Course Objective:

- To learn various data storage methods on cloud.

UNIT II
Data Storage and Cloud Computing

POINTS TO COVER
PART I : DATA STORAGE

- Introduction to Enterprise Data Storage
  - Direct Attached Storage
  - Storage Area Network
  - Network Attached Storage
- Data Storage Management
- File System
- Cloud Data Stores
- Using Grids for Data Storage.

Reading Material :
Srinivasan, J. Suresh, "Cloud Computing: A Practical Approach for Learning and Implementation",
Chap 11, pg 139 - 147

POINTS TO COVER
PART II CLOUD STORAGE

- Data Management
- Provisioning Cloud storage
- Data Intensive Technologies for Cloud Computing.

Reading Material :
Srinivasan, J. Suresh, "Cloud Computing: A Practical Approach for Learning and Implementation",
Chap 12, pg 150- 156

POINTS TO COVER
PART III CLOUD STORAGE FROM LANs TO WANs

- Cloud Characteristics
- Distributed Data Storage

- **Exemplar/Case Studies** Online Book Marketing Service, Online Photo Editing Service

Reading Material :
Srinivasan, J. Suresh, "Cloud Computing: A Practical Approach for Learning and Implementation",
Chap 13, pg 159 – 165

CLOUD COMPUTING

A practical approach for learning and implementation

A. Srinivasan | J. Suresh

ALWAYS LEARNING

PEARSON

- Storage is a resource to be allocated to organizations to add more value.
- Data storage management includes a set of tools to configure, backup, assign to users according to defined policies.
- Service level agreements (SLA) support clear business objectives, reduced risk mitigation levels and legal issues.

POINTS TO COVER

PART I : DATA STORAGE

- Introduction to Enterprise Data Storage
  - Direct Attached Storage
  - Storage Area Network
  - Network Attached Storage
- Data Storage Management
- File System
- Cloud Data Stores
- Using Grids for Data Storage.

Introduction to Enterprise Data Storage

- To build an effective storage system understanding storage system is must.
- This will yield cost effective, high performance and ease in managing the systems.
- The various types of storage subsystems are:
  1. Direct Attached Storage [DAS]
  2. Storage Area Network [SAN]
  3. Network Attached Storage [NAS]

Note

- DAS is the basic in a storage system and employed in building SAN and NAS either directly or indirectly.
- NAS is the top most layer, having SAN and DAS as its base.
- SAN lies between a DAS and a NAS.

STORAGE TYPES COMPARISON

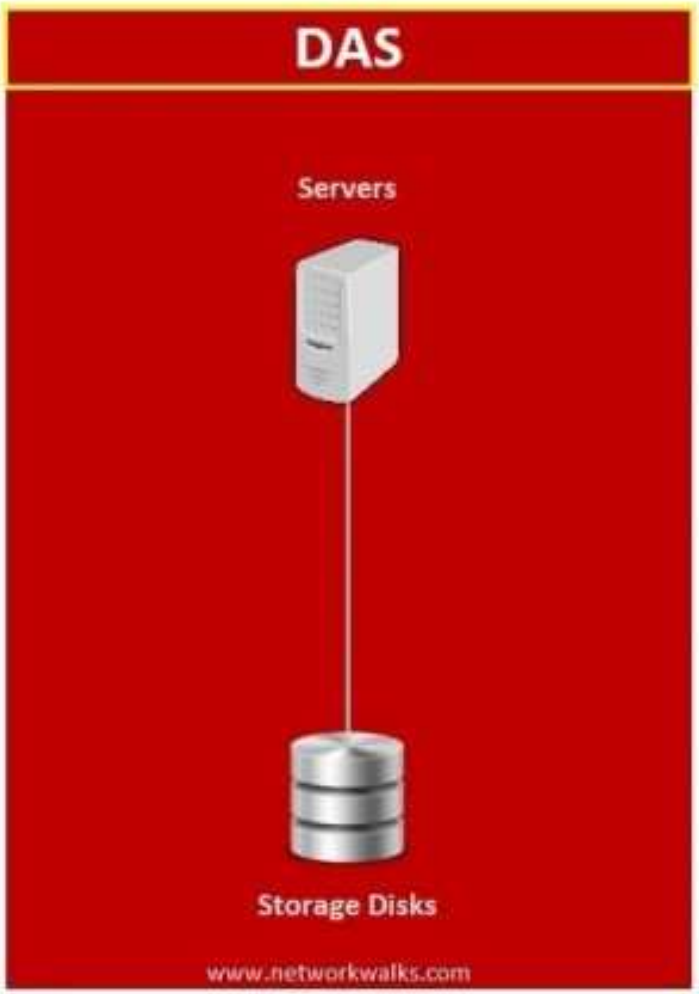

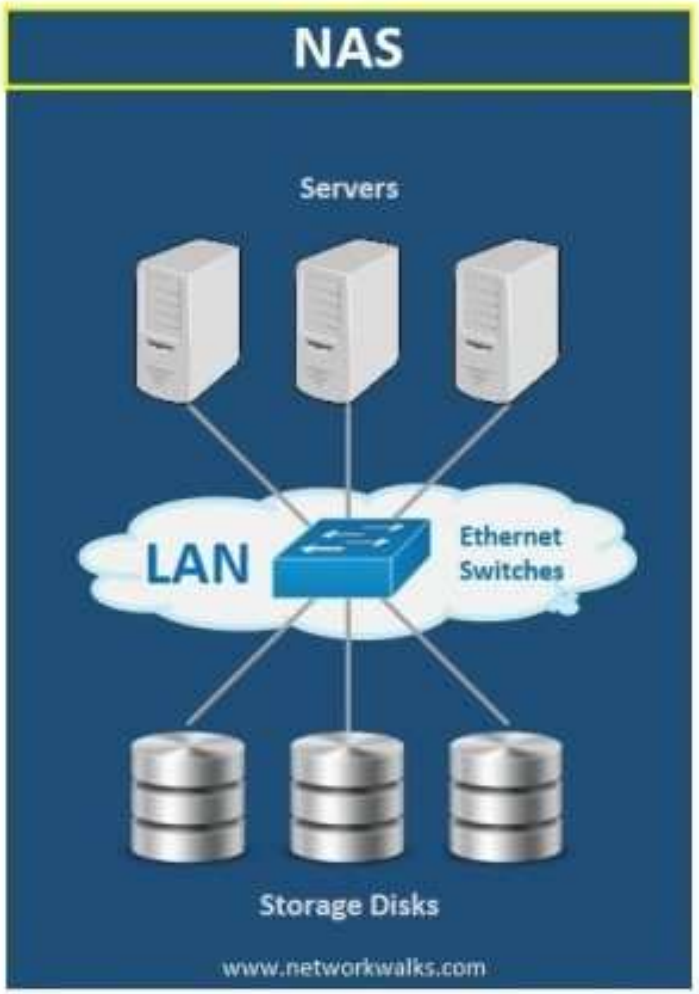

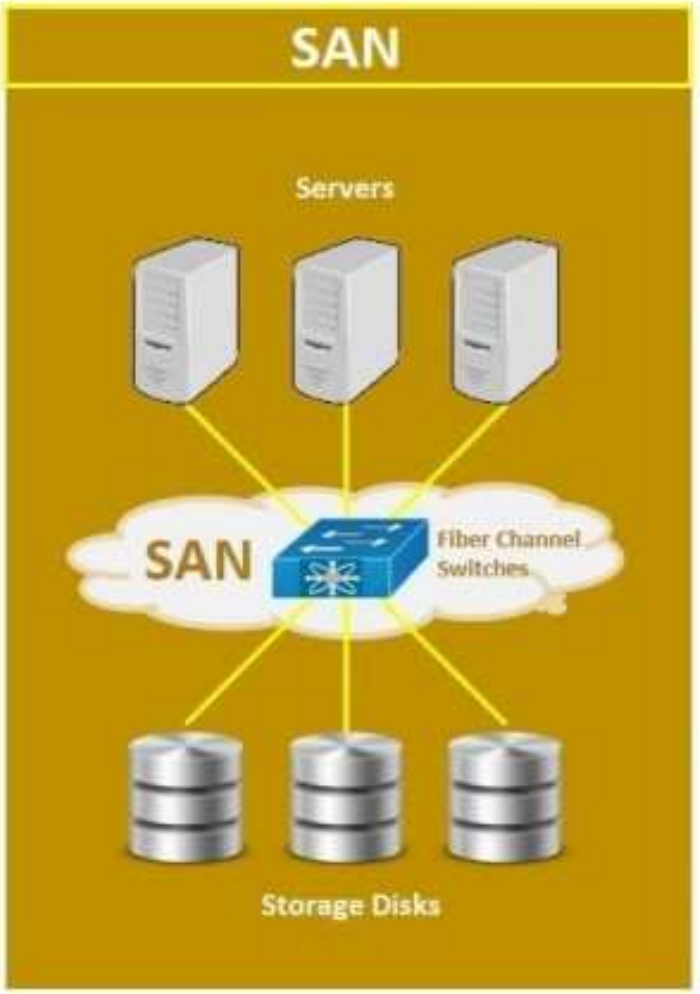

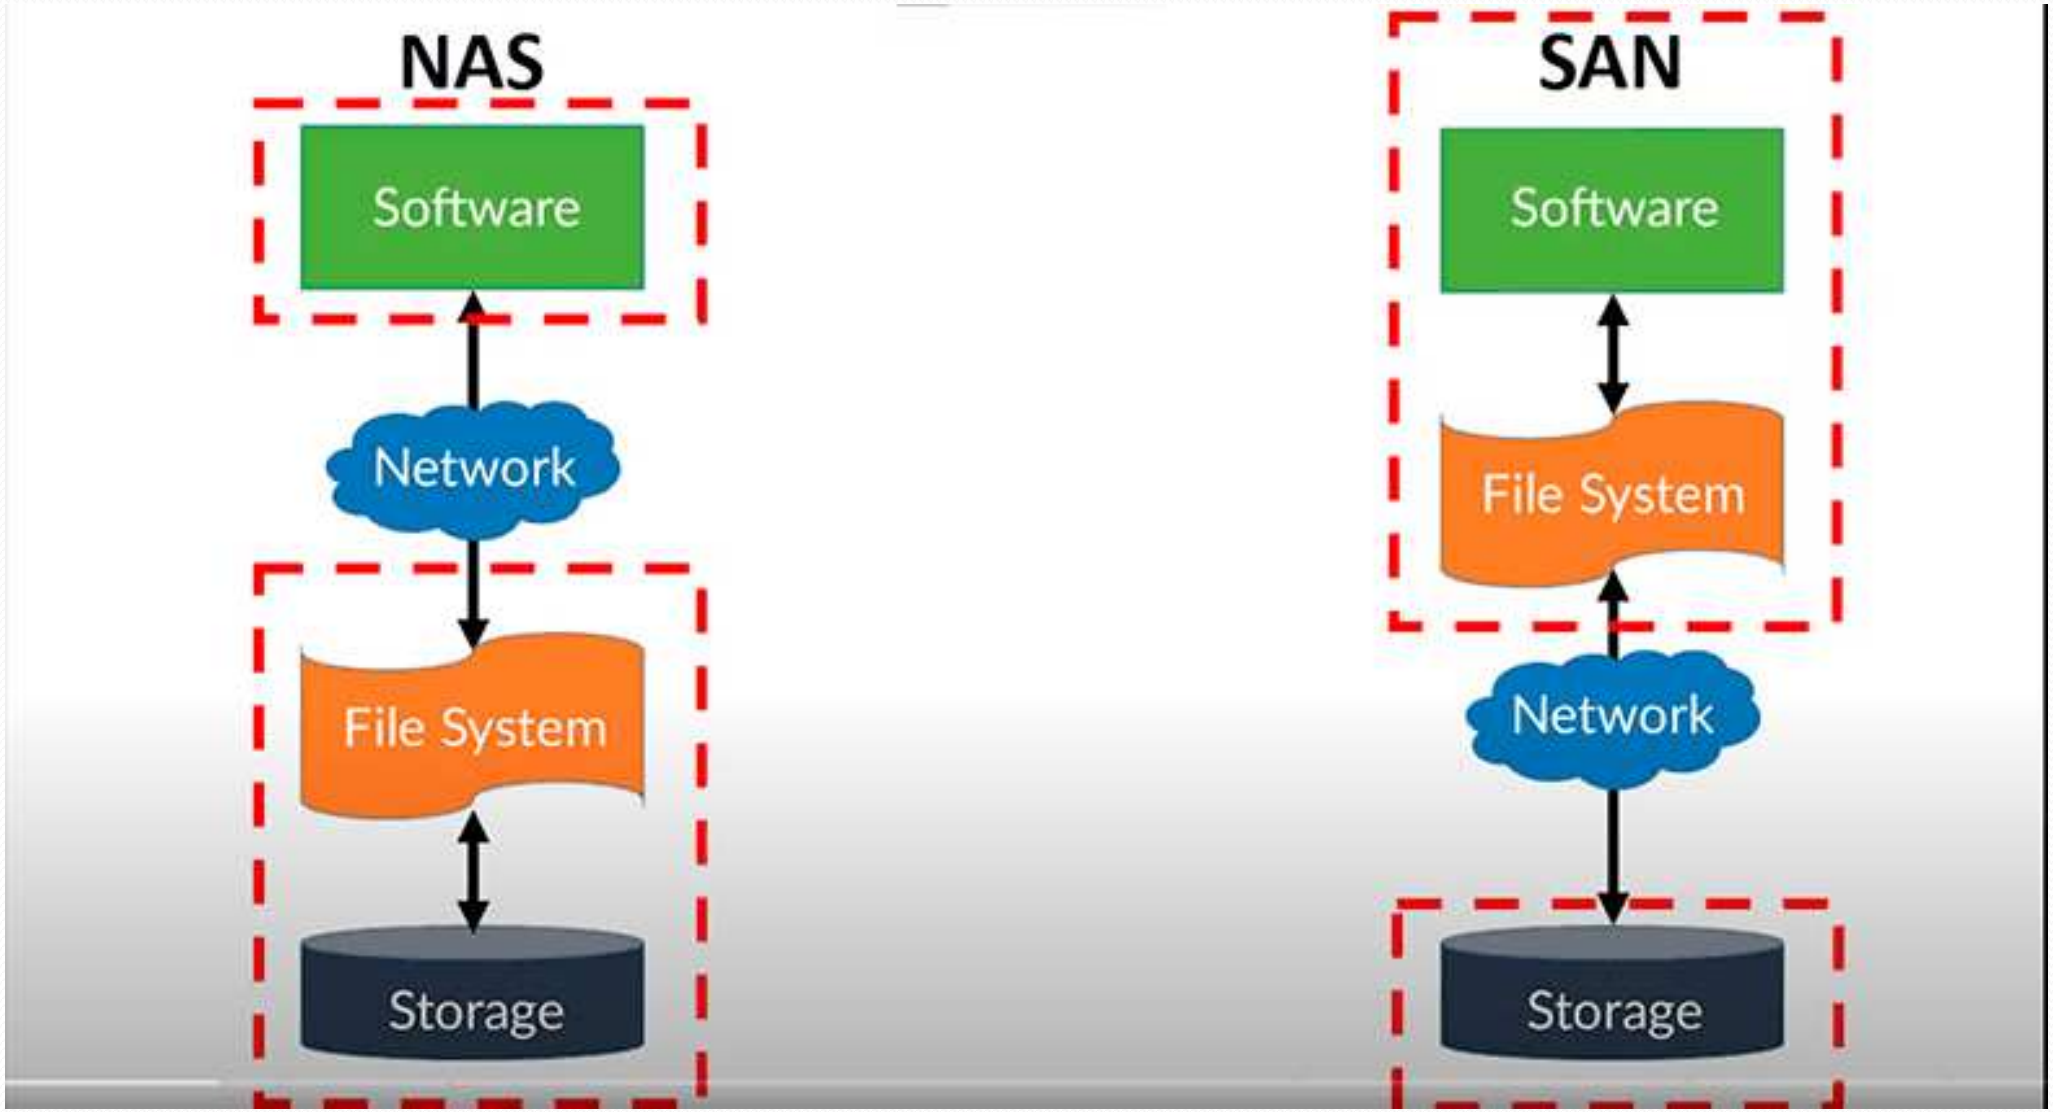
NAS
Software
Network
File System
Storage
SAN
Software
File System
Network
Storage

DAS: Direct Attached Storage

- Basic storage system :
  - To provide block-level storage
  - Use for building SAN and NAS
- The performance of SAN and NAS depends on DAS.
- Performance of DAS will always be high
  - it is directly connected to the system
- Storage devices used to build a DAS storage subsystem are
  - SCSI [Small Computer System Interface]
  - PATA [Parallel Advanced Technology Attachment]
  - SATA [Serial Advanced Technology Attachment]
  - SAS FC, Flash and RAM.

SCSI (Small Computer System Interface)

- *Used to connect and communicate between computers and peripheral devices, such as hard disk drives, tape drives, CD/DVD drives, and scanners.*
- *Originally developed as both a protocol and a parallel physical interface.*
- *Today, SCSI has retained and expanded the protocol, but replaced the parallel physical interface with multiple different types of serial interfaces.*
- *The SCSI protocol is now transported over a Serial Attached SCSI bus (using SAS and SPL), in Fibre Channel environments (using FCP), and over IP based networks (LAN/WAN, using iSCSI).*
- *SCSI was originally developed in the 1980s as a standard for connecting peripherals to computers, especially in high-performance and server environments.*
- *It has gone through various revisions and updates and is now used primarily for block storage devices and tape storage devices.*
- *SCSI is known for its flexibility, scalability, and performance, making it popular in enterprise and professional settings where high-speed data transfer and reliable storage connectivity are required.*
- *The SCSI protocol continues to be used in enterprise and professional environments where high-performance storage connectivity is crucial.*

PATA, SATA

- **PATA** stands for **Parallel Advanced Technology Attachment** which is a bus interface used for connecting secondary storage devices like hard disks, optical drives. It was first introduced in the year 1986 by Western Digital and Compaq. It was later replaced by SATA.
- **SATA** stands for **Serial Advanced Technology Attachment** is a bus interface that connects hard disks, optical drives. It was introduced in 2001 after **PATA** was slowly declining its demand by Serial ATA Working Group. SATA has more advantages than PATA making its demand more.

S.NO	PATA	SATA
1.	PATA stands for Parallel Advanced Technology Attachment.	SATA stands for Serial Advanced Technology Attachment.
2.	It is a 40 pin connector.	It is a 7 pin connector.
3.	It is high in cost.	It is cheaper in cost.
4.	The speed of data transfer is lower.	The speed of data transfer is higher.
5.	Power consumption is more.	Power consumption is less.
6.	The cable size is bigger.	The cable size is smaller.
7.	It doesn't have the feature of hot swapping.	It has the feature of hot swapping.
8.	External hard drives cannot be used.	External hard drives can be used.
9.	Multi bits data transfer at a time	1-bit data transfer at a time

Serial-Attached SCSI (SAS)

- Serial-Attached SCSI (SAS) is a method used to access computer peripheral devices that employs a serial -- one bit at a time -- means of digital data transfer over thin cables.
- In the business enterprise, Serial-Attached SCSI is especially of interest for access to mass storage devices, particularly external hard disk drives and magnetic tape drives.
- SAS is specified in the American National Standards Institute standard called Serial-Attached SCSI (Small Computer System Interface), also known as ANSI/InterNational Committee for Information Technology Standards (INCITS) 376-2003.

SAS…

- SAS is a protocol for point-to-point serial transmissions between storage devices and the computers they are storing data for.
- Point-to-point means that all data transfers across SAS are sent directly between the two communicating entities -- storage device and computer -- which are connected by a physical cable.
- Serial means that all data sent using SAS is transmitted a single bit at a time, in sequence.
- The SAS protocol is implemented on computers to use a dedicated link among the computer and disk drives, tape drives and any other SCSI storage devices that are connected to the computer's host bus adapters (HBAs) over a serial interface.

- A SAS system includes the following four basic components:

1. **Targets** are storage drives that are connected to the SAS system.
2. **Initiators** are used to send requests and handle responses from target devices. They are connected to the SAS device using a dedicated link.
3. **Service delivery subsystem** is the component of the SAS system -- typically a cable, but it could include expander hardware for port expansion or a system backplane -- that links the targets, or storage devices, with the initiator of the SAS system.
4. **Expanders** are an optional component of port multiplier hardware that enable port expansion in SAS systems so that multiple devices can be included in the SAS system. Expanders enable SAS systems to address as many as 65,535 storage devices.

Flash

- Flash memory is secondary memory and so it is not volatile which means it persists the data even if there is not an electrical supply provided.
- This flash memory works on the principle of EEPROM. EEPROM stands for Electrical Erasable Programmable Read-Only Memory.
- ROM operation can only one time write and many times read and we can’t erase it.
- But Flash Memory can be erased multiple times and update the data or program integrated into it. So it gives flexibility to the updation of the program but ROM has no such type of feature.

Features of Flash Memory

- Non-volatile: There is no loss of data when there is no electricity supply.
- Solid-state: It is SS technology so it is faster than HDD type storage.
- Fast access times: It supports solid-state technology so it has faster access time.
- Large storage capacity: Flash memory devices can store large amounts of data, from a few GB (Gigabytes) to several TB(Terabytes).
- Low power consumption: It is not based on header like HDD so no mechanical components in flash memory so it uses less amount of electricity from read the data.
- Flexibility towards Erase and write operations: Flash memory can be erased electrically multiple times and read multiple times so flexibility towards read/write operation is more in flash memory.

Applications of Flash Memory

- Used in SSDs: Flash memory is used in SSDs to increase the speed of read/write of operations.
- Embedded systems: Flash memory is used in embedded systems. Examples: digital cameras, camcorders, MP3 players etc.
- Smartphones and tablets: Flash memory is used in smartphones and tablets.
- USB drives: Flash memory is commonly used in USB drives.

SAN: Storage Area Network

- When multiple hosts want to connect a single storage device, then SAN is used.
- SAN provides block-level storage
  - Simultaneous access is not permitted
    - hence it is suitable for clustering environment.
- SAN technologies are
  - FC (Fibre Channel),
  - iSCSI (Internet SCSI) and
  - AoE (ATA over Ethernet).

NAS: Network Attached Storage

- For fi le-level storage, NAS is used.
- SAN and DAS act as base system for NAS.
- NAS is also called as ‘File Server’.
- The main advantage of NAS :
  - multiple hosts can share a single volume at the same time
- *When using SAN or DAS only one client can access the volume at a time.*

DATA STORAGE MANAGEMENT

- Data storage is expensive; therefore, storage administrators are trying to use tiered storage.
- Using fiber channel for storing data for a network user gives better performance but storage devices used are small and are expensive.
- SAN or DAS is cost effective performance-wise it is of lower grade.
- Today IT organizations are implementing tiered storage as a mix of storage technologies that meet the performance needs and are cost effective.

DATA STORAGE MANAGEMENT

1. Data Storage Management Tools
2. Storage Management Process
3. Data Storage Challenges
4. Unified Storage

1.Data Storage Management Tools

- Maintaining storage devices is a tedious job for storage administrators.
- They adopt some utilities to monitor and manage storage devices.
- Management level tasks are configuration, migration, provisioning, archiving and storage monitoring/reporting.
- Storage Resource Management (SRM) tools include configuration tools, provisioning tools and measurement tools.
  - **Configuration tools** handle the set-up of storage resources. These tools help to organize and manage RAID devices by assigning groups, defining levels or assigning spare drives.
  - **Provisioning tools** define and control access to storage resources for preventing a network user from being able to use any other user's storage.
  - **Measurement tools** analyze performance based on behavioral information about a storage device. An administrator can use that information for future capacity and upgrade planning.

Top 6 Storage Management Software

2.Storage Management Process

- Data storage management tools must rely on policies which governs the usage of storage devices and its procedures.
- Storage management encompasses three areas—change management, performance and capacity planning and tiering (tiered storage).
- The process used to request, schedule, implement and evaluate adjustments to the storage infrastructure is called **change management**.
  - The **change management** process defines the way a request is made and approved and documents the steps used to configure and provision the requested space on a storage array or server.
  - **Change management** may also document processes such as data migration and maintains the integrity and availability of that data for network users.
- **Performance and capacity planning** are used to measure the performance of a system in-terms of storage and utilization.
  - The result of performance and consumption analysis is used to make sensible decisions about subsequent storage purchases.

3.Data Storage Challenges

- In depth, understanding of storage devices
  - will minimize the risks, and
  - an administrator can easily handle challenges like finding out the reason for performance degrading, cost check, etc.
- Managing traditional storage devices is a complicated task because of high operations cost, performance and scalability issues.
- Some challenges are
  - massive data demand
  - performance barrier
  - power consumption and cost

Massive Data Demand

- An industry survey estimates the digital world to increase by 45 zettabytes by 2020,
  - one terabyte is equal to 1024 gigabytes,
  - one petabytes is equal to 1024 terabytes,
  - one exabytes is equal to 1024 petabytes
  - one zettabytes is equal to 1024 exabytes.

- **Reality: 149 zettabytes (As of 2024)**

Performance Barrier

- Rapid growth in data has caused a parallel increase in the size of databases.
- In the traditional storage method, the response time taken for queries is slow and it should be increased.
- Be it a social networking site, an enterprise database or a web application, all requires faster disk access to read and write data.

Power Consumption and Cost

- Because of increase in storage demands, IT organizations and data centres need larger storage with minimal cost.
- Performance lags with minimal cost but has other expenses like licensing and maintenance.
- Apart from this, other factors such as power consumed by storage devices, cooling systems, man power for managing it and space for data centres are to be considered.

Solution : Unified Storage

- A new innovative solution ‘Unified Storage’ is developed to address these issues.
- Basically this type of storage solution is a combination of NAS and SAN and termed as NUS (network unified storage).
- This type of storage system handles both file and block level accessing and hence storage devices can be accessed by single and multiple hosts.
- The main advantage of this system is reduced cost and it supports fibre channel and iSCSI.

POINTS COVERED
PART I : DATA STORAGE

- ✓ Introduction to Enterprise Data Storage
  - ✓ Direct Attached Storage
  - ✓ Storage Area Network
  - ✓ Network Attached Storage
- ✓ Data Storage Management
- File System
- Cloud Data Stores
- Using Grids for Data Storage.

Transparency

Goals Of DS

2. Transparency :

- How to achieve the single system image?
- How to "fool" everyone into thinking that the collection of machines is a "simple" computer?
- To hide the fact that DS's processes and resources are physically distributed across multiple computers
- A DS that is able to present itself to users and applications as if it were only a single computer system is said to be ***transparent***.

Transparency in a Distributed System

Transparency	Description
Access	Hide differences in data representation and how a resource is accessed
Location	Hide where a resource is located
Migration	Hide that a resource may move to another location
Relocation	Hide that a resource may be moved to another location while in use
Replication	Hide that replicas of a resource exist
Concurrency	Hide that a resource may be shared by several competitive users
Failure	Hide the failure and recovery of a resource
Persistence	Hide whether a (software) resource is in memory or on disk

FILE SYSTEMS

- A structure used in computer to store data on a hard disk
- When we install a new hard disk, we need to partition and format it using a file system before storing data.
- There are three file systems in use in Windows OS:

1. NTFS
2. FAT32
3. FAT[rarely-used]

FAT File System

- File Allocation Table, 1977
- first devised for personal computer environment in the early years
- was the default file system for MS-DOS and Windows 9x operating systems
- Planned for systems with very small RAM and small disks.
- It required much less system resources compared to other file systems like UNIX.
- Three major variants: FAT12, FAT16 and FAT32
- FAT was replaced with NTFS as the default file system on Microsoft operating systems starting with Windows XP
- Essentially, the FAT system has made a come back.
- Thumb or flash drives have become very common and have smaller size that makes the FAT system useful.
- The smaller sizes are even formatted in FAT16.

NTFS [New Technology File System], 1993

- In the 1990s, Microsoft recognized that DOS based Windows was inadequate because of demands in business and industry.
- They started working for better software which can suit larger systems.
- Support Linux and BSD
- NTFS is much simpler than FAT.
- While files are used, the system areas can be customized, enlarged, or moved as required.
- NTFS has much more security incorporated.
- NTFS is not apt for small-sized disks.

Cloud File System

- In cloud file systems, the considerations are:
  - It must sustain basic file system functionality.
  - It should be an open source.
  - It should be developed enough that users will at least think about trusting their data to it.
  - It should be shared, i.e., available over a network.
  - It should be paralleling scalable.
  - It should provide honest data protection, still on commodity hardware with only internal storage.

- Organizations that use cloud computing outsource massive amounts of data and workloads to cloud providers.
  - Due to its low cost, lower management overhead and elasticity, organizations move towards using cloud computing.
- In cloud storage, systems host or consumers can find only corruption or loss of data from their service provider's report, when a system failure occurs.
  - This consumer–provider gap creates business risk and complicates compliance SLAs.

Expectations from Cloud File System….

- A cloud file system should be scalable enough to adopt large organizations file systems under different workloads with good performance requirements.
- Cloud file systems should have high throughputs then local file systems.
- Cloud file system should have minimal operation latency.
- The system should also be scalable to multiple hosts operating in parallel.
- Transparency and backwards compatibility is important to facilitate migration to the cloud with less effort.

Some Cloud File Systems….

- Following are some of the cloud file systems:

1. Ghost File System
2. Gluster File System
3. Hadoop File System
4. XtreemFS: A Distributed and Replicated File System
5. Kosmos File System
6. CloudFS

1.Ghost File System

- Ghost cloud file system is used in Amazon Web Services (AWS).
- It gives high redundant elastic mountable, cost-effective and standards-based file system.
- A fully featured scalable and stable cloud fi le systems is provided by ghost cloud file system.
- GFS (Ghost File System) run over Amazon's S3, EC2 and SimpleDB web services.
- When using GFS, user can have complete control of the data and can be accessed as a standard network disk drive.

Benefits of Ghost CFS

- Elastic and cost efficient: Pay for what you use from 1 GB to hundreds of terabytes.
- Multi-region redundancy: Aiming to take advantage of AWS's 99.99% availability
- Highly secure: Uses your own AWS account (ghost cannot access your data).
- No administration: Scales elastically with built in redundancy—no provisioning or backup.
- Anywhere: Mount on a server or client or access fi les via a web page or from a mobile phone.

Features of Ghost CFS

- Mature elastic file system in the cloud.
- All files and metadata duplicated across multiple AWS availability regions.
- WebDAV (**Web Distributed Authoring and Versioning**) for standard mounting on any Linux, Windows or Mac server or client in the world.
- FTP access.
- Web interface for user management and for fi le upload/download.
- File name search.
- Side-loading of fi les from torrent and from URL.

Content Beyond Syllabus

- Gossip Protocol :

Gluster File System

- An open source, distributed file system capable of handling multiple clients and large data.
- Clusters storage devices over network, aggregating disk and memory resources and managing data as a single unit.
- Based on a stackable user space design and delivers good performance for even heavier workloads.
- Supports clients with valid IP address in network.
- Users no longer locked with legacy storage platforms which are costly and monolithic.
- Gives users the ability to deploy scale-out, virtualized storage, centrally managed pool of storage.
- Attributes of GlusterFS include scalability and performance, high availability, global namespace, elastic hash algorithm, elastic volume manager, gluster console manager, and standards-based.

Hadoop File System

- A distributed file system designed to run on commodity hardware is known as Hadoop Distributed File System (HDFS).
- In HDFS, files are stored in blocks ranging from 64 MB to 1024 MB.
- The default size is 64 MB.
- The blocks will be distributed across the cluster and replicated for fault tolerance.

XtreemFS: A Distributed and Replicated File System

- XtreemFS is a distributed, replicated and open source.
- XtreemFS allows users to mount and access files via WWW.
- Engaging XtreemFS a user can replicate the files across data centres to reduce network congestion, latency and increase data availability.
- Installing XtreemFS is quite easy, but replicating the files is bit difficult..

Kosmos File System

- Kosmos Distributed File System (KFS) gives high performance with availability and reliability.
- For example, search engines, data mining, grid computing, etc.
- It is deployed in C++ using standard system components such as STL, boost libraries, aio, log4cpp.
- KFS is incorporated with Hadoop and Hypertable

CloudFS

- CloudFS is a distributed file system to solve problems when file system is itself provided as a service.
- CloudFS is based on GlusterFS, a basic distributed file system, and supported by Red Hat and hosted by Fedora.
- There are really three production level distributed/parallel file systems that come close to the requirements for the cloud fi le systems: Lustre, PVFS2 and GlusterFS.

POINTS COVERED
PART I : DATA STORAGE

- ✓ Introduction to Enterprise Data Storage
  - ✓ Direct Attached Storage
  - ✓ Storage Area Network
  - ✓ Network Attached Storage
- ✓ Data Storage Management
- ✓ File System
- Cloud Data Stores
- Using Grids for Data Storage.

What is a data store? [AWS…..]

- A data store is a digital repository that stores and safeguards the information in computer systems.
- A data store can be network-connected storage, distributed cloud storage, a physical hard drive, or virtual storage.
- It can store both structured data like information tables and unstructured data like emails, images, and videos.
- Organizations use data stores to retain, share, and manage information across business units.

- You can use a data store to reliably save information in computer systems and prevent data loss.
  - Computer systems store information on persistent storage devices.
  - Persistent storage is nonvolatile, i.e. the storage retains the data even after a device's power is turned off.
  - This ensures that the computer system has access to the same data after it is powered on again.
- Businesses use data stores to manage, categorize, and streamline data for operations, analysis, reporting, and data retention, which is important for regulatory compliance.

Why is a data store important?

  - Data stores have several use cases, such as data created and consumed by applications, data archiving, data analytics, and disaster recovery.
- Due to the complexities in data requirements, companies use different types of data storage infrastructure to provide accessibility, redundancy, governance, and transparency.
  - For example, organizations use Amazon Elastic File System (**Amazon EFS**) for a serverless file system and Amazon Simple Storage Service (**Amazon S3**) for object storage.

Some Terms Related To Data Stores

- **Database**
- **Data stores compared to databases**
- **Data warehouse**
- **Data stores compared to. data warehouses**

Database

- A database is an organized storage system.
- Most databases are based on the relational database architecture.
- The relational database management system (RDBMS) allows users to store data in tables associated with specific data points.
- Organizations use databases to store transactional data, such as accounting, sales, and administrative logs.

Data stores compared to databases

- Discussions on data stores involve different methods to store and retrieve information.
- A database is one method that allows applications to store, share, and retrieve data easily.
- Unlike file systems, a database adheres to specific rules of how data is organized, formatted, and stored in the database.

Data warehouse

- A data warehouse is an extensive collection of business-related information acquired from various sources.
- Companies use data warehouses to support business intelligence and analytics.
- Business analysts and data scientists derive actionable insights from a data warehouse.

Data stores compared to data warehouses

- Data store is an umbrella term that includes the different hardware, technologies, formats, and architectures for storing and retrieving information.
- A data warehouse is a specific type of data store for consolidating analytical data for businesses.
  - For example, GE Renewable Energy uses **AWS Redshift** to gain new insights into its collected data.

How does a data store work?

- A physical data storage device is the underlying technology behind a data store.
- You can read and write information to the device in specific formats such as files, tables, or blocks.
- The device can be local, remote, or in the cloud.
- Large data stores are typically distributed across multiple physical devices in different geographic locations.
- Software systems and services abstract the underlying operations of the data store.

Some Examples Of Physical Devices

- **Flash and SSD drives**
  - A solid state drive (SSD) is a semiconductor technology that allows the writing and reading of data in flash memory chips. Flash storage technology was commercially available in pen drives before becoming an alternative to hard disk drives (HDD). Compared to an HDD, a physical SSD has no moving parts, which means it has faster performance and a longer lifespan.

- **Hybrid storage array**
  - Hybrid storage array is a physical storage setup that consists of an SSD and an HDD. While an SSD offers a low-latency operation, it costs much more per-unit storage than an HDD. Therefore, organizations use a hybrid storage array to balance performance, capacity, and cost.

- **RAID**
  - RAID stands for a redundant array of independent disks. It is a technology that keeps the same data in multiple places on an SSD.

AWS document reference ends….

CLOUD DATA STORES

- Distributed Data Store
- Types of Data Stores

CLOUD DATA STORES

- A data store is a data repository where data are stored as objects.
- Data store includes data repositories, flat files that can store data.
- Data stores can be of different types:
  - Relational databases (Examples: MySQL, PostgreSQL, Microsoft SQL Server, Oracle Database)
  - Object-oriented databases
  - Operational data stores
  - Schema-less data stores, e.g. Apache Cassandra or Dynamo
  - Paper files
  - Data files (spread sheets, flat files, etc)

Distributed Data Store

- A Distributed Data Store is like a distributed database where users store information on multiple nodes.
- These kinds of data store are non-relational databases that searches data quickly over a large multiple nodes.
  - Examples for this kind of data storage are Google's BigTable, Amazon's Dynamo and Windows Azure Storage.
- Some Distributed Data Stores use to recover the original file when parts of that file are damaged or unavailable by using forward error correction techniques.
- Others download that file from a diverse mirror.

Types of Data Stores

- Established IT organizations have started using advanced technologies for managing large size data, which come from social computing and data analysis applications.
- BigTable

BigTable

- A compressed, high performance and proprietary data storage system construct on Google File System, Chubby Lock Service, SSTable and a small number of other Google technologies.
- Developed in 2004 and is used in number of Google applications such as web indexing, Google Earth, Google Reader, Google Maps, Google Book Search, MapReduce, Blogger.com, Google Code hosting, Orkut, YouTube and Gmail.
- Advantages of developing BigTable : scalability and better performance control
- BigTable charts two random string values (row and column key) and timestamp into an associated random byte array.
- Designed to scale into the petabyte range across multiple machines and easy to add more machines and automatically start using resources available without any configuration changes.

Other similar softwares are as follows:

- Apache Accumulo: Construct on top of Hadoop, ZooKeeper. Server-side programming mechanism deployed in Java environment.
- Apache Cassandra: a distributed database management system built to manage a large amount of data over several cloud data centers. It uses Java.
- Hbase: Supports BigTable and Java programming language. designed to handle scaling across thousands of servers and managing access to petabytes of data. With the elasticity of Amazon EC2, and the scalability of Amazon S3, HBase is able to handle online access to massive data sets.
- Hypertable: Designed for cluster of servers especially for storage and processing.
- KDI: Kosmix stab to make a BigTable clone and is written in C++.

Dynamo: A Distributed Storage System

- A vastly offered, proprietary key-value structured storage system or a dispersed data store.
- It can act as databases and also distributed hash tables (DHTs).
- It is used with parts of Amazon web services such as Amazon S3.
- The most powerful relational database available in World Wide Web.
- Relational databases have been used a lot in retail sites, to make visitors browse and search for products easily.
- It is difficult to create redundancy and parallelism with relational databases which is a single point failure.
- Replication is also not possible.
- Dynamo is a distributed storage system and not a relational database.
- Similar to a relational database it stores information to be retrieved; however, it stores the data as objects and not as tables.
- The advantage: responsive and consistent in creating a distributed storage solution.

AWS……

- What is Amazon DynamoDB?

Handle more than 10 trillion requests per day and can support peaks of more than 20 million requests per second.

Secure your data with encryption at rest, automatic backup and restore, and guaranteed reliability with an SLA of up to 99.999% availability.

Focus on innovation and optimize costs with a fully managed serverless database that automatically scales up and down to fit your needs.

Integrate with AWS services to do more with your data. Use built-in tools to perform analytics, extract insights, and monitor traffic trends.

USING GRIDS FOR DATA STORAGE

- Grid Storage for Grid Computing
- Grid Oriented Storage (GOS)

Grid Storage for Grid Computing

- Grid computing established architecture which provides users and applications to use shared pool of resources.
  - The compute grid connects computers both desktops and servers and storage across an organization.
  - It virtualizes heterogeneous and remotely located components into a single system.
  - Allows sharing of computing and data resources for multiple workloads and enables collaboration both within and across organizations.
- Demand for storage requirement prevails in gird computing.
  - Storage for grid computing requires a common file system to present as a single storage space to all workloads.
  - Presently grid computing system uses NAS type of storage.
  - NAS provides transparency but limits scale and storage management capabilities.
- To set the unique demands of the compute grid on its storage infrastructure, storage for the grid must be abnormally flexible.
  - DAS is basically not an option.
  - Virtualization is a start, providing the single unit behavior where the global filing system requires data compute grid. Due to this, SAN architectures are used.
  - However, the scale of these SANs is beyond the capabilities of fibre channel.

Grid Oriented Storage (GOS)

- A dedicated data storage architecture connected directly to a computational grid
- Supports and acts as a data bank and reservoirs for data, which can be shared among multiple grid clients.
- A successor of Network-Attached Storage (NAS) products in the grid computing era
- Accelerates all kinds of applications in terms of performance and transparency.
- A GOS system contains multiple hard disks, arranged into logical, redundant storage containers like traditional file servers.
- Deals with long-distance, heterogeneous and single-image file operations.
- Acts as a file server and uses file-based GOS-FS protocol.
- Similar to GridFTP, GOS-FS integrates a parallel stream engine and Grid Security Infrastructure (GSI).
- GOS-FS can be used as an underlying platform to utilize the available bandwidth and accelerate performance in grid-based applications.

POINTS COVERED
PART I : DATA STORAGE

- ✓ Introduction to Enterprise Data Storage
  - ✓ Direct Attached Storage
  - ✓ Storage Area Network
  - ✓ Network Attached Storage
- ✓ Data Storage Management
- ✓ File System
- ✓ Cloud Data Stores
- ✓ Using Grids for Data Storage.

POINTS TO COVER
PART II CLOUD STORAGE

- Data Management
- Provisioning Cloud storage
- Data Intensive Technologies for Cloud Computing.

Reading Material :
Srinivasan, J. Suresh, "Cloud Computing: A Practical Approach for Learning and Implementation",
Chap 12, pg 150- 156

WHAT IS CLOUD STORAGE?

- Most of the organizations in an effort to cut cost are switching to taking **cloud computing** and **cloud storage solutions**.
- Cloud computing is a model which wraps around current technologies, for example, server virtualization to make use of resources more efficiently.
- The benefits of cloud storage are: a high point of **scalability** and **elasticity** alongside management.
- When virtualized storage is offered on demand in a network, an organization can be free from the need to buy its storage capacity before opting for data storage.

- Predicting growth of storage is **not possible**.
- IT organization has to provide enough capacity for storing data that are generated.
- With traditional storage hardware, expanding the capacity immediately is **difficult** and also **expensive**.
- Apart from this, maintaining these storage devices is **tedious** and **time consuming**.

Solution……

For solving these problems, the best practices adopted are :

- **Unpredictable storage growth**: IT organizations should constantly monitor storage consumption to track whether the actual growth rates are in line with initial projections.
- **Cost and complexity of conventional storage**: Enterprises must think Storage-as-a- Service solutions for remote and branch offices when it is possible.
- **Security**: As employees move between offices all over the world and take their data with them, enterprise should ensure that in-house and customer data is always protected and safe.

- IT organizations with less staff members cannot depute staff for their remote offices.
- Such IT organizations can end up with a series of problems in terms of structures that operate differently and inefficiently.
- To avoid these problems, the following solutions can be applied:
  - IT organizations should aim to centralize data storage and protection.
  - IT organizations should eliminate the need for personnel on-site and establish a single point-of-control.
  - Need a clear service level agreement between remote organization and the central organization.

- The abstract meaning for cloud is **any-to-any connectivity using network**.
- The cloud abstraction acts as a base upon which other features are built.
- In general, cloud model expands this base by adding **a pool of resources**.
- An important part of the cloud model is the recent innovation called **virtualization** that made the sharing of pool of resources **effective** and with **reduced complexity**.
- Cloud storage is nothing but virtualized storage on demand called as **Data storage as a Service (DaaS)**.

Cloud storage

- The data storage hosted remotely using data storage devices in WWW
- Maintained by the third party (service provider)
- A part of cloud computing
- Deployed using WAN infrastructure
  - includes hardware components such as switches and routers

Cloud Storage Deployment

- Cloud storage can be deployed in many ways.
- For example:
  - Local data (desktop/laptop) can be backed up to cloud storage.
  - A virtual disk can be ‘sync’ to the cloud and distributed.
  - The cloud can be used as a reservoir for storing data.

OVERVIEW OF CLOUD STORAGE

- A subset of cloud computing
- Standards and services pertaining to cloud storage have to be understood before its implementation.
- Resources that are exposed to clients are called as functional interfaces, that is, **data paths**.
- Resources maintained by the service providers are called as management interfaces, that is, **control paths**.
- A standard model is to be developed and proposed for both interfaces, that is, **consumers and providers**.
  - That standard should be mapped to various services rendered by the provider.
  - This standard should act as a base for cloud storage interfaces.

- Cloud storage became popular because of the following attributes available in cloud computing:
  - pay-as-you-use
  - elasticity
  - simplicity (management)
- It is important that any provider providing storage as a service should also provide these attributes to the consumer.

Additional Cloud Storage Attributes

- **Resource pooling and multi-tenancy**: Multiple consumers can use shared single storage device. Storage resources are pooled and consumers can be assigned and unassigned resources according to their needs.
- **Scalable and elastic**: Virtualized storage can be easily expanded on need basis.
- Accessible **standard protocols** including HTTP, FTP, XML, SOAP and REST.
- **Service-based**: Consumers no need to invest, that is, no CAPEX (Capital Expenditure) and only pay for usage, that is, **OPEX** (Operational Expenditure).
- **Pricing based on usage**
- **Shared and collaborative**
- **On-demand self-service**

- Cloud storage can be accessible through web-based applications or through web services Application Programming Interfaces (APIs), and using this data are stored.
- IT organizations have started developing personalized web applications for easy access of cloud storage services.

DATA MANAGEMENT FOR CLOUD STORAGE

- In the early stages, cloud storage focused on the **best effort service**.
- To support enterprise applications, **quality of service** has to be increased and extra services deployed.
- Cloud storage will lose its abstraction and its benefits such as simplicity, heterogeneity and good performance, if complex management services are added.
- Cloud storage should incorporate new services according to change of time.

- For cloud storage, a standard document is placed by SNIA Storage Industry Resource Domain Model (SIRDM).
- https://www.snia.org/sites/default/files/technical-work/sirdm/SIRDM_v2.0.pdf
- It states the importance of simplicity for cloud storage.
- Diagram : the SIRDM model which uses CDMI standards

Cloud Storage Usage of SIRDM Model

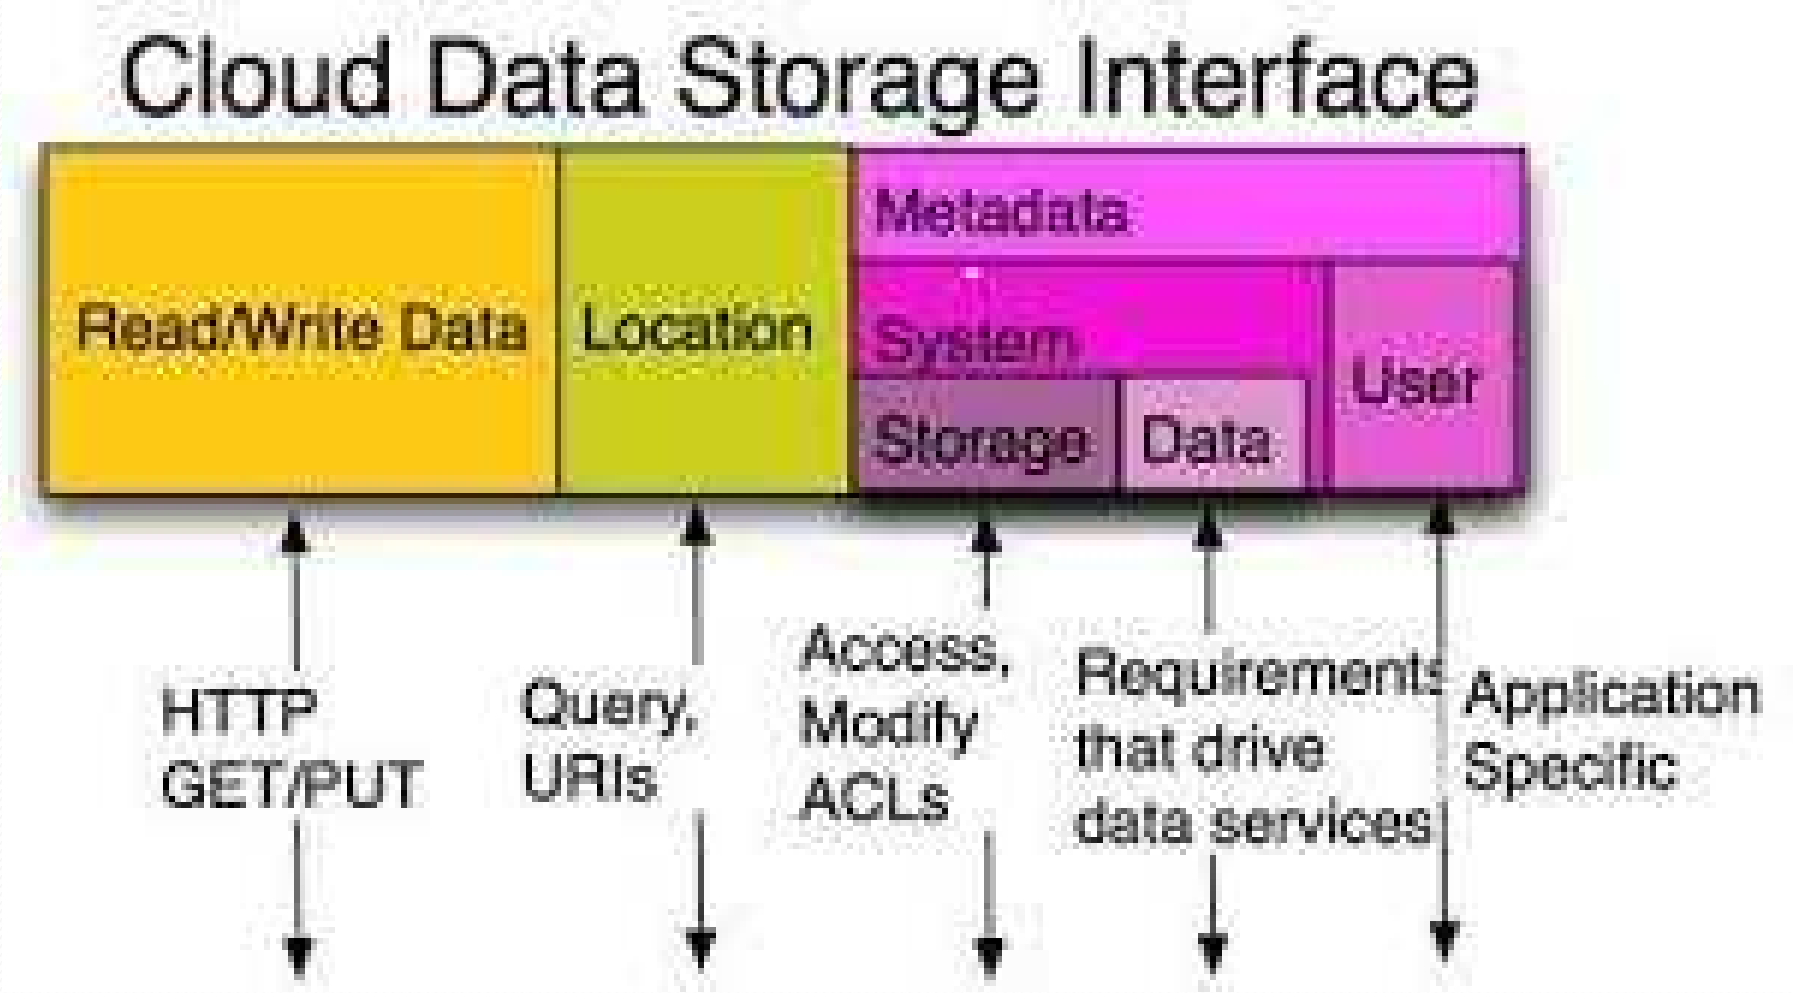

Cloud Storage Usage of SIRDM Model

- SIRDM model adopts three metadata: system consisting of storage metadata, data metadata and user metadata.
- By using these metadata, cloud storage interface can offer services without adding unnecessary complexity in managing the data.
- Storage system and data system metadata are used to meet the requirements of the data and the simplicity required is maintained.
- User metadata is used by the cloud to find the data objects and containers.
- Storage system metadata is used by the cloud to offer basic storage functions like assigning, modifying and access control.
- Data system metadata is used by the cloud to offer data as a service based on user requirements and controls the operation based on that data.

1.Cloud Data Management Interface (CDMI)

- To create, retrieve, update and delete objects in a cloud the cloud data management interface (CDMI) is used.
- The functions in CDMI are:
  - Cloud storage offerings are discovered by clients
  - Management of containers and the data
  - Sync metadata with containers and their objects
- CDMI is also used to manage containers, domains, security access and billing information.
- CDMI standard is also used as protocols for accessing storage.
- CDMI defines how to manage data and also ways of storing and retrieving it.
- 'Data path' means how data is stored and retrieved.
- 'Control path' means how data is managed.
- CDMI standard supports both data path and control path interface.

2. Cloud Storage Requirements

1. Multi-tenancy
2. Security
3. Secure Transmission Channel
4. Performance
5. Quality of Service (QoS)
6. Data Protection and Availability
7. Metering and Billing

2.1.Multi-tenancy

- In a multi-tenancy model,
  - resources provided are pooled
  - may be shared by multiple customers
  - based on their needs
- Due to the elasticity property in cloud computing, shared pool of storage model makes the provider cost effective and billing is made easy.

2.2.Security

- Secure cloud storage requires a secure transmission channel and methods.
- Securing data can be done using encryption, authentication and authorization.
  - *Encryption* is the process of scrambling data in such a manner as to make it unreadable without special information, called a key, to make it readable again.
  - *Authentication* is the process of determining their identity. Authentication can employ passwords, biometrics, identifying tokens and other means.
  - *Authorization* determines access rights on the data and the levels of authorization. To provide secure cloud storage, access must be restricted for the communication channel, the data source and the cloud storage sites.

2.3.Secure Transmission Channel

- The four primary methods used to secure network communications are:
  1. Transport Layer Security (TLS) and Secure Sockets Layer (SSL)
  2. Hypertext Transfer Protocol Secure (HTTPS)
  3. Private Networks
  4. Virtual Private Networks (VPNs)

2.4.Performance & 2.5.Quality of Service (QoS)

Performance

- Cloud storage performance can be categorized into two: speed and latency.
- Factors that affect cloud storage performance are:
  - available network bandwidth,
  - types of systems available in provider's end,
  - method adopted for compression and caching.

Quality of Service (QoS)

- Quality of service (QoS) refers to levels of performance and efficiency of the system that they can provide.

2.6.Data Protection and Availability

- To ensure that data is protected from loss and theft, providers must take some precautionary measures:
  - Physical site security
  - Protection against power loss
  - Protection against loss of network access
  - Data redundancy
  - Server redundancy and server fail-over
  - Redundant data sites
  - Levels of redundancy
  - Versioning and data retention
  - Accessibility of cloud storage as live data
  - Backup to tape or other media
  - Data availability, when contract disputes

2.7.Metering and Billing

- Metering and billing in cloud storage are done based on:
  - data uploaded,
  - data downloaded,
  - data stored
- depends on requests and types of request.

PROVISIONING CLOUD STORAGE

- **Cloud** means sharing third party resources via the Internet.
  - sharing be done on need basis and no need to invest any infrastructure at consumers end.
- **Storage clouds** increase the efficiency of storing data in remote places, by sharing the storage devices provided by the service providers.
  - Capacity of storage can be increased on need basis and can be done using multi-tenancy methods.
- **Private storage clouds** reside at the back of an organization's firewall that is deployed for in-house customers and is designed for providing elasticity and simplicity in cloud model.
- By adopting **Cloud Data Management Interface (CDMI)**, standard service providers can implement the method for metering the storage and data usage of consumers.
  - This interface also helps the providers for billing to the IT organizations based on their usage.
  - Advantage : IT organizations need not write/use different adapters used by the service providers.
  - By using this interface, they can connect with different service providers.

DATA-INTENSIVE TECHNOLOGIES FOR CLOUD COMPUTING

- Processing Approach
- System Architecture
  1. MapReduce
  2. HPCC

Compute & Data-intensive computing

- Parallel processing approaches are divided into two types:

1. **compute-intensive**

- Applications which need more execution time for computational requirements are termed as compute-intensive.
- Application programs that are compute bound described using compute-intensive requires more execution time.
- Parallel processing of this type of application involves running individual algorithms parallel within a process.

2. **data-intensive**

- Data-intensive computing is a lateral type of computing which use parallelism concept for processing large volumes of data, called big data.
- Data-intensive is a term used to describe applications which seeks large volume of data and time in process.

3.1.Processing Approach

- Data-intensive computing platforms use a parallel computing approach.
- This approach combines multiple processors and disks as computing clusters connected via high-speed network.
- The data needed are processed independently by computing resources available in the clusters.
- This improves the performance and scalability.
- A cluster can be defined as a parallel and distributed system, consisting of multiple inter-connected standalone computers working as a single computing resource.
- In parallel computing, this move is suitable for data-intensive computing.

Characteristics of data-intensive computing

- The principle mechanism used for collection of the data and programs or algorithms to perform the computation
- Programming model used
- Reliability and availability
- Scalability of both hardware and software

3.2.System Architecture

- For data-intensive computing an array of system architectures have been implemented.
- A number of solutions have come out, one among them is MapReduce concept which is developed by Google and available as open-source implementation known as Hadoop.
  - This project is used by Yahoo, Facebook and others.
- Apart from this, proprietary system architecture for data-intensive computing is developed by LexisNexis Risk Solutions called LexisNexis.

3.2.1. MapReduce

- The MapReduce architecture and programming model is an example for data-intensive computing, pioneered by Google.
- To create a map function, the MapReduce architecture uses a functional programming style using key-value pair.
- This pair is connected with the input data to produce a set of intermediary key-value pairs.
- Reduce function merges all intermediate values using intermediate keys.
- System takes care of particulars like partitioning the input data, scheduling and executing automatically.
- Hence programmers who do not have experience in parallel programming can simply use a large distributed processing environment without any problem.

3.2.2. HPCC

- LexisNexis Risk Solutions independently developed and implemented a solution for data-intensive computing called the HPCC (High-Performance Computing Cluster).
- The LexisNexis method structure clusters with commodity hardware that runs in Linux OS.
- Custom system software and middleware parts were created and layered to provide the execution environment and distributed file system support that is essential for data-intensive computing on the base of Linux operating system.
- A new high-level language for data-intensive computing called ECL [**E**nterprise **C**ontrol **L**anguage] is also implemented by LexisNexis.

POINTS COVERED
PART II CLOUD STORAGE

- ✓ Data Management
- ✓ Provisioning Cloud storage
- ✓ Data Intensive Technologies for Cloud Computing.

POINTS TO COVER
PART III CLOUD STORAGE FROM LANs TO WANs

- Cloud Characteristics
- Distributed Data Storage

- **Exemplar/Case Studies** Online Book Marketing Service, Online Photo Editing Service

Reading Material :
Srinivasan, J. Suresh, "Cloud Computing: A Practical Approach for Learning and Implementation",
Chap 13, pg 159 – 165

CLOUD CHARACTERISTICS

- There are three characteristics of a cloud computing natural environment that are most pertinent to be considered before choosing storage in cloud.

1. Compute Power
2. Data
3. Data Duplication

1. Computer power

- Computer power is elastic, when it can perform parallel operations.
- In general, applications conceived to run on the peak of a shared-nothing architecture are well matched for such an environment.
- Some cloud computing goods, for example, Google's App Engine, supply not only a cloud computing infrastructure, but also an entire programs stack with a constrained API so that software developers are compelled to compose programs that can run in a shared-nothing natural environment and therefore help elastic scaling.

2. Data

- Data is retained at an unknown host server.
- In general, letting go off data is a threat to many security issues and thus suitable precautions should be taken.
- The very title ‘loud computing’ implies that the computing and storage resources are being operated from a celestial position.
- The idea is that the data is physically stored in a specific host country and is subject to localized laws and regulations.
- Since most cloud computing vendors give their customers little command over where data is stored, the customers has no alternative but to expect the least that the data is encrypted utilizing a key unavailable with the owner, the data may be accessed by a third party without the customer’s knowledge.

3. Data Duplication

- Data is duplicated often over distant locations.
- Data accessibility and durability is paramount for cloud storage providers, as data tampering can be impairing for both the business and the organization's reputation.
- Data accessibility and durability are normally accomplished through hidden replications.
- Large cloud computing providers with data hubs dispersed all through the world have the proficiency to provide high levels of expected error resistance by duplicating data at distant locations across continents.
- Amazon's S3 cloud storage service replicates data over 'regions' and 'availability zones' so that data and applications can survive even when the whole location collapses.

POINTS COVERED
PART III CLOUD STORAGE FROM LANs TO WANs

- ✓ Cloud Characteristics
- Distributed Data Storage

- **Exemplar/Case Studies** Online Book Marketing Service, Online Photo Editing Service

Distributed storage

- Are evolving from the existing practices of data storage for the new generation of WWW applications through organizations like Google, Amazon and Yahoo
- Reasons for distributed storage means to be favored over traditional relational database systems encompassing **scalability, accessibility and performance**.
- The new generation of applications require processing of data to a tune of terabytes and even petabytes.
- This is accomplished by distributed services.
- Distributed services means distributed data.
- This is a distinct massive compared to traditional relational database systems.
- Several studies have proposed that this is an end of an architectural era and relational database systems have to take over.
- Emerging answers are Amazon Dynamo, CouchDB and ThruDB.

DISTRIBUTED DATA STORAGE

- Amazon Dynamo
- CouchDB
- ThruDB

Amazon Dynamo

- Amazon Dynamo is a widely used key-value store.
- It is one of the main components of Amazon. com, the biggest e-commerce stores in the world.
- It has a primary-key only interface.
- This demands that data is retained as key-value in twos, and the only interface to get access to data is by identifying the key.
- Values are anticipated to be barely there (less than 1 MB).
- Dynamo is said to be highly accessible for composing as opposed to reading, since malfunction of composing inconveniences the end-user of the application.
- Therefore any data confrontations are finalized at the time of reading than writing.

Additional Content to Refer

DynamoDB

amazon
web services™

Dynamo

- Another distributed data system
- Developed at Amazon
- Underlies its SimpleDB key-value pair database
- Designed specifically for supporting a large volume of concurrent updates, each of which could be small in size

Dynamo

- Dynamo's data model is that of simple key-value pairs
- It is expected that applications read and write such data objects fairly randomly.
- This model is well suited for many web-based e-commerce applications that all need to support constructs such as a 'shopping cart.'

Dynamo

- Also replicates data for fault tolerance
- Makes use of distributed object versioning and quorum-consistency
  - To enable writes to succeed without waiting for all replicas to be successfully updated
- Managing conflicts if they arise is relegated to reads which are provided enough information to enable application dependent resolution.

- *Because of these features, Dynamo does not rely on any underlying distributed file system and instead directly manages data storage across distributed nodes.*

The architecture of Dynamo

- Objects are key-value pairs with arbitrary arrays of bytes.
- An MD5 hash of the key is used to generate a 128-bit hash value.
- The range of this hash function is mapped to a set of virtual nodes arranged in a ring, so each key gets mapped to one virtual node.
- The object is replicated at this primary virtual node as well as $N - 1$ additional virtual nodes (where N is fixed for a particular Dynamo cluster).

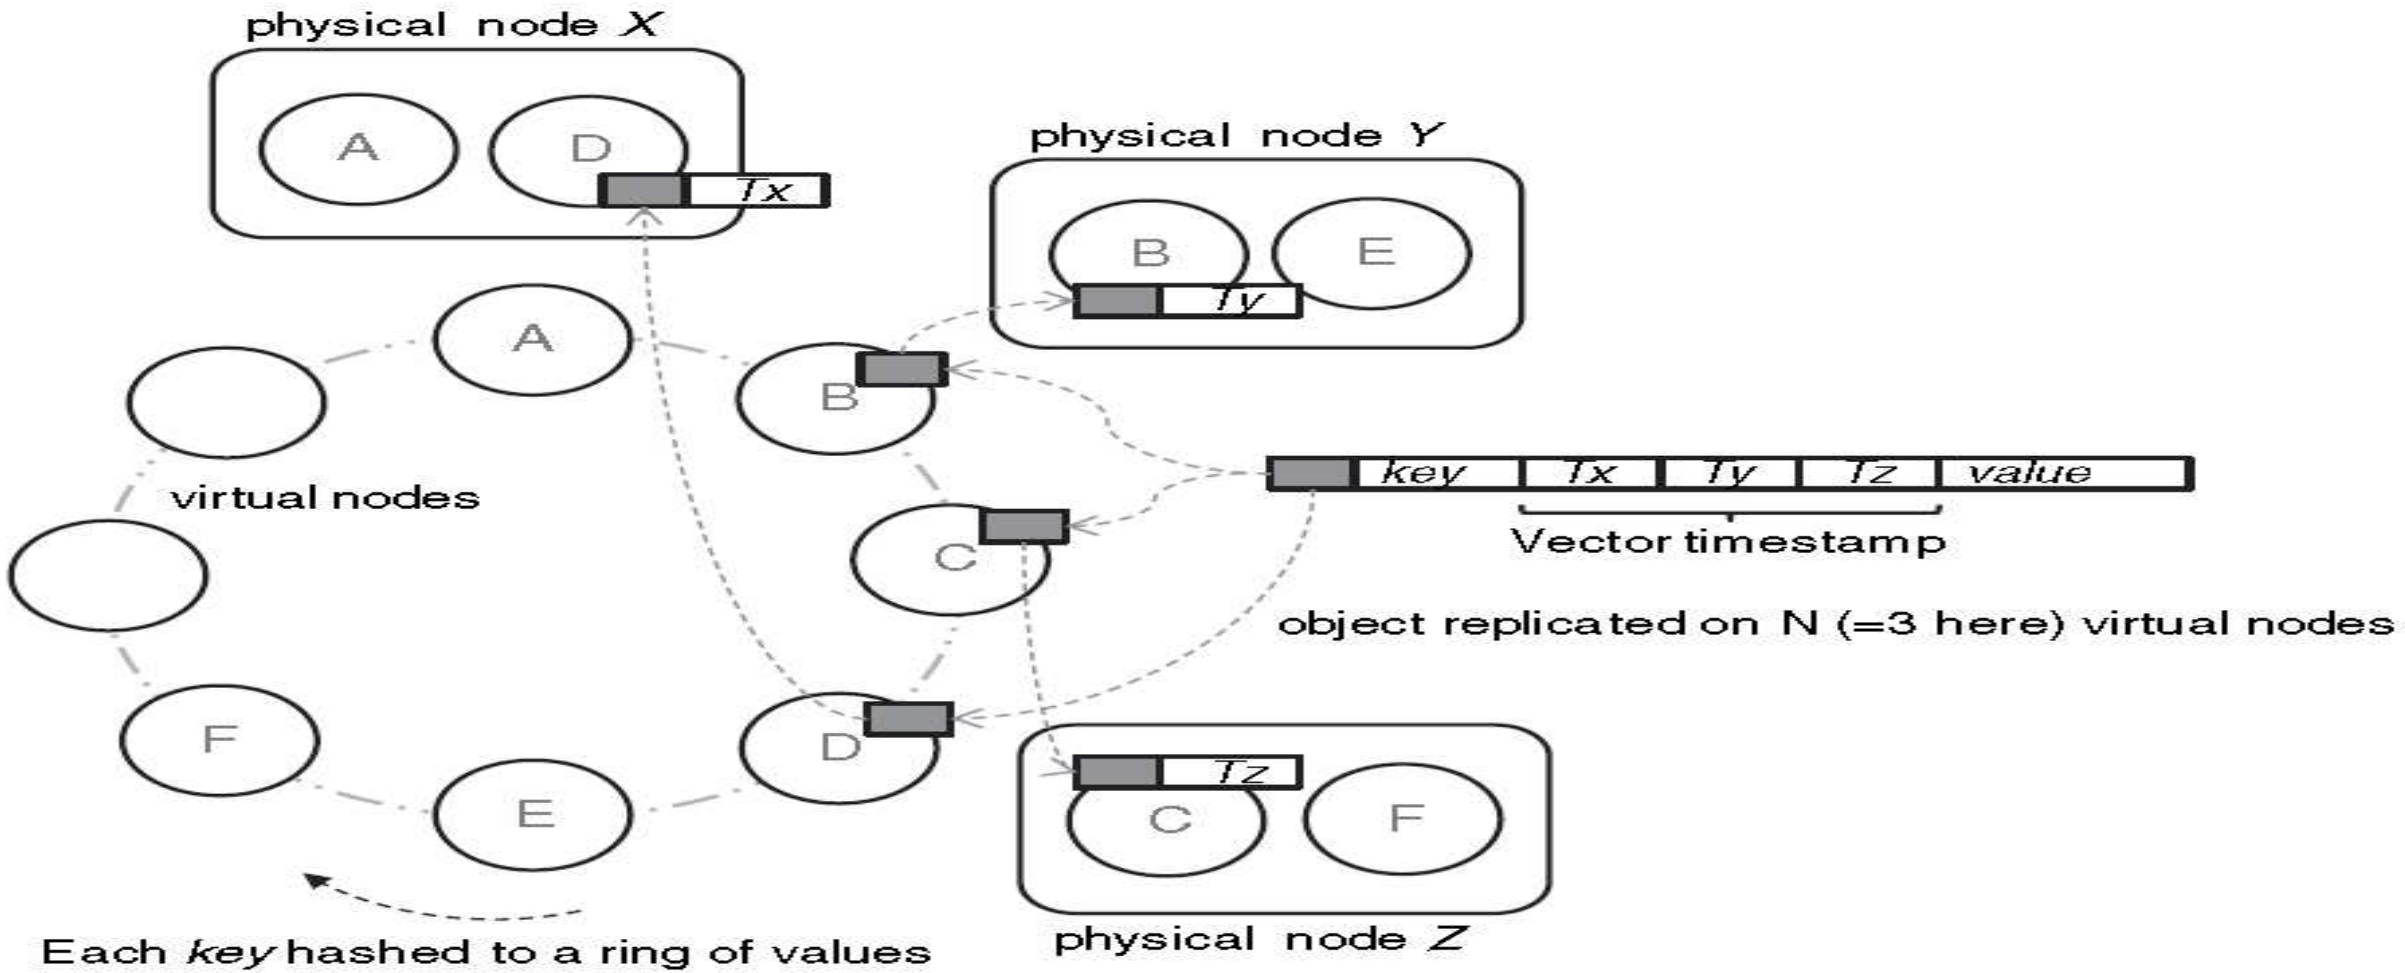

FIGURE 10.6. Amazon Dynamo

The architecture of Dynamo

- Each physical node (server) handles a number of virtual nodes at distributed positions on the ring
  - To continuously distribute load evenly as nodes leave and join the cluster because of transient failures or network partitions.
- The Dynamo architecture is completely symmetric with each node being equal

Write….

- A write request on an object is first executed at one of its virtual nodes which then forwards the request to all nodes having replicas of the object.
- Objects are always versioned, so a write merely creates a new version of the object with its local timestamp (*Tx* on node *X*) incremented.
- Thus the timestamps capture the history of object updates
- versions that are superseded by later versions having a larger *vector* timestamp are discarded.

- In Dynamo write operations are allowed to return even if all replicas are not updated.
- However a quorum protocol is used to maintain *eventual* consistency of the replicas when a large number of concurrent reads and writes take place

- Each read operation accesses R replicas and each write ensures propagation to W replicas;

  as long as

$R + W > N$ the system is said to be quorum consistent

- Thus, if we want very efficient writes, we pay the price of having to read many replicas, and vice versa.

- In practice Amazon uses $N = 3$, with R and W being configurable depending on what is desired
- for a high update frequency one uses $W = 1$, $R = 3$,
- for a high-performance read store $W = 3$, $R = 1$ is used.

Conclusion

- Dynamo is able to handle transient failures by passing writes intended for a failed node to another node temporarily.
- Such replicas are kept separately and scanned periodically with replicas being sent back to their intended node as soon as it is found to have revived.
- Finally, Dynamo can be implemented using different storage engines at the node level, such as Berkeley DB or even MySQL
- Amazon is said to use the Berkeley DB in production

AWS….

- Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability.
- DynamoDB lets you offload the administrative burdens of operating and scaling a distributed database so that you don't have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling.
- DynamoDB also offers encryption at rest, which eliminates the operational burden and complexity involved in protecting sensitive data.

AWS DynamoDB

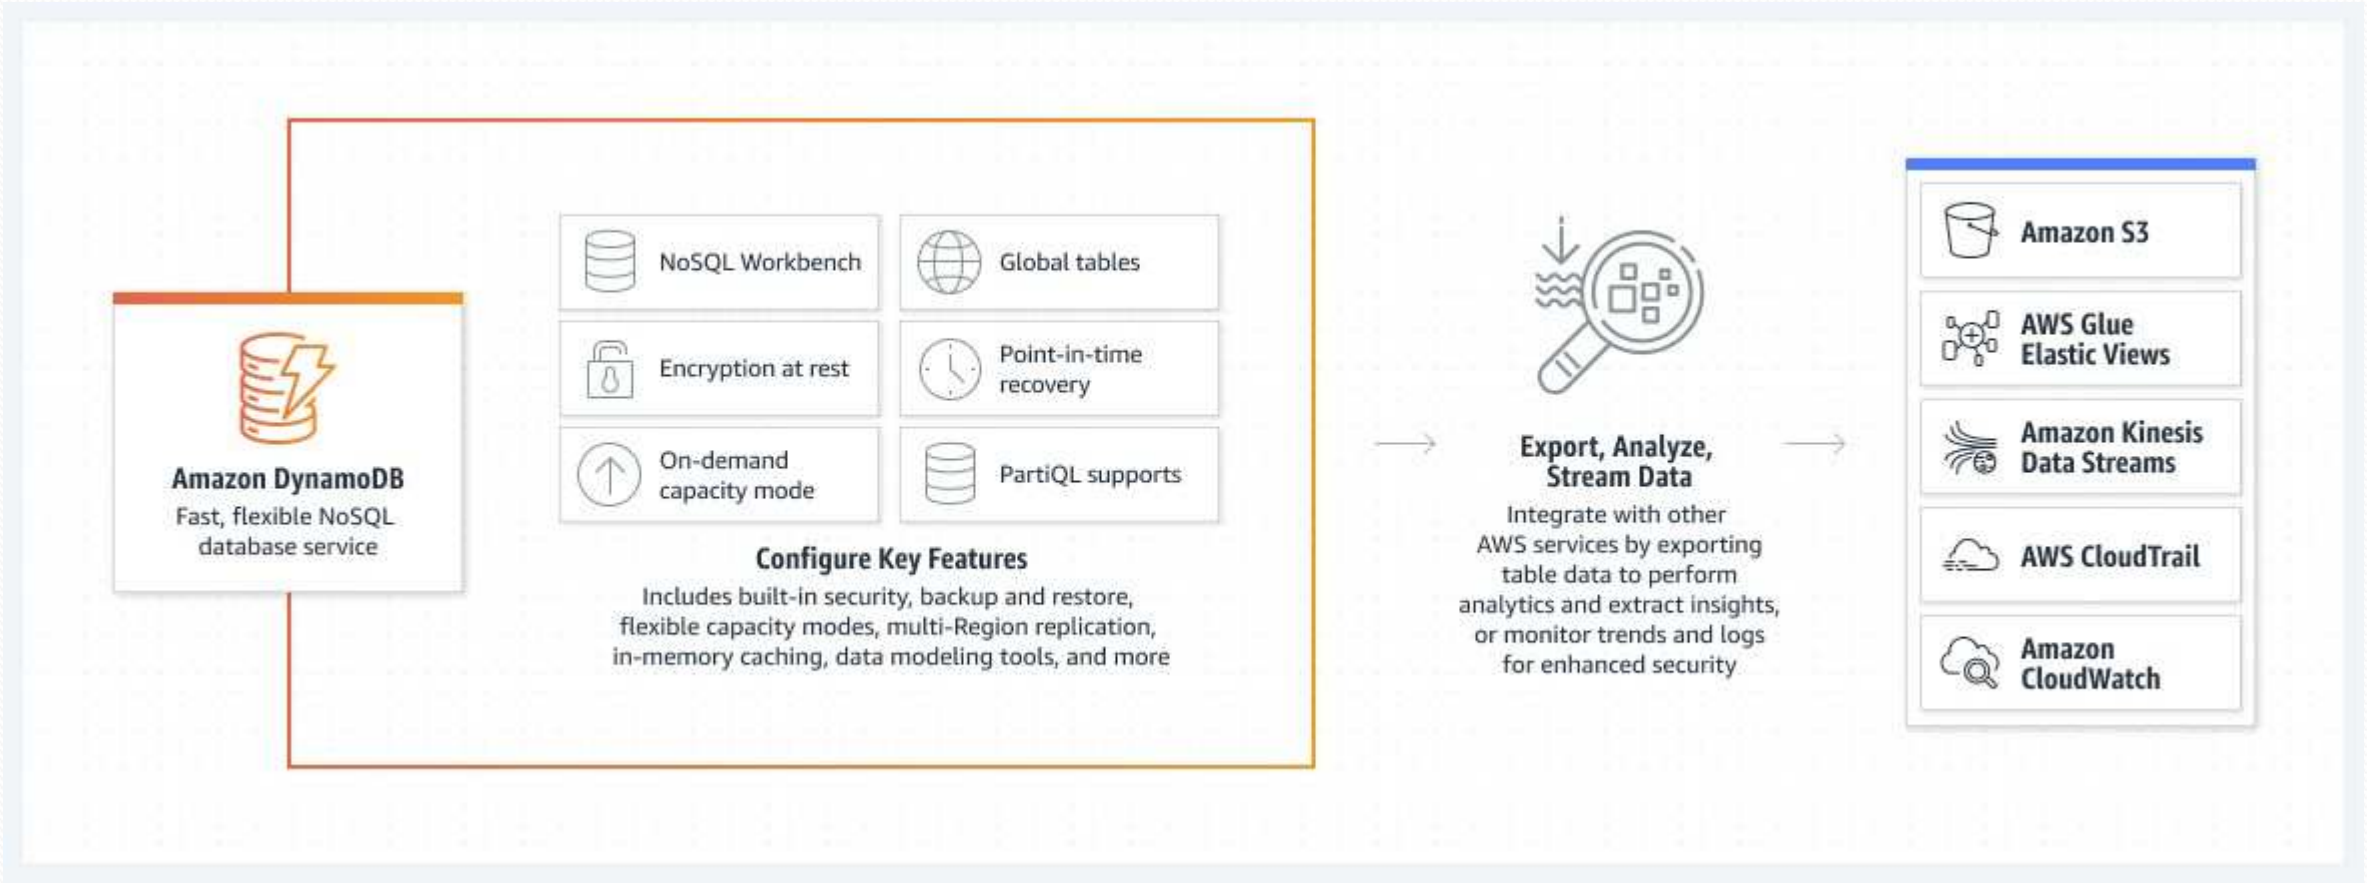

AWS reference ends…..

2. CouchDB

- A document-oriented database server, accessible by REST APIs
- An acronym for ‘Cluster Of Unreliable Commodity Hardware’, emphasizing the distributed environment of the database
- Designed for document-oriented applications, for example, forums, bug following, wiki, Internet note, etc.
- CouchDB is ad-hoc and schema-free with a flat address space

- CouchDB aspires to induce the 4 Pillars of Data Management by these methods:
  1. Save: ACID compliant, save efficiently
  2. See: Easy retrieval, straightforward describing procedures, fulltext search
  3. Secure: Strong compartmentalization, ACL, connections over SSL
  4. Share: Distributed means

- The storage form is a Multiversion Concurrency Control (MVCC) scheme with hopeful locking.
- A purchaser sees a snapshot of the data and works with it even if it is altered at the same time by a distinct client.

- CouchDB actually has no apparent authentication scheme, i.e., it is in-built.
- The replication is distributed.
- A server can revise others once the server is made offline and data is changed.
- If there are conflicts, CouchDB will choose a survivor and hold that as latest.
- Users can manually suspend this surviving alternative later.
- Importantly, the conflict tenacity yields identical results comprehensively double-checking on the offline revisions.
- This also promises to compose a storage motor for MySQL founded on CouchDB.

Wikipedia…

- Unlike a relational database, a CouchDB database does not store data and relationships in tables.
- Instead, each database is a collection of independent documents.
- Each document maintains its own data and self-contained schema.
- An application may access multiple databases, such as one stored on a user's mobile phone and another on a server.
- Document metadata contains revision information, making it possible to merge any differences that may have occurred while the databases were disconnected.

- CouchDB implements a form of multiversion concurrency control (MVCC) so it does not lock the database file during writes.
- Conflicts are left to the application to resolve.
- Resolving a conflict generally involves first merging data into one of the documents, then deleting the stale one.

3.ThruDB

- ThruDB aspires to be universal in simplifying the administration of the up-to-date WWW data level (indexing, caching, replication, backup) by supplying a reliable set of services:
  - Thrucene for indexing
  - Throxy for partitioning and burden balancing
  - Thrudoc for article storage

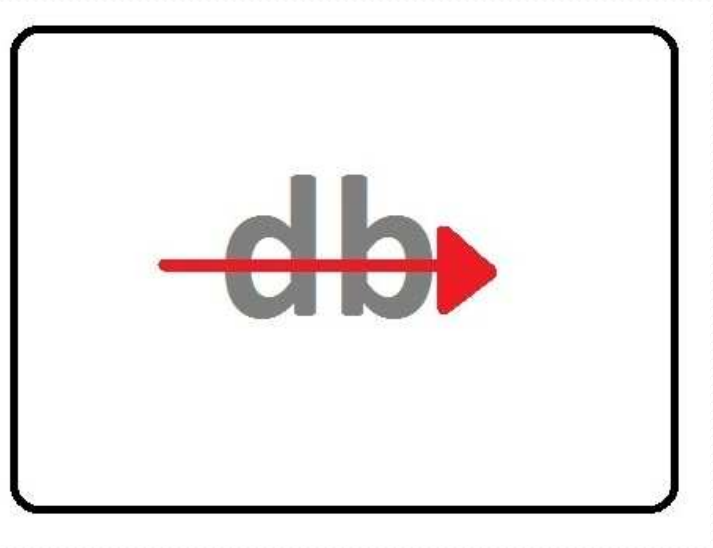

- ThruDB builds on top of some open source projects: Thrift, Lucene (indexing), Spread (message bus), Memcached (caching), Brackup (backup to disk/S3) and also values Amazon S3.

http://www.dbtalks.com/article/introduction-to-thrudb/

- ThruDB is an open source database built on Apache's Thrift framework and is a set of simple services such as scaling, indexing and storage which is used for building and scaling websites.
- It provides flexible, fast and easy-to-use services that simplify the management of the modern web data layer and provides developers with features and tools most web developers need.
- These features can be easily configured or turned off.

Thrudb contains two services

- *Thrudoc* - Document storage service
- *Thrudex* - Indexing and search service

Features

- Multi-master replication
- Built for horizontal scalability
- Incremental backups and redo logging
- Multiple storage back-end client libraries for most languages
- Simple and powerful search API

Services

ThruDB provides web-scale data management by providing these services:

- *Thrucene* - For Lucene-based indexing
- *Throxy* - For partitioning and load balancing
- *Thrudoc* - For document storage
- *Thruqueue* - For a persistent message queue service
- *Thrift* - For cross-language services framework

- Compared to SimpleDB, ThruDB removes many of the former's restrictions- no 1024 byte limit per attribute, ability to use proprietary data formats, no query time limits, and more. Since it also runs on EC2 and public Thrift/ThruDB *Amazon Machine Images* (AMIs) exist, ThruDB is in a sense a competitor to SimpleDB.
- Thrudb is an attempt to simplify the process of the modern web data layer which comprises indexing, caching, replication & backup. It provides a consistent set of services. Thrucene provides indexing, Throxy provides partitioning and load balancing, and Thrudoc provides document storage.
- The entire stack is implemented on top of Thrift, which has a great access for easy language interoperability and API access.

More Systems…

- ***Amazon Simple Storage Service*** is a straightforward data storage scheme with a hash- table like API. It is a hosted service with interior architecture minutia not available. It is proclaimed that the conceive obligations of S3 are scalable, reliable, fast, inexpensive and simple.
- ***Amazon SimpleDB*** is a hosted WWW service for running queries on organized data in real time. It has the prime functionality of a database, real-time lookup and straightforward querying of organized data.
- ***MemcacheDB*** is a distributed key-value storage scheme conceived for persistence. It conforms to the memcache protocol. Memcachedb values Berkeley DB as a saving backend, so allotments of characteristics encompassing transaction and replication are supported.
- Distributed storage has numerous anxieties, scalability, hardware obligations, query form, malfunction management, data consistency, durability, reliability, effectiveness, etc.
- The future of data storage can be viewed in addition, for example, CouchDB Integration with Abdera, an Atom store.

Case Studies

- Online Book Marketing Service : Amazon KDP

https://www.junglescout.com/blog/amazon-kindle-direct-publishing/#:~:text=Kindle%20Direct%20Publishing%2C%20or%20KDP,upfront%20costs%20or%20inventory%20orders.

- Online Photo Editing Service

Online Photo Editing Service

Online Photo Editors

- While the Cloud and SaaS signal are growing, a numerous user-friendly Photo Editors accessible on a SaaS form absolutely free.
- Generally, they are neither as unique nor very fast as Photoshop, but they are developing, and in most cases they have the usual characteristics and much more.
- Another good thing about Online Photo Editors is that they can be accessed from any location and any computer with Internet connection.

Photoshop Express Editor

- It has been constructed on the convention of Photoshop minus the technicalities.
- It is ideal for unskillfulphotographers who don’t wish to get involved into the complicated features of the Photoshop.
- In spite of all this, it has its own limitations.
- The publishing choices are absent.
- Also, it does not support photographs from high mega pixel cameras.

Picnik

- It has the responsibility of extraordinary consequences, a variety of fascinating fonts and shapes.
- It enables red eye decrease and also edits the exposure which is most challenging for photographers.
- It is very speedy and works well on distinct platforms such as Mac, Windows and Linux.
- One of the most compelling characteristics of Picnik is the support of the photograph distributing sites and the community networking sites.

Splashup

- A free online tool for editing photos.
- It is browser friendly and carries various photograph sharing services as Picasa, Flickr and Facebook.
- Splashup comprises of numerous photograph revising tools such as lasso, distort, brush load up, crop, etc.
- Multiple windows are permitted, which entrusts its demonstration compared to the other tools.
- Besides, it also presents us a Splashup Light, a free offline photograph reviewer, which works flawlessly on our desktop as well as on wireless PCs.

FotoFlexer

- It is one of the best choices for photograph editing.
- It has all the rudimentary characteristics and supplements numerous sophisticated tools.
- It offers upfront animations that most of the online devices don’t offer.
- sOther characteristic that makes Fotoflexer to stand out is that it has 25 filters and it can make flat images.

Pixer.us

- Pixer.us is a straightforward and direct device for revising quickly.
- It does not need signup. It has all the rudimentary devices such as crop, rotate, flip, resize with hue rectify, and resize.
- A positive feature is that it permits unlimited undo options.

Photoshop Express Editor

- https://www.adobe.com/in/products/photoshop-express.html

Course Outcome:

CO2:

On completion of the unit, learners should be able to Use appropriate data storage technique on Cloud, based on Cloud application