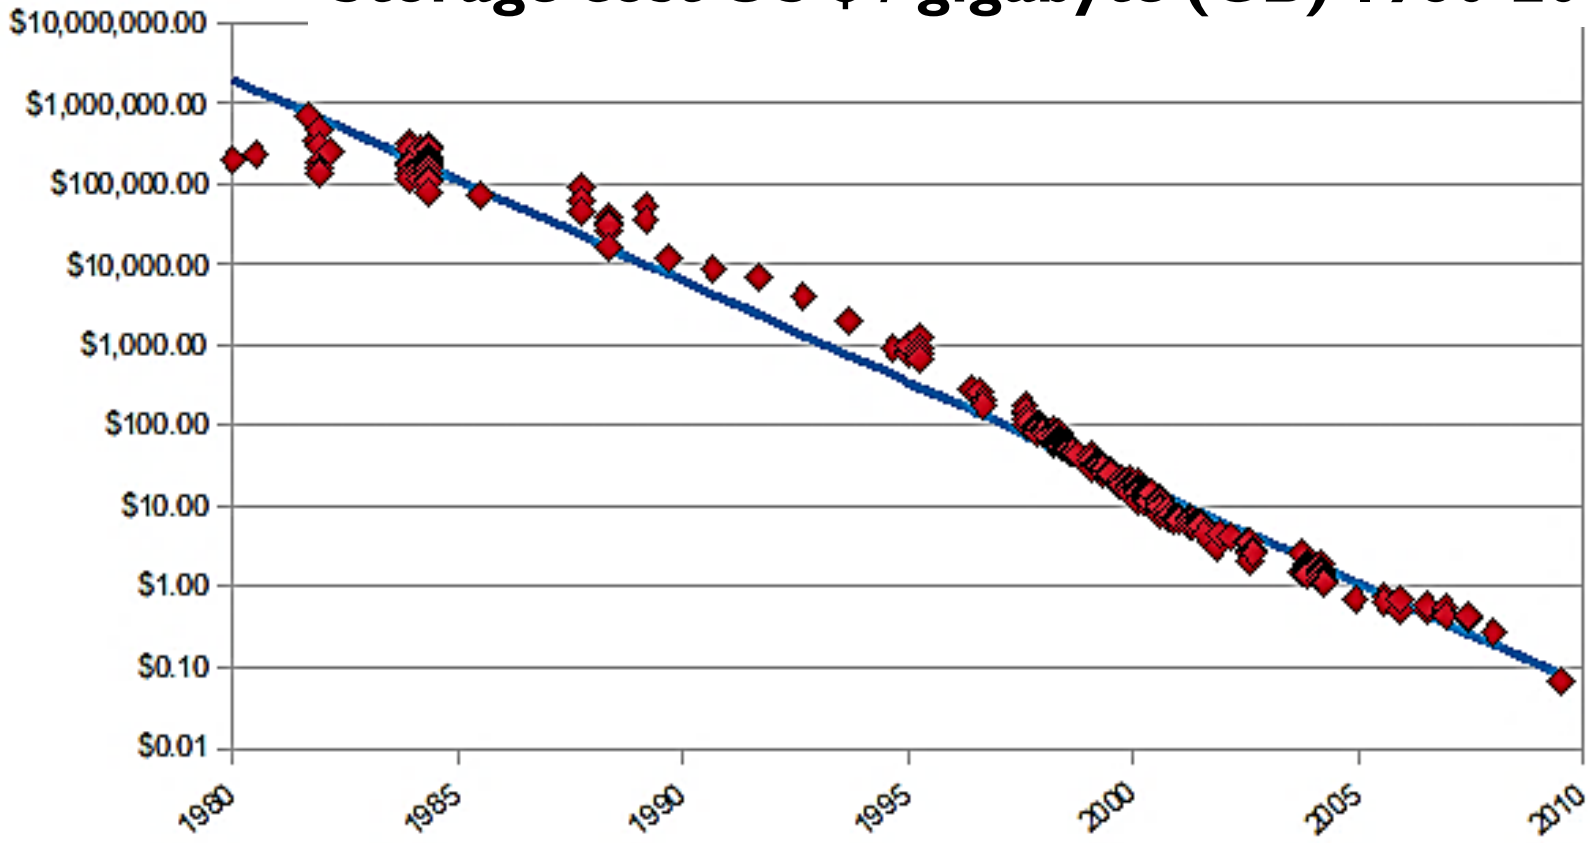

http://www.mkomo.com/cost-per-gigabyte

Storage density (存储密度)

- Over the years, **storage density (存储密度)** has also increased.
- **Density** is the **amount of data** that can be stored on some **surface**.
- In 1980, the density was approximately 0.01 GB / inches^2 (平方英寸)
- In 2011, the density reached approximately 744 GB / inches^2 (平方英寸)

About 10,000 times more dense!

Access time

- While the density of storage devices has increased and the cost has decreased dramatically, the **access time has improved only slightly**.

- Some recent advances:
  SSD drives…
  (一旦固态硬盘)

How to store big data?

- Often:
  cost of managing data > cost of storing data.
- Storing data in the **cloud** can help to **reduce management costs**.
- How data is stored is important.
- Applications should also be designed in a way that tries to reduce the number of accesses to the data.
- This helps processing the data faster.

How data is stored in the cloud?

- **In the cloud**, normal **computer components** are used to store data such as hard drives (硬盘驱动器).
- In the past, for a **desktop computer**, the design of storage devices was « **performance** at any cost ».
- For the **cloud**, it is « **reliability** at the lowest cost ».
- **Reliability** (可靠性) means the ability of a device to work without failure for a long period of time.
- In the cloud, we wish that the cloud continue to work normally **even if some computer(s) fail**.

Cloud storage: how to choose?

- Cost
- Compatibility with your devices
- Free storage? How much for more?
- Maximum file size? Types of files?
- Reputable company?
- Transfer speed (bandwith)?
- Read the privacy policy / license agreement.

STORAGE MODELS (存储模型), FILE SYSTEMS (文件系统) AND DATABASES (数据库)

8.2

How data is stored in a computer?

A computer stores data on some storage device(s).

We can distinguish between:

- how the data is stored **physically** (物理存储)
- how the data is stored **logically** **(逻辑存储)**, that is how the data is organized in terms of files, folders, databases, etc.

Most storage devices provides two operations: **reading data** and **writing data**.

Many different kind of storage devices

Characteristics of storage devices

- **access time:** time to locate a unit of data on a storage device
  (**e.g.** 10 ms to read 1 byte of data from a hard drive)
- **transfer rate:** the speed for transferring data from and to a device.
  (**e.g.**10 GBs = 10 Gigabytes / second)

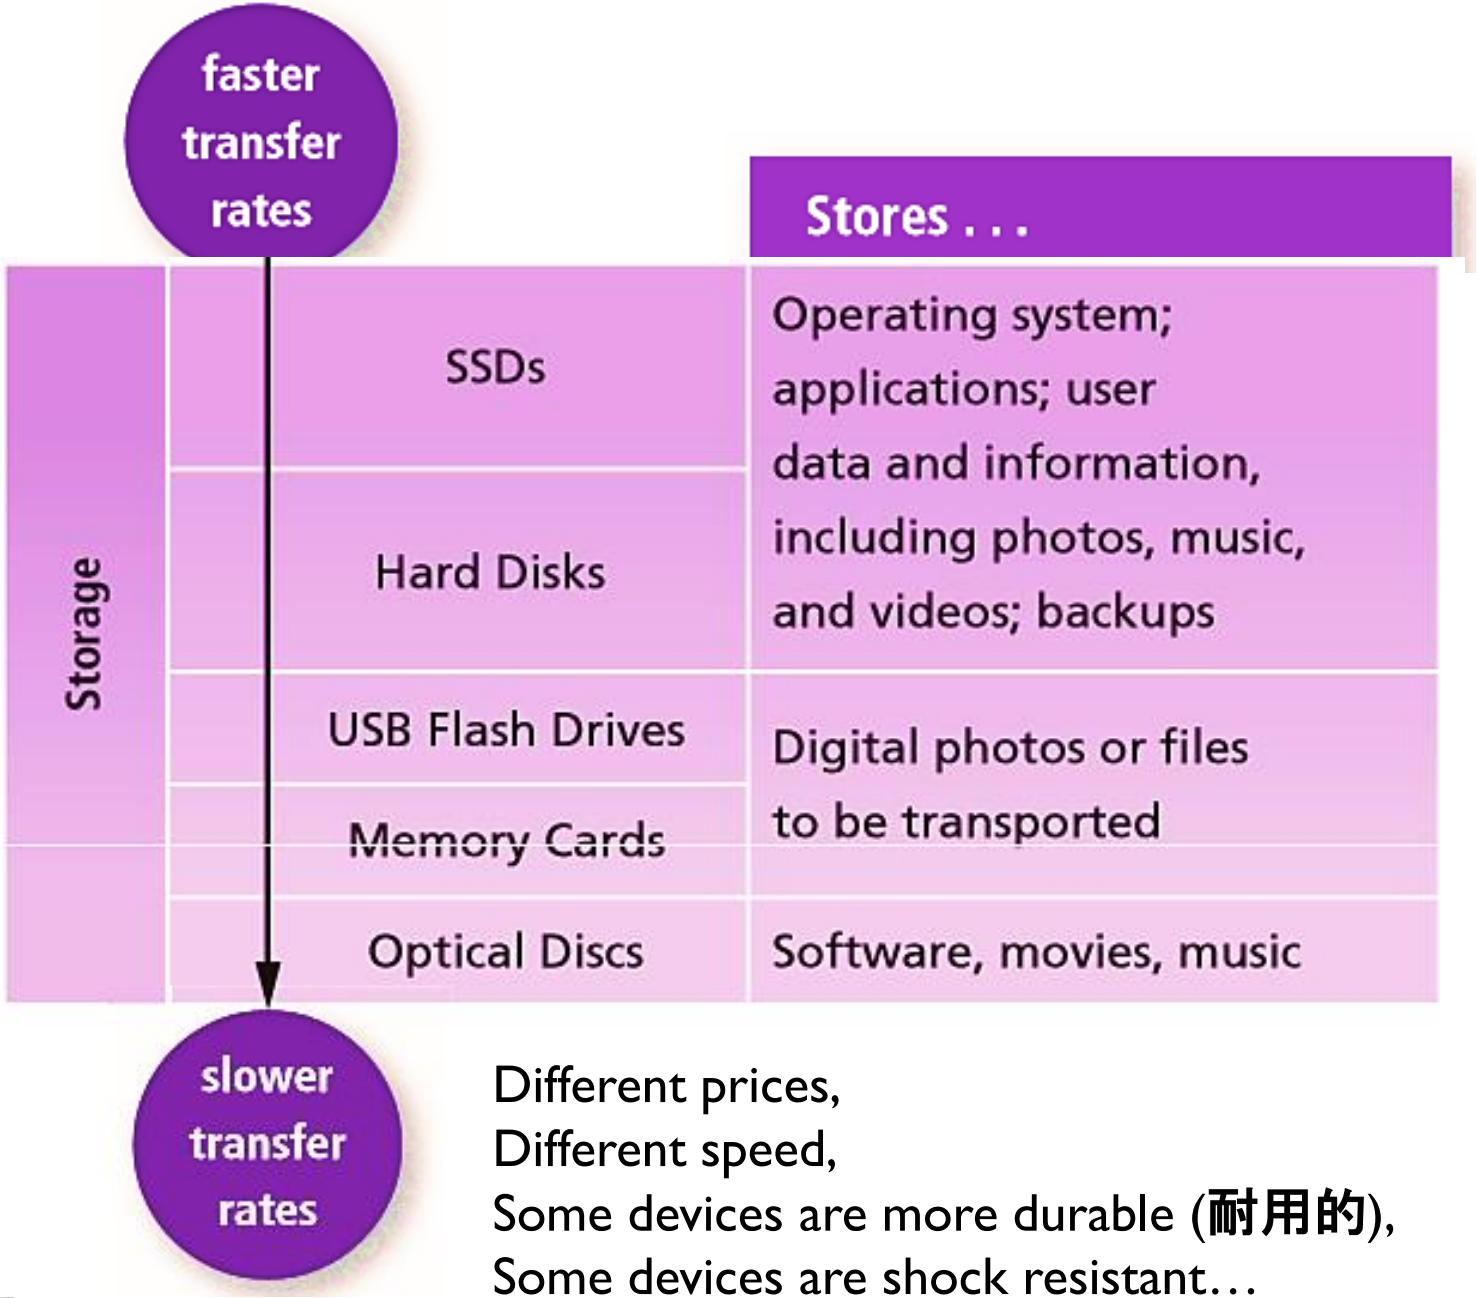

Different prices,
Different speed,
Some devices are more durable (**耐用的**),
Some devices are shock resistant…

How data is stored? – Cell storage

- **Physically**, most storage devices **stores data in memory cells (存贮单元)**.
- A **memory ell** is a memory location where data can be stored.
- All cells typically have the **same size**.
- A cell has a name (called its **address (内存地址)**):
  e.g. cell 1, cell 2, cell 3…
- An application can read or write data in a cell.

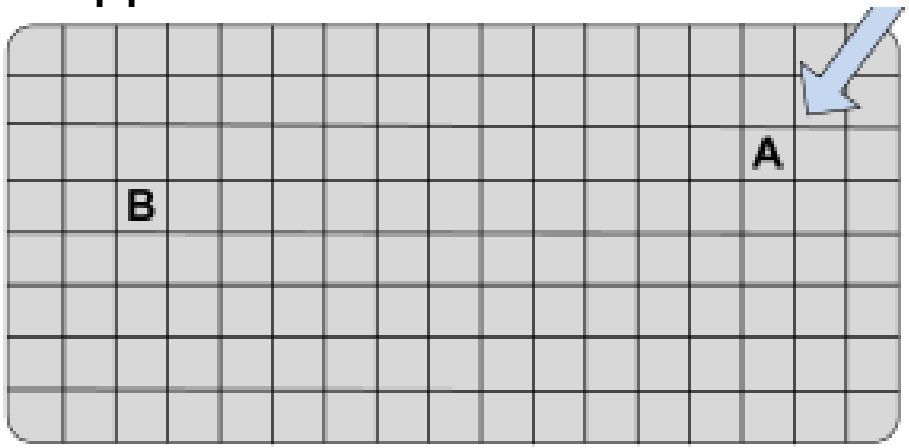

Cell storage

Journal (日志) storage

- **An improvement of cell storage**, designed to better cope with failures and the loss of data, is **journal storage**.
- The idea is simple.
- A **journal (log - 日志)** is used to record all the changes that are made to the **memory cells**.
- If some **error happens** and some data is not saved, it is possible to read the **journal** to restore the data.

Example →

Example

The computer wants to write A into a cell. Before changing the value in the cell, the change is recorded in the journal.

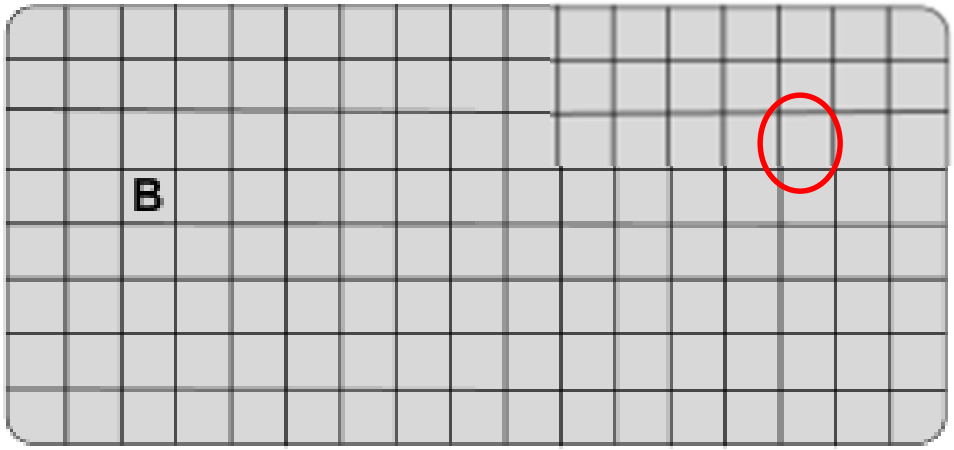

Cell storage

Example

The computer wants to write A into a cell. Before changing the value in the cell, the change is recorded in the journal.

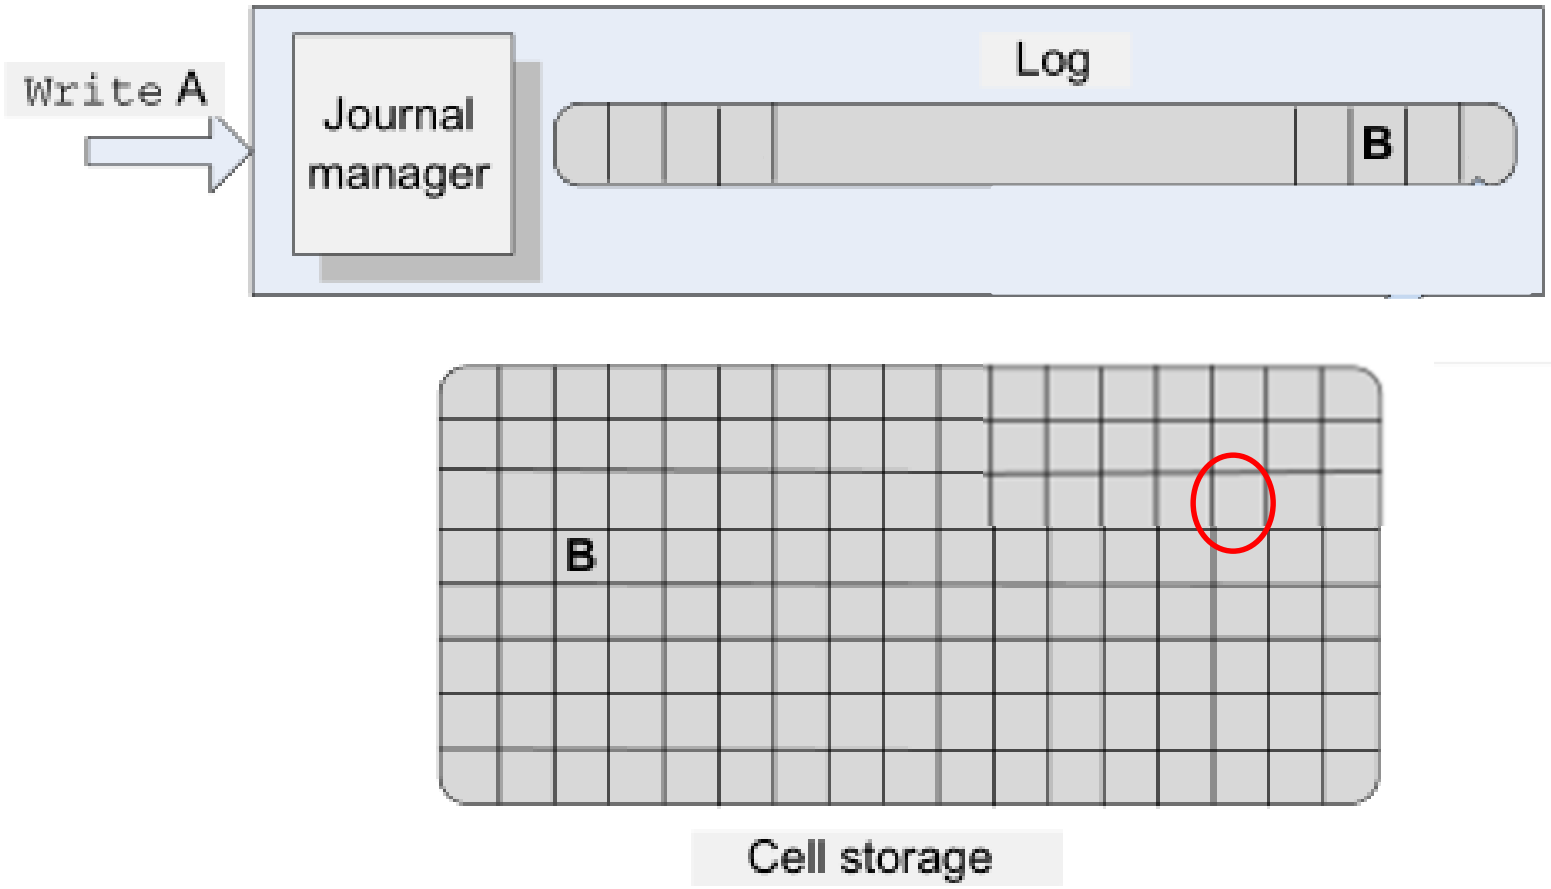

Example

The computer wants to write A into a cell. Before changing the value in the cell, the change is recorded in the journal.

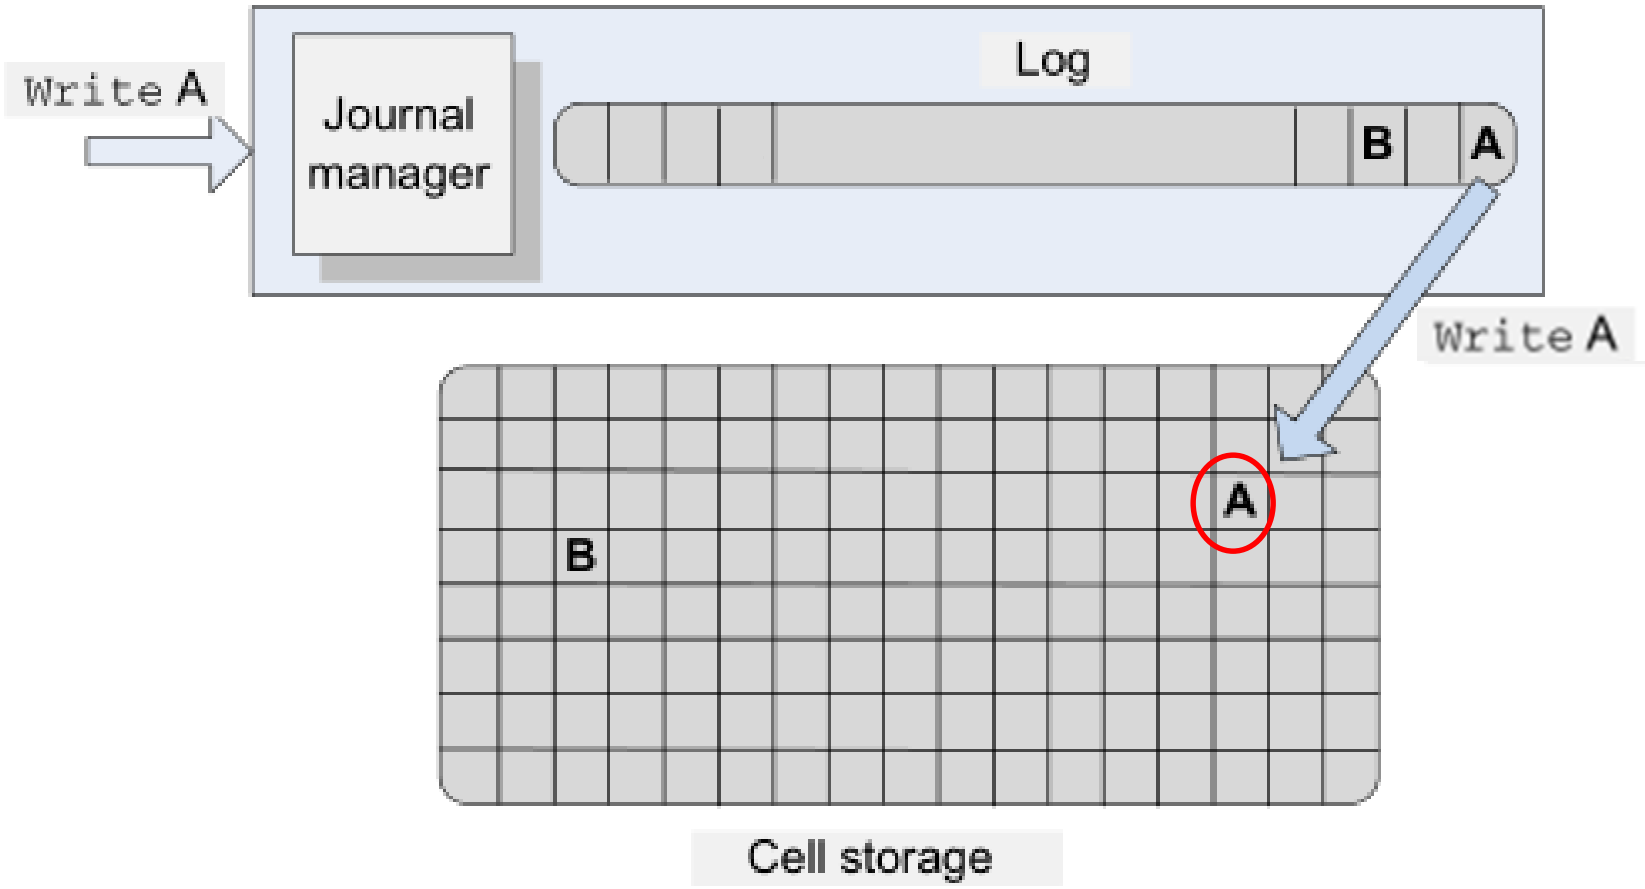

Journal (日志) storage

- **Journal storage** is useful for applications such as transferring money between banks.

- To ensure that no money is lost if a transaction fails, the transaction is first recorded in the journal before the data is modified.
- As a result, if the transaction fails, it is possible to undo the changes that have been made.

File system (文件系统)

It is a collection of directories (目录 - folder), which can contains files (文件).

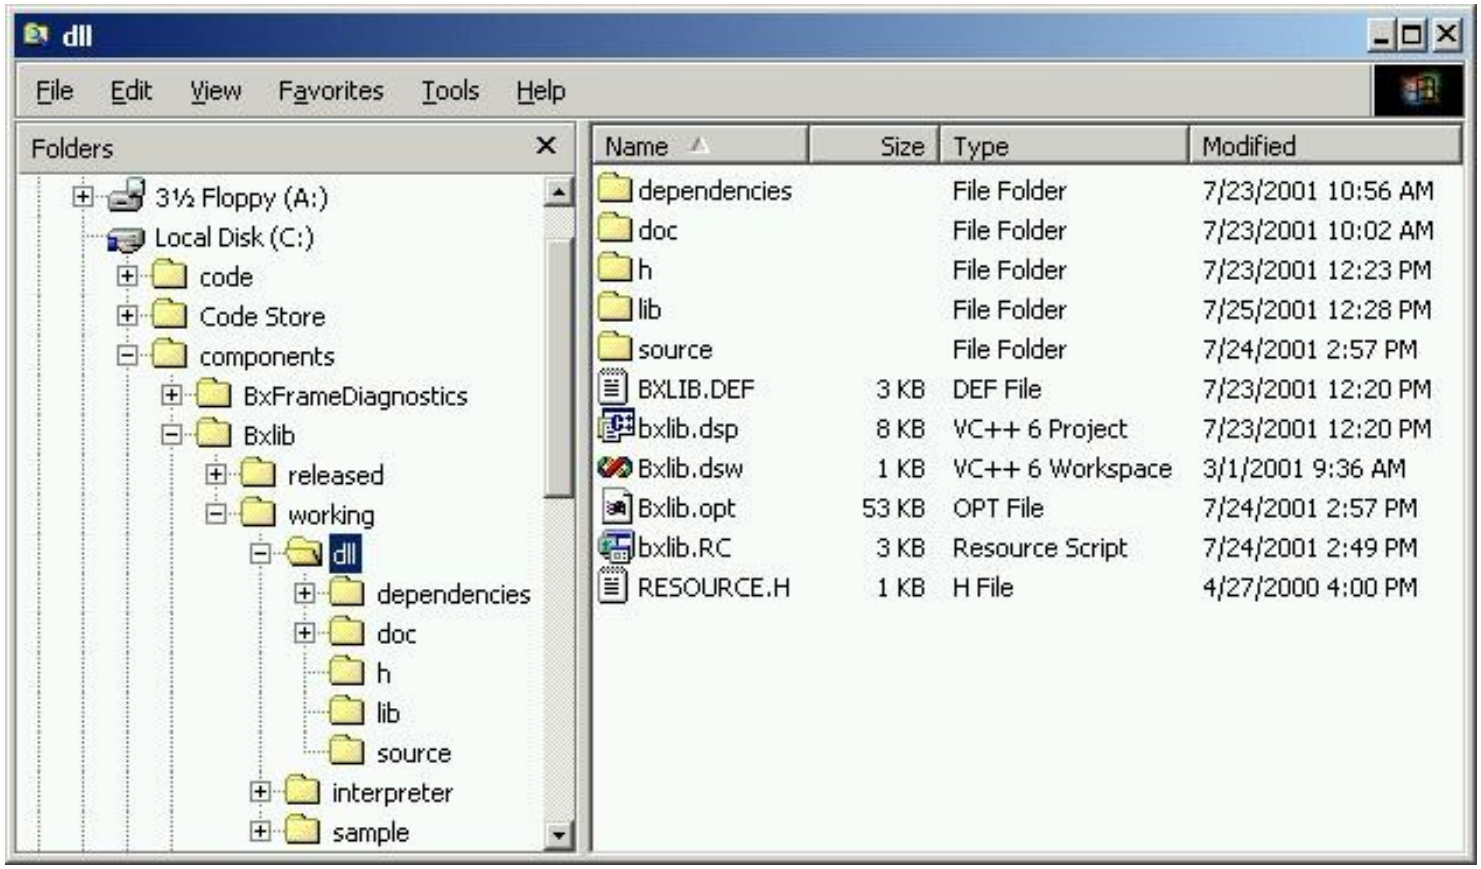

File system (文件系统)

- **A file system** for multiple computers is called a **distributed file system** (分布式文件系统).
- They are:
  - scalable (可伸缩的),
  - capable of distributing files between many computers,
  - reliable, etc.
- In the cloud, computers used to store data are often not the computers that perform the calculations. **Why?**

Database (数据库)

- Some cloud applications do not directly use files but instead use databases.
- **Database systems (数据库系统)** provides many **useful features** such as:
  - **access control** (访问控制): controlling who can access the data,
  - **concurrency control** (并发控制): controlling if several computers can read/write the same data at the same time.
  - **recovery after failure** (数据恢复): providing features to recover data when some failure occurs.
  - **query languages** (查询语言): some languages can be used such as SQL to search for information in a database.

Database (数据库)

Many databases use the « **relational model** » (关系模型), where data is stored into **tables**.

School Table

ID	Name
S001	University of Technology
S002	University of Applied Science

Student Table

School ID	ID	Name	DOB
S001	UT-1000	Tommy	05/06/1995
S001	UT-1000	Better	16/04/1995
S002	UAS-1000	Linda	02/09/1995
S002	UAS-1000	Jonathan	22/06/1995

e.g. Microsoft Access

This model is popular but it does not work well for storing and processing large amount of data in the cloud. Thus, other solutions have been designed for the cloud.

Requirements for storing data in the cloud

- **Data replication** must be performed.
- It requires mechanisms to ensure **fault-tolerance** (容错) (that the system works correctly even when failures occur).
- All copies of the data must. **consistent** with one another (彼此一致)
- Accessing the data must be **fast**.
- Other requirements may be **specific** to some particular applications

THE GFS FILE SYSTEM (GFS文件系统)

8.5

Introduction

- It was developed in the late 1990s.
- **Goal**: store data from Web search engines.
  - Store data on thousands of computers using cheap storage devices.
  - **Reliable (可靠的)**. It tolerates hardware failures, software errors, applications errors and human errors.
  - Store **large files** (**gigabytes** or even **terabytes**)
- There is an open-source (开放源码) version of GFS called **CloudStore**

How data is stored?

- Each **file** is divided into one or more **chunks** (parts of a file)

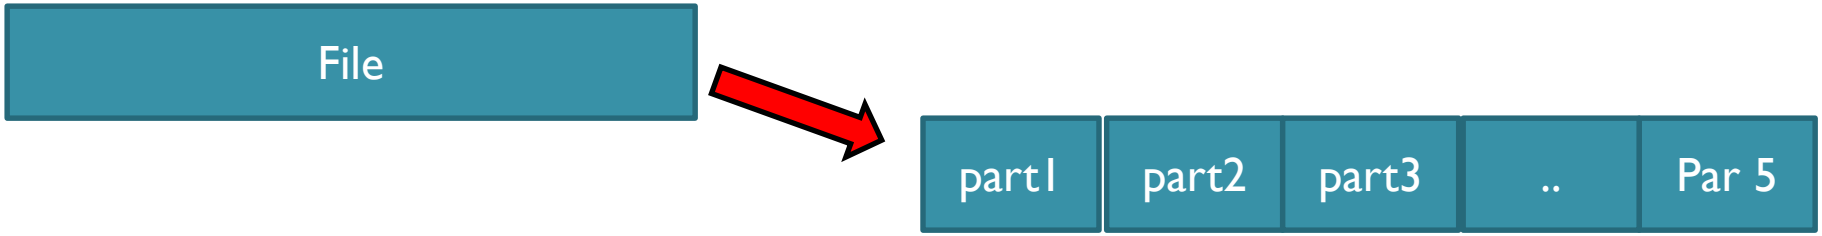

How data is stored?

- Each **file** is divided into one or more **chunks** (parts of a file)

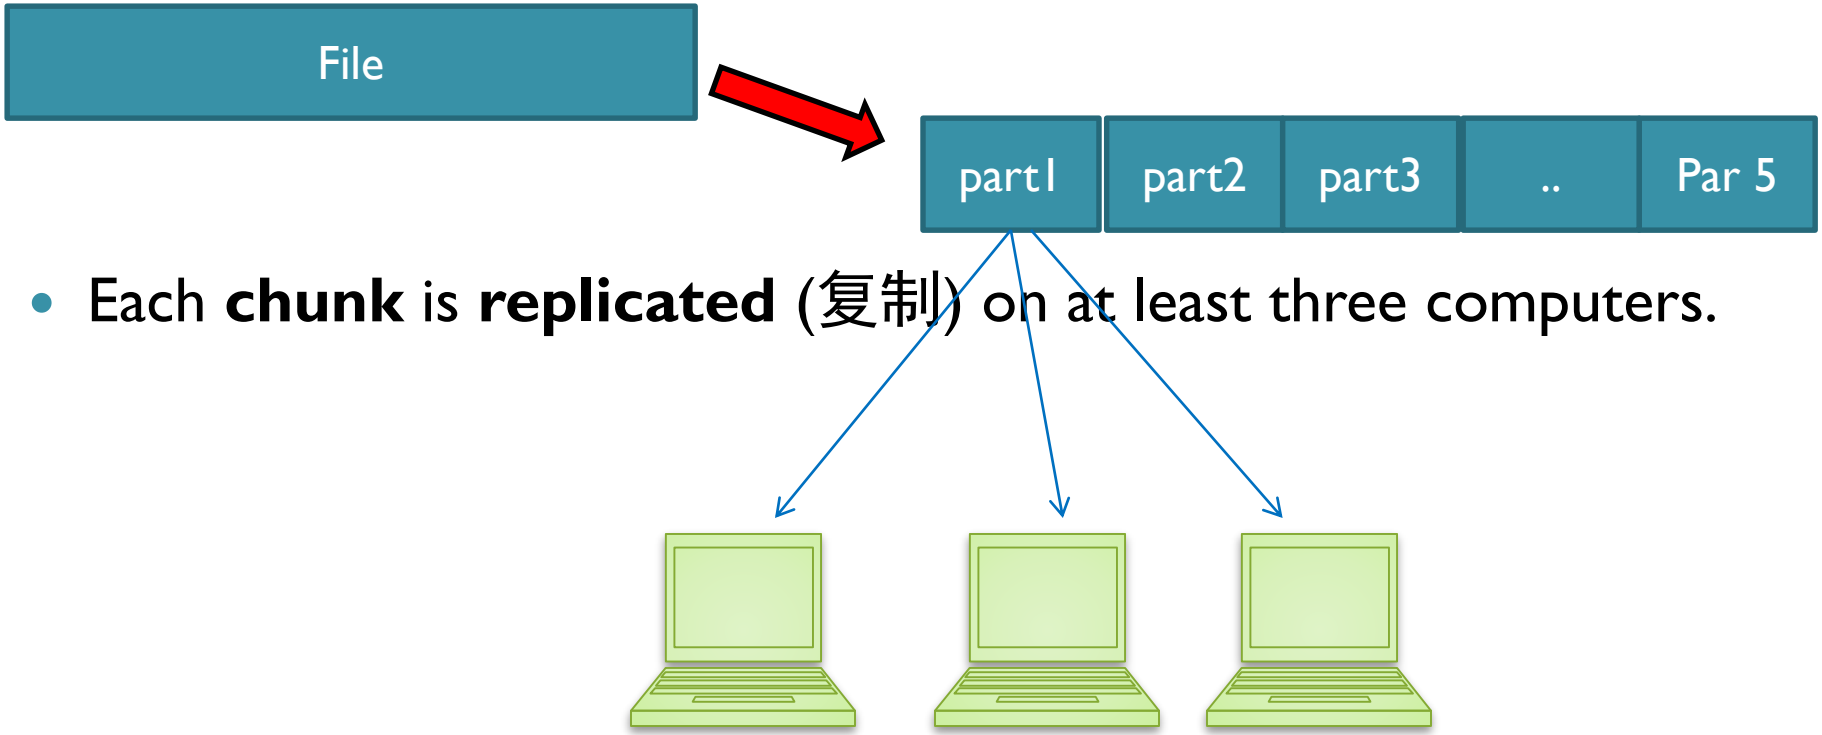

- Each **chunk** is **replicated** (复制) on at least three computers.

The user choose how many copies should be done for each chunk.

Two types of computers

- **The "master" computer**:
  - which supervises all the other computers used for storing data,
  - it **knows where the data** (chunks) **is stored**,
  - **it monitors the state of all computers** used for storing data.

- **The chunk servers**:
  - they store the data (chunks)

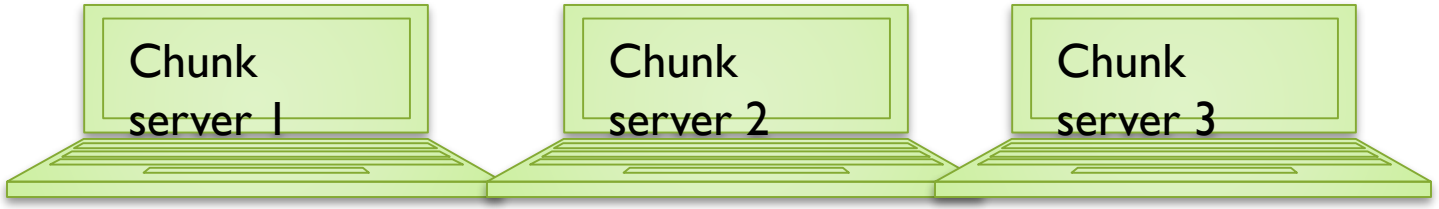

Architecture of a GFS system

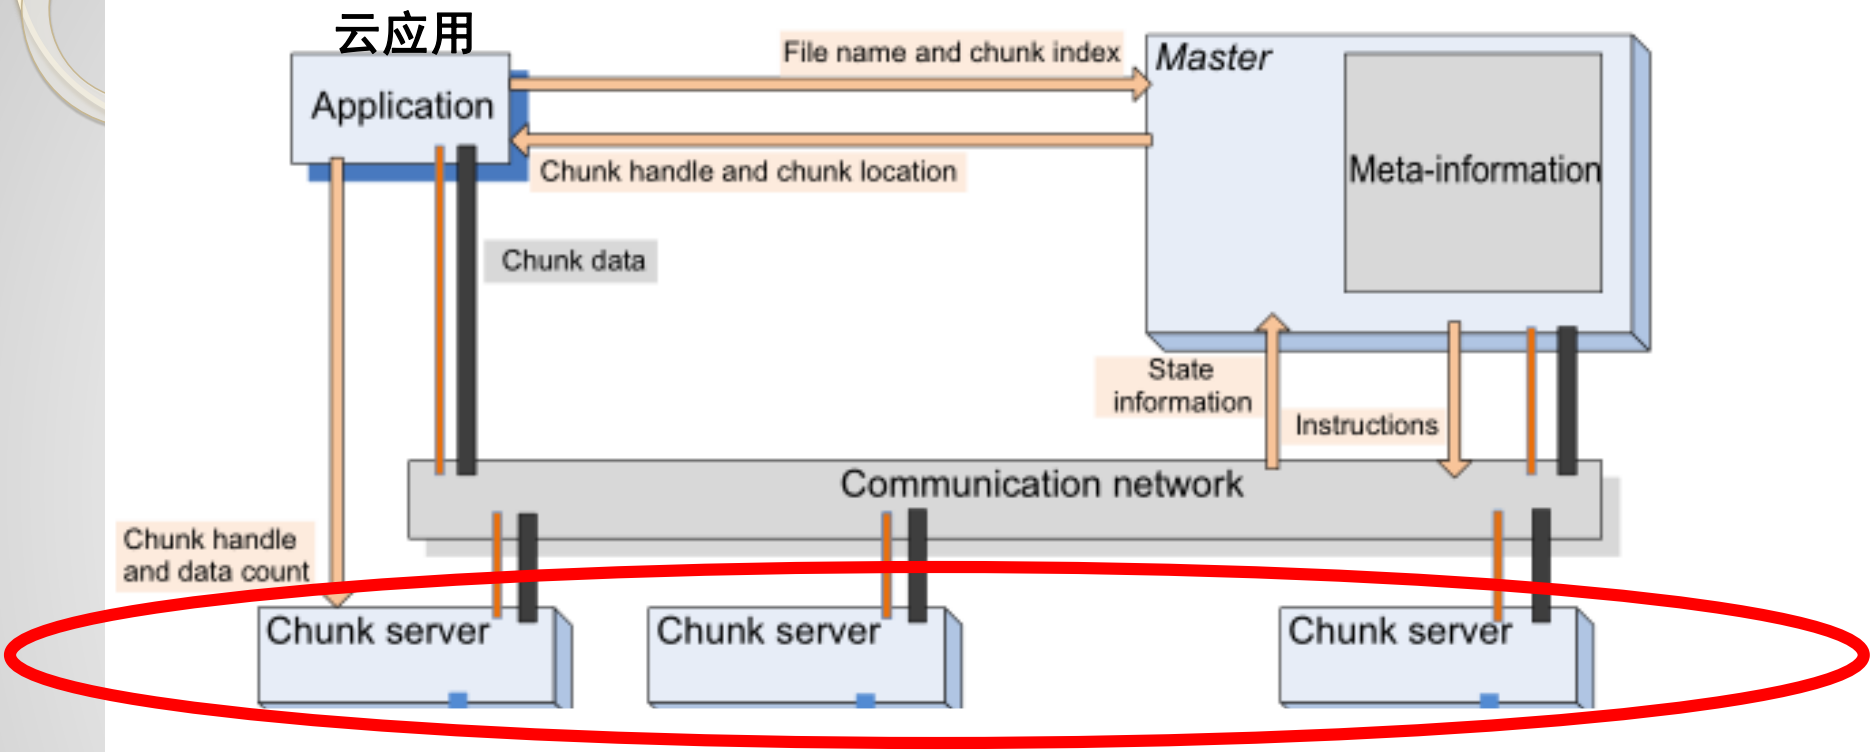

- The chunks are stored on some computers called the « **chunk servers** ».
- The **chunk servers** run the Linux operating system.

Architecture of a GFS system

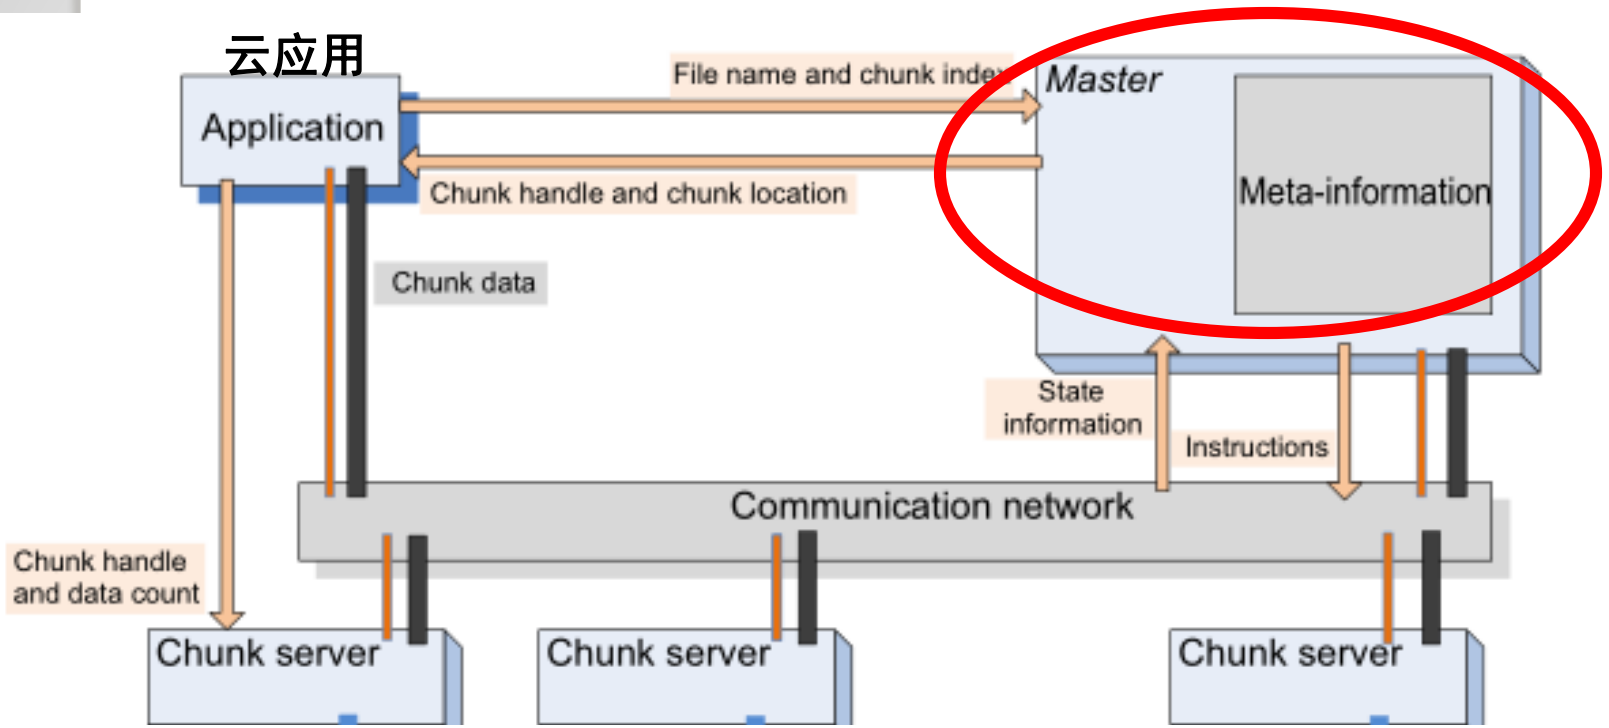

- The **Master** computer keeps information about which **files** are contained in each chunk, and which **chunk server** has a **copy** of each **chunk**.
- It keeps information about the **states of chunk servers (e.g.** failure?**)**.
- It keeps information for **access control** (访问控制)

How to access data?

Three steps:

1. a **cloud application** asks the **master** computer where the chunk is located,
2. The **master** computer tells the location of the chunk on a chunk server,
3. Then, the **cloud application** contacts the chunk server to access the data.

A **cloud application (云应用)** always read/send data from/to the closest storage computer (最近的存储计算机)in the network (to ensure fast access to the data).

How to write data?

Three steps

1. a **cloud application** asks the **master** computer for the permission to modify a chunk,
2. The **master** computer says **Yes!** or **No!**
   If **Yes**, it gives **the permission to modify the data for a limited amount of time.**
   It also indicates the location of the chunk (on a chunk server).
3. Then, the **cloud application** directly contacts the **chunk server** to modify the chunk.

Note that the **master** computer never gives the permission for modifying a chunk to two cloud applications at the same time to avoid **conflicts**.

How to ensure consistency (一致性)?

Consistency: all copies of a **chunk** are the same (一致性)

How we do that?

- When **a modification** is made to a **chunk**, it is performed on all copies of the chunk (on different chunk servers).
- **Then,** the **modification** becomes **visible** to other cloud applications
- If a **failure** occurs during modification, the modification is cancelled on all chunk servers.

How is the performance?

- The master server does not participate much to the process of reading/writing data.
- Why?
  - to make sure that reading and writing is **fast**.
  - If the master server would check all the read/write operations, it would be **slow**.
  - Using GFS, **reading data** can be **almost as fast** as reading data from a desktop computer.
  - **Writing** is slower than reading.
    It can be **five times** slower.

How big is a chunk?

- The size of chunks is **64 MB** by default (默认情况下) but it can be changed.
- **Why 64 MB?**
  - **If small**, there could be too many chunks! It can be **difficult to manage** them.
  - **If large**, it is more likely **that a computer will perform many operations** on the same chunk. But accessing a chunk may be slower.

More details

- The master server creates the chunks and files.
- The master server uses a **journal**.
- This **journal** is used to restore data in the case of a failure.
- GFS has a « **recovery mechanism** » to restore data when a failure occurs, that we will not discuss.

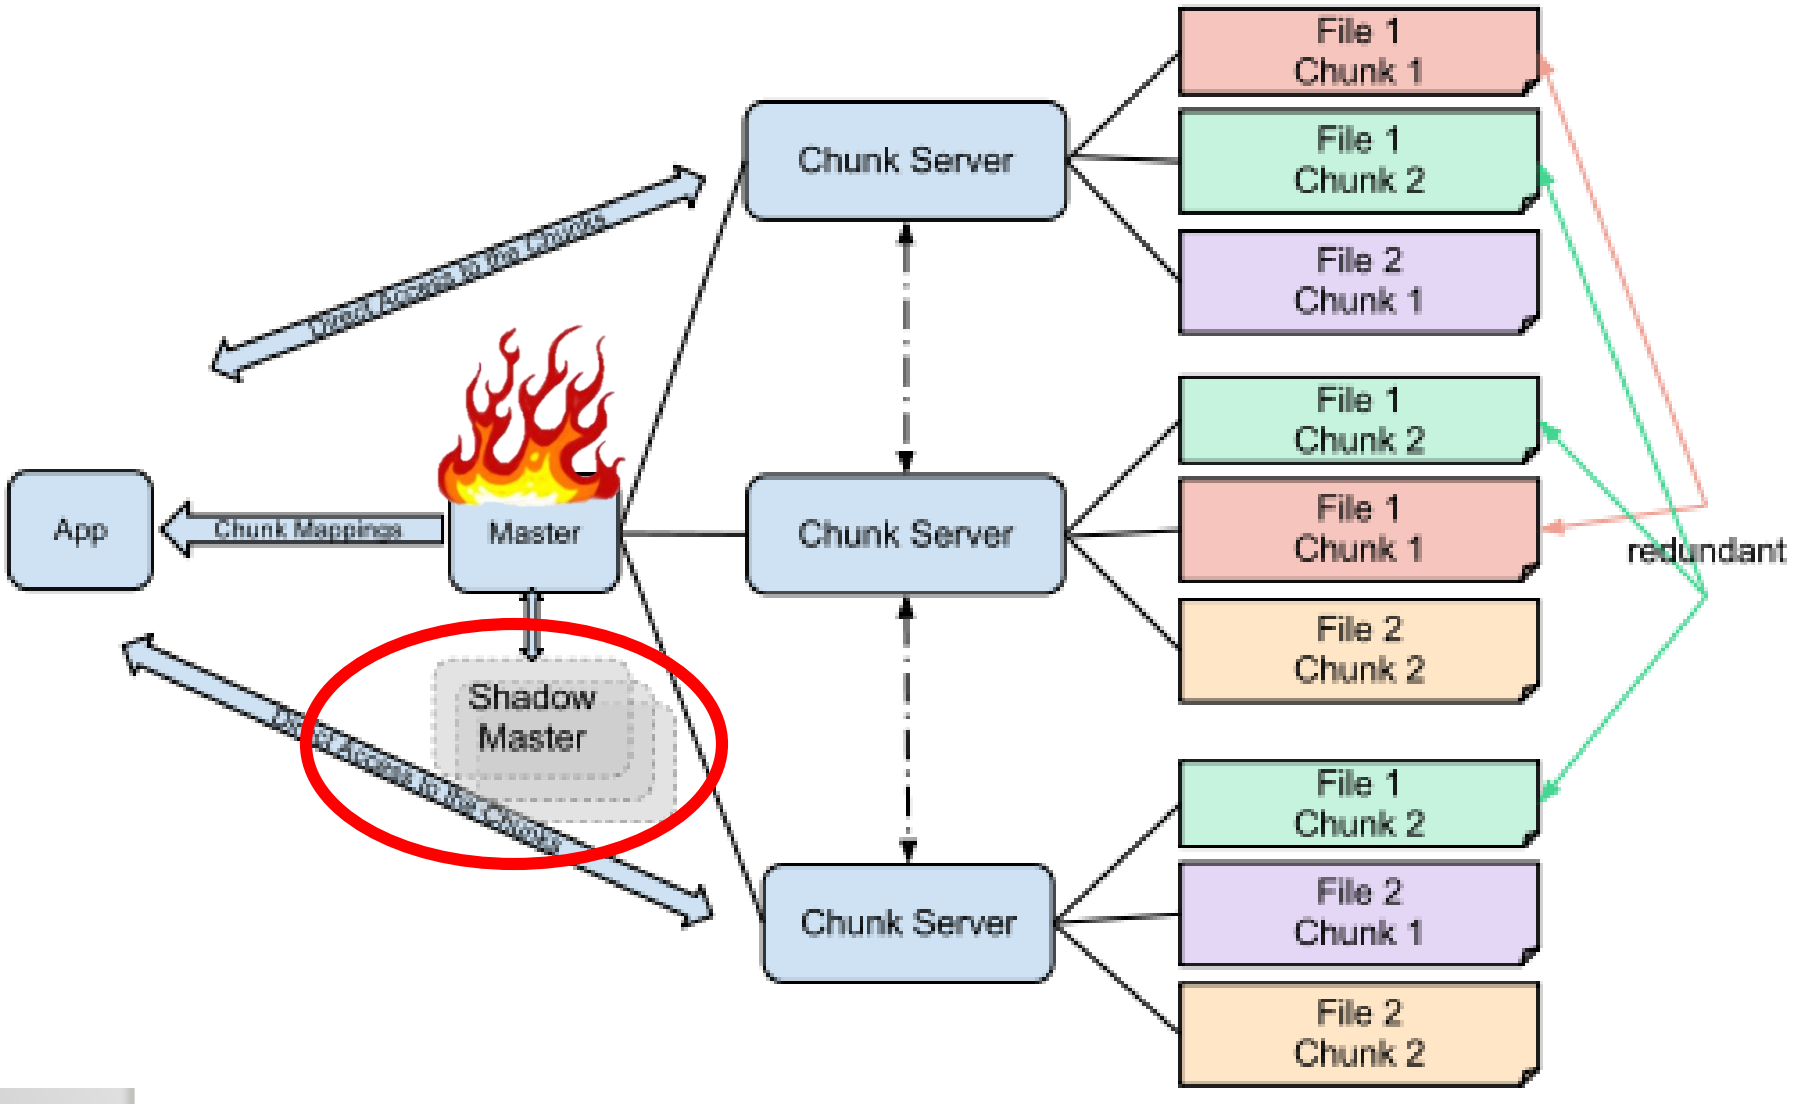

- There is always one master server.
- But there is also be a « **shadow master server** » that keeps a copy of the data of the master server in case that it fails.

Picture: wikipedia

APACHE HADOOP

8.5

Introduction

- **Hadoop** is an open-source system for cloud computing based on **MapReduce**.
- **It has two components:**
  - A Map Reduce engine
  - A storage system to store the data
- Different **storage systems** can be used with Hadoop:
  - Hadoop File System (HDFS)
  - Cloudstore
  - Amazon S3
  - …

The Hadoop File System (HDFS)

- **HDFS** was written using the **Java** programming language (java编程语言).
- The **Hadoop file system** is similar to GFS.
  - **Chunk servers** are called **Data Nodes** or **Slave Nodes**
  - The **master server** is called a **Name Node**
  - **Chunks** are called **blocks**
- How data is written in the cloud is slightly different.
  - **GFS** allows to write small amount of data quickly.
  - **HDFS** allows only to read/write one **full** block at a time.
- HDFS is designed for MapReduce, while GFS is more general.

Architecture of a HDFS cloud

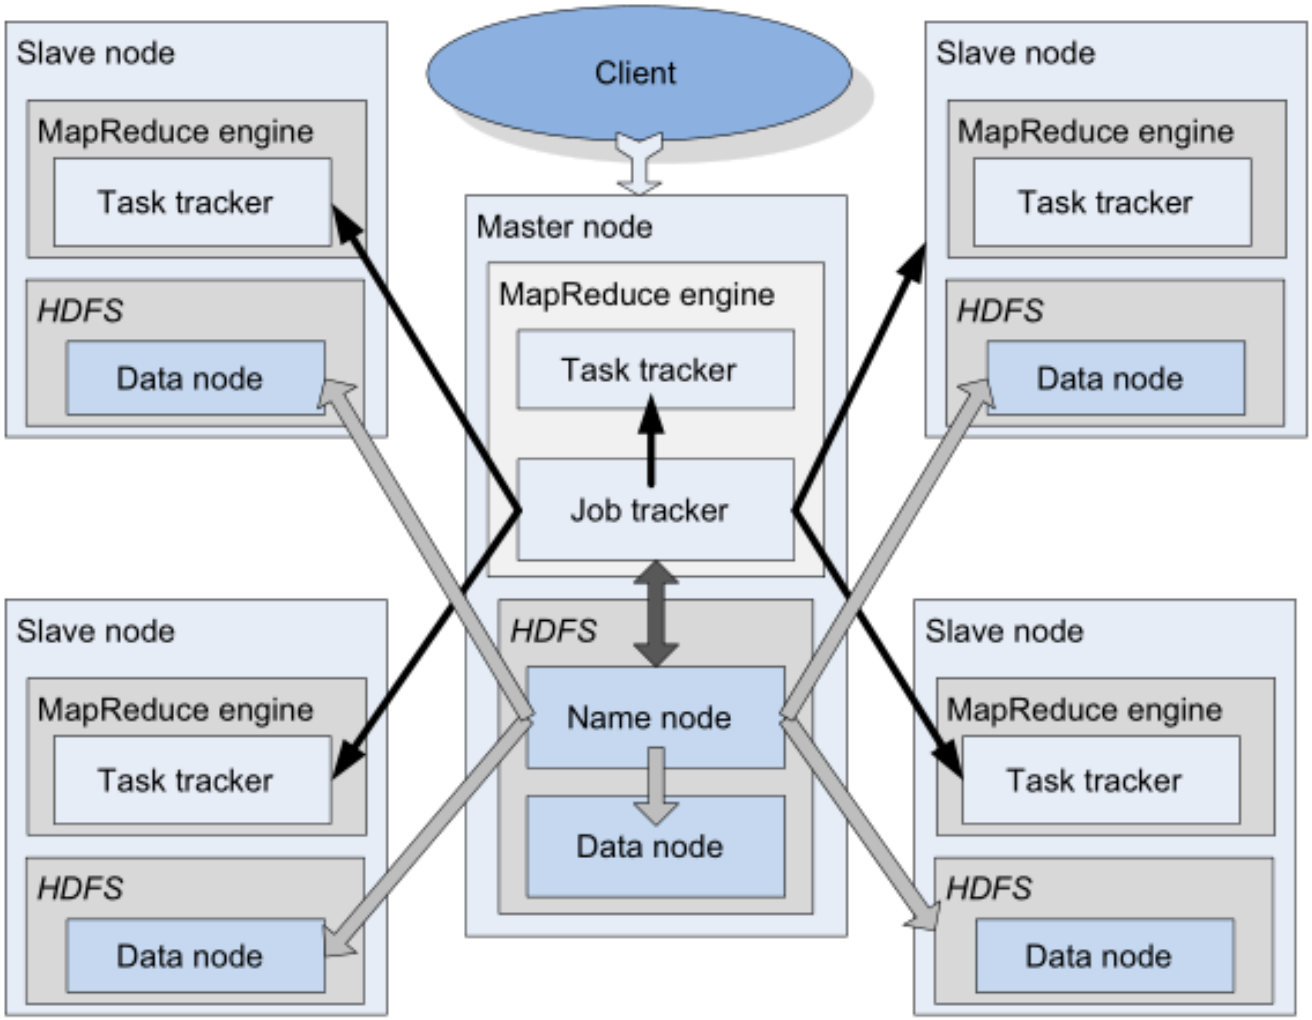

Here, there are 1 master computer and 4 slave computers.
A cloud application (**client**) can send map reduce jobs to the **Name Node** (master). The **Name Node** then executes the job on the **Slave Nodes**.

Architecture of a HDFS cloud

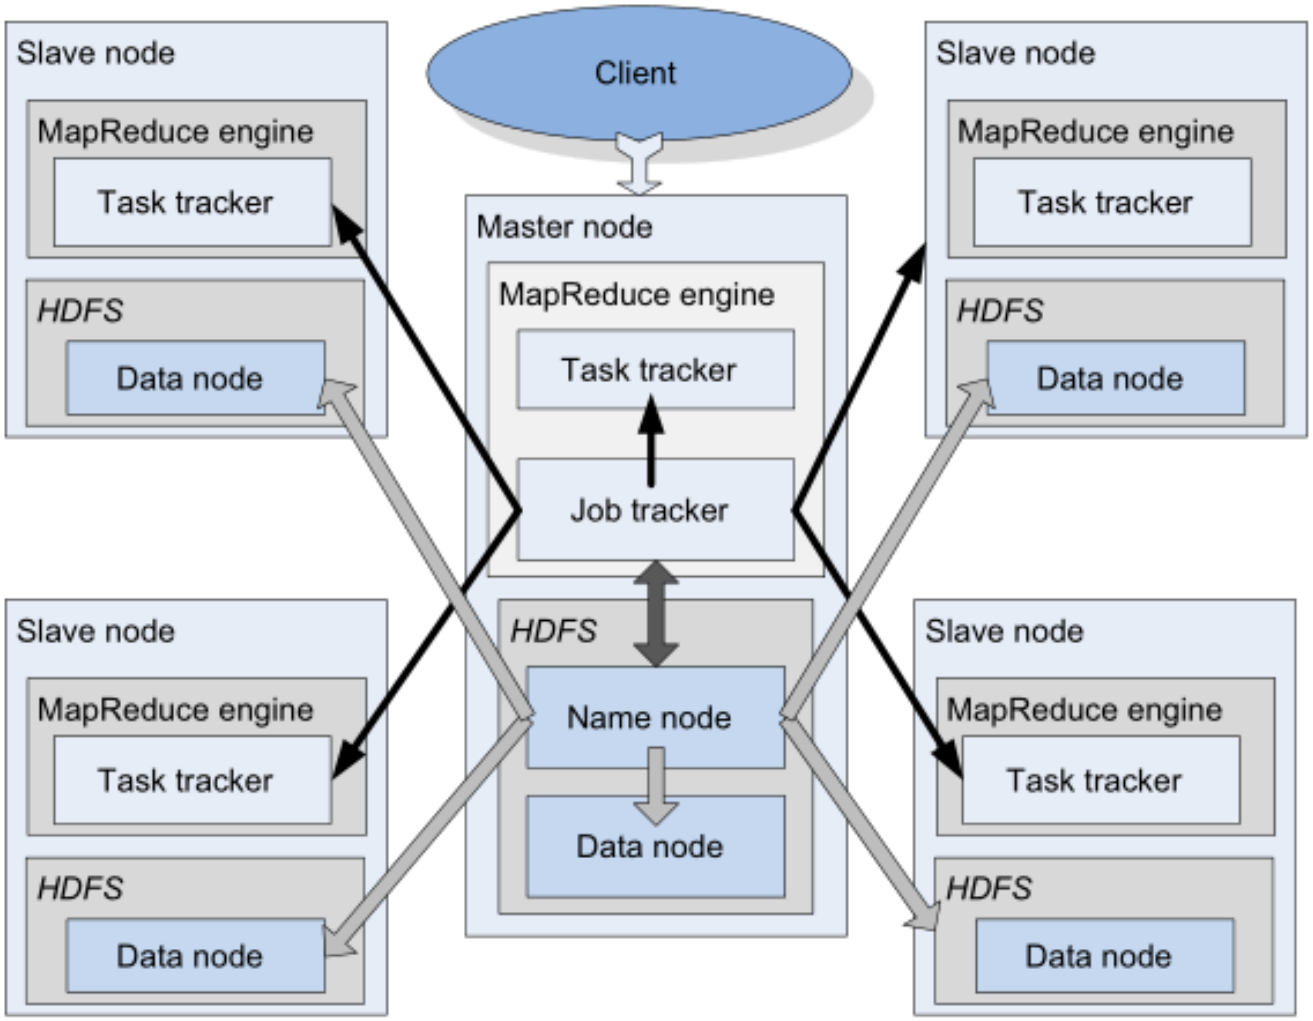

It can be observed in this illustration that the node performing the calculation are also the node that store the data.
This is different from GFS.

Conclusion

- We discussed the **importance of how data is stored** in the cloud.
- We discussed some **cloud storage systems** as examples: **GFS, HADOOP**...
- There exists other ways of storing data.
- **Big cloud companies** such as **Amazon, Alibaba** and others often develop their own storage solutions to have more control over performance.
- Many of these technologies are not public.

Conclusion

- If you are interested by cloud computing, there are a lot of things to learn about the various storage solutions.
- Today we have discussed cloud storage.
- Next week, it will be the **last lecture for this course**.
  - We will conclude
  - We will discuss more about the upcoming final exam.

References

- Chaptre 8. D. C. Marinescu. Cloud Computing Theory and Practice, Morgan Kaufmann, 2013.

Final exam

- **Date to be announced**
- **120 minutes**.
- It is a **closed-book exam.**
- Questions will be approximately evenly distributed between the different topics that we have discussed.

  8 lectures = 8 to 10 questions

Final exam

- Answers must be written in English.
- Some typical questions in my exams:
  - What is the advantages/disadvantages of using X instead of Y ?
  - When X should be used?
  - How X works ? or why X is designed like that?
  - There could be 1 question that is similar to assignments.

Final exam

- **If you are not sure about the meaning of a question in the final exam** because of English, you may raise your hand to ask me.

- **No electronic devices** are allowed.
- A pen/pencil/eraser can be used during the exam.
- Bring your **student ID card**.