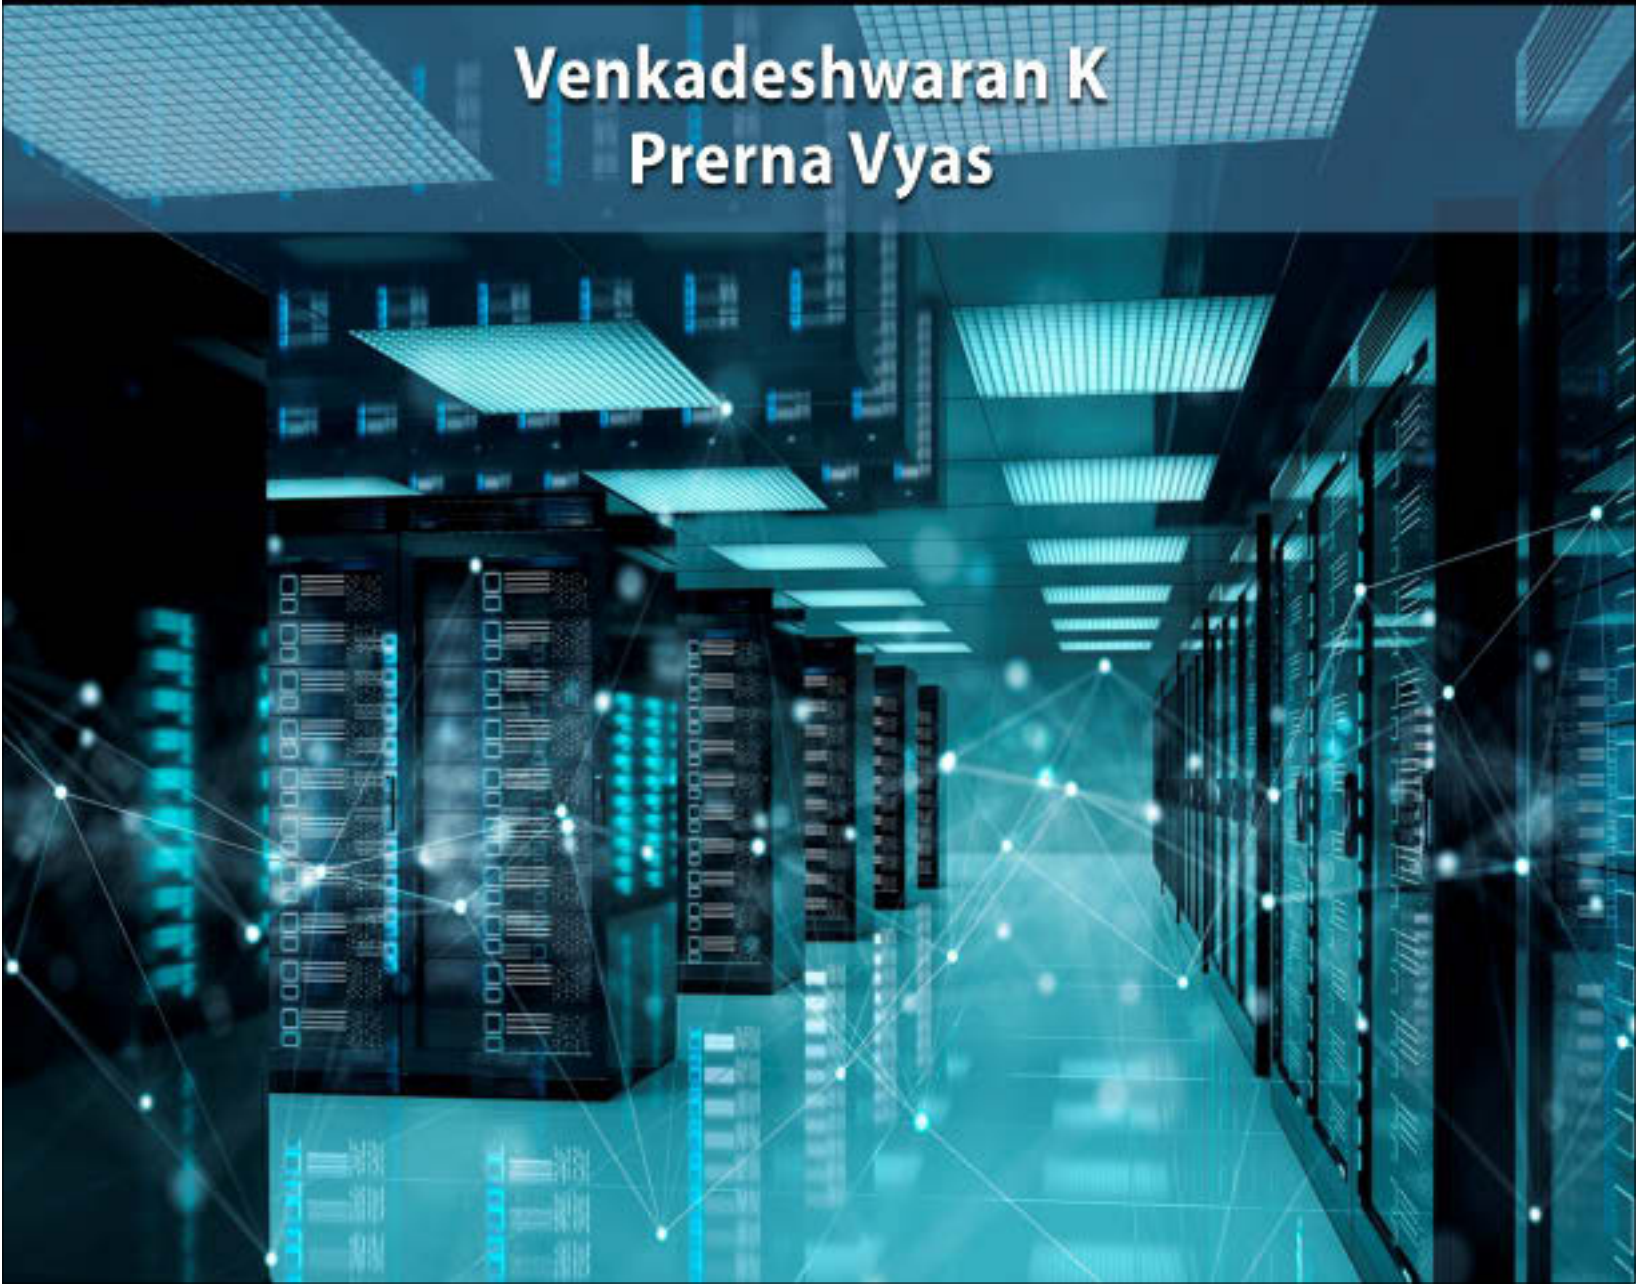


INFORMATION STORAGE MANAGEMENT

Venkadeshwaran K
Prerna Vyas



Information Storage Management

Information Storage Management

Venkadeshwaran K

Prerna Vyas



BOOKS ARCADE

KRISHNA NAGAR, DELHI

Information Storage Management

Venkadeshwaran K
Prerna Vyas

© RESERVED

This book contains information obtained from highly regarded resources. Copyright for individual articles remains with the authors as indicated. A wide variety of references are listed. Reasonable efforts have been made to publish reliable data and information, but the author and the publisher cannot assume responsibility for the validity of all materials or for the consequences of their use.

No part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereinafter invented, including photocopying, microfilming and recording, or any information storage or retrieval system, without permission from the publishers.

For permission to photocopy or use material electronically from this work please access booksarcade.co.in

BOOKS ARCADE

Regd. Office:

F-10/24, East Krishna Nagar, Near Vijay Chowk, Delhi-110051

Ph. No: +91-11-79669196, +91-9899073222

E-mail: info@booksarcade.co.in, booksarcade.pub@gmail.com

Website: www.booksarcade.co.in

Year of Publication 2023

International Standard Book Number-13: 978-93-90762-74-3



CONTENTS

Chapter 1. Introduction to Information Storage and Management	1
— <i>Venkadeshwaran K</i>	
Chapter 2. Storage System Environment.....	12
— <i>Pradeepa P</i>	
Chapter 3. Redundant Array of Independent Disks.....	25
— <i>Dimple Bahri</i>	
Chapter 4. Intelligent Storage System	36
— <i>Dr. Uthama Kumar A</i>	
Chapter 5. Storage Networking Technologies	48
— <i>Dr. Thirukumaran Subbiramani</i>	
Chapter 6. Storage Area Network.....	58
— <i>Dr. Santosh S Chowhan</i>	
Chapter 7. Network-Attached Storage.....	75
— <i>Purna Vyas</i>	
Chapter 8. Advanced Storage Networking	88
— <i>Purna Vyas, Assistant Professor</i>	
Chapter 9. Storage Virtualization.....	99
— <i>Amita Kashyap</i>	
Chapter 10. Business Continuity	113
— <i>Amita Kashyap</i>	
Chapter 11. Backup and Recovery.....	124
— <i>Brijraj Singh Solanki</i>	
Chapter 12. Local Replication.....	141
— <i>Brijraj Singh Solanki</i>	

CHAPTER 1

INTRODUCTION TO INFORMATION STORAGE AND MANAGEMENT

Venkadeshwaran K, Associate Professor

Department of Mechanical Engineering, Faculty of Engineering and Technology, JAIN (Deemed-to-be University),
Karnataka – 562112

Email Id- k.venkadeswaran@jainuniversity.ac.in

Information technology's main tenet is information storage. Every second, individuals and companies produce a considerable amount of digital information. It is necessary to manage, optimize, and store this data in traditional, virtualized, and dynamically changing cloud infrastructures. Data storage was thought to consist solely of a collection of discs or tapes hooked to the computer's back.

Information storage technology is still crucial to the availability, performance, integration, and optimization of the overall IT infrastructure, but few people outside of the storage business are aware of this. Information storage has advanced significantly over the past ten years, offering a wide range of options for maintaining, linking, safeguarding, sharing, and maximizing digital information. Modern storage technologies are becoming increasingly crucial and pertinent for the success of businesses and other organizations due to the widespread acceptance of virtualization, the rise of cloud computing, the multifold increase in data volume year over year, and a variety of data types and sources. More than ever, IT administrators face difficulties finding and training highly qualified technical personnel with experience in storage systems across traditional, virtualized, and cloud environments.

Many Types of data storage

- A. **Cloud storage:** By storing data in the cloud, businesses can make it more easily accessible to authorized users through the internet. Google, Microsoft, and other well-known cloud storage providers are just a few examples.
- B. **Software-defined storage:** This method of managing data through abstraction uses software. It functions by removing data from the actual storage that has been set up for network use. Additionally, it functions nicely with micro services and containers that use unstructured data.
- C. **File storage:** One of the most popular methods of data storage employed by corporations is file storage. It maintains a single piece of data in a hierarchical structure. This makes it easier for users to access data using special identifiers or pathways like names, places, and URLs.
- D. **Block storage:** Storage is divided into separate blocks using block storage. Each block has a distinct identity that guarantees data security and allows for the free placement of small amounts of information for easier retrieval. Blocks are quicker and are perfect for rich media databases. Additionally, they can give consumers total configuration freedom.
- E. **Object storage:** Data or files are divided up into units of information known as objects for object storage. Each object has an own identity and is a self-contained repository. Users may find and access information even on distributed systems because to this.

The Functions of Data Storage Management:

Performance and dependability: Managing data so that it is easily accessible for business activities is the goal of data storage management. Employee performance, effectiveness, and productivity rise with fast and easy access to data, which also enhances the user experience. Teams may employ media and automated tiering to optimize various storage tiers to speed up the process.

Data security and protection: It's critical to comprehend the significance of data protection while utilizing cloud storage. Use data backup services, encryption for both saved and in-transit data, multivariable identification to prevent unwanted access, and other security measures to safeguard business-critical data.

Control and compliance: To store the most important data assets, use different tiering levels or automated tiering. This aids in data management and archiving as well as helping businesses show regulatory compliance.

Data center infrastructure

Data center infrastructure refers to the essential physical or hardware-based resources and elements that make up a data center, including all IT infrastructure tools, gadgets, and technologies. It is modelled and named in a design plan that contains an exhaustive description of all the infrastructural parts required to build a data center.

Basic Components of a Data Center

The majority of data centers are located in buildings or other types of physical structures. However, not all of them are like this, since others are situated in underwater caverns or subterranean bunkers. Most data centers are located in office buildings or other structures of a similar kind. Although it sometimes has an aperture to bring in adequate air to assist prevent IT and computer equipment from overheating, the actual structure may not always have windows. HVAC equipment is often found on roofs, and data centers may use solar or wind power to generate energy. Major natural disasters like floods, hurricanes, earthquakes, tornadoes, typhoons, blizzards, heat waves, etc. may be physically handled by many data centers. For this reason, a lot of data center buildings are made of sturdy materials like steel-reinforced concrete. One story, or more often, numerous stories, inside the building, might house a data center. Some buildings have raised floor architecture, which creates a space between the actual floor and another floor that houses additional IT equipment. This space is then equipped with the electrical wiring, cabling, cooling equipment, and other resources necessary to support the ongoing operation of the primary computing equipment. Greater accessibility to connection and cooling infrastructure, as well as more efficient use of vertical space, are advantages of raised floor design. Typically, access to the floors is controlled by closed doors, which are monitored by security staff and security infrastructure (such as cameras and alarms). A data center may also contain separate, lockable server rooms, depending on the sort of building in issue.

Data Center Equipment: The server, which acts as the data center's central processing infrastructure, is the primary piece of computing hardware often found in data centers. Servers are often kept in cabinets and racks. Servers are linked together via cabling as well as to a larger network, such the internet. For a similar function, routers are found in plenty of racks. For the best connection for servers, the majority of data centers include a strong infrastructure for connectivity, including easy access to fiber optic cable.

Energy for a Data Center: Electrical infrastructure is also included in data centers to power computers and other equipment. Most data centers include backup electrical infrastructure in addition to having plenty of electrical sockets for servers and cable to link the data center to the larger municipal electrical grid. Data center facilities often include backup generators, solar panels, and wind turbines (as well as fuel for the generators).

Cooling a Data Center: The majority of data centers include some kind of cooling infrastructure to maintain servers and other components of computing infrastructure at an ideal temperature and prevent them from overheating (plus the infrastructure needed to keep the cooling infrastructure functioning and running smoothly). Fans, HVAC systems, air conditioners, and pipes carrying cold water from the outside that are run alongside heated infrastructure to cool it down are all examples of this.

Key Obstacles to Information Management

Businesses must take into account the following major information management difficulties in order to create an effective information management policy:

- A. The digital cosmos is expanding. Information is expanding at an exponential rate. The growth of information has multiplied due to repurposing and data duplication to ensure high availability.
- B. A growing reliance on information a key factor in determining a company's success is how strategically it uses information to get an edge over rivals.
- C. Information's value changes with time, thus what is significant today might not be as essential tomorrow. Information's worth fluctuates a lot throughout time.

Benefits and difficulties of Managing Data Storage

The management of data storage offers benefits and drawbacks. Positively, it enhances performance and guards against data loss. Storage systems function well over time, space, and users with good management. Additionally, it makes sure that data is protected from external dangers, human mistake, and system flaws. This data protection strategy includes elements such as appropriate backup and disaster recovery. Users receive the appropriate amount of storage capacity thanks to an efficient management method. Storage space can be increased or decreased as needed by organizations. The storage plan takes into account the applications' and needs' continual change. By centralizing administration, storage management also makes it simpler for administrators to manage a range of storage systems.

Challenges:

Persistent cyber threats, data management laws, and a distributed workforce are all difficulties in managing data storage. These difficulties highlight the need of putting in place a thorough plan: Lack of compliance could result in significant fines, and remote workers need to be confident they'll have access to information and applications just as they would if they were in a traditional office setting. A storage management strategy should ensure organizations protect their data against data breaches, ransomware, and other malware attacks. Complex and distributed systems are a challenge for managing data storage. Workers are dispersed, and systems operate both locally and on the cloud. HDDs, SSDs, and tapes might be used in an on-site storage environment. Organizations frequently employ several clouds. New technology, like AI, can be advantageous to

businesses but also increase complexity. These advantages result in lower expenses as well because administrators may make better use of storage resources.

The Value of Information Management and Storage:

Data storage and management has recently been the hot trend that all businesses must pursue. Data analytics is challenging, but managing and storing information is even more challenging. However, the management, storage, and analyses of data have all been made simpler by the AccelOps platform. Information management and storage are critical components of any organization. If the business will never use all that data, there is no point in gathering it. So, effectively handling information to optimize its use within the company is what data management actually is. With the AccelOps platform, information storage and management are carried out automatically, freeing up a significant amount of time for the IT team to perform other tasks. Even though the terms "data" and "information" are sometimes used interchangeably, "information" refers to data that has already undergone modifications to make it usable and comprehensible by people. Since this information will be utilized to evaluate a company's previous performance and forecast future events and happenings, information management is a crucial component of an organization's operations. Information is also utilized to communicate with clients, particularly if a business solicits their opinions frequently.

Types of Data

Depending on how it is handled and kept, data may be classed as structured or unstructured. Structured data is arranged in rows and columns according to a strict format specification so that programs may quickly obtain and handle it. Typically, structured data is kept in a database management system (DBMS).

Data is unstructured and hence challenging for business applications to query and retrieve if its components cannot be represented in rows and columns. Customer contacts, for instance, may be kept on sticky notes, in emails, on business cards. It is challenging to extract using a customer relationship management tool because of its unstructured nature. Unstructured data could lack the elements necessary to uniquely identify itself for any kind of processing or interpretation. Since over 80% of company data is unstructured and takes a substantial amount of storage space and work to maintain, businesses are mainly concerned with handling this data.

Information

Data, whether it is organized or unstructured, is useless to people or organizations if it is not represented in a meaningful way. Data must be analyzed by businesses in order to be useful. The wisdom and understanding obtained from data is known as information. Businesses examine unprocessed information in order to spot important patterns. A corporation may plan or adjust its strategy based on these patterns. For instance, a merchant may identify the brands and items that its consumers favor by examining their purchasing behavior and keeping track of the things they often buy.

By using the data in innovative ways, effective data analysis not only helps already-existing firms but also has the ability to open up brand-new business prospects. An example is a job portal. Job searchers upload their resumes on several websites that provide job search tools in order to access a larger pool of potential employers. These websites gather resumes and put them in easily found places for potential employers. Companies also advertise open opportunities on job-search

websites. The keywords on resumes are matched with the keywords in job advertisements using job matching software. In this way, the job search engine takes data and transforms it into knowledge for both job seekers and companies.

Information security and availability are constant concerns since they are essential to a company's success. These worries are further exacerbated by the legal, regulatory, and contractual duties relating to data availability and protection. Millions of dollars per hour are lost due to outages in vital sectors including electricity, telecommunications, manufacturing, and financial services.

Storage

Data produced by people or organizations must be archived and made readily available for processing in the future. Devices created to store data are referred to as storage systems or simply storage in a computer environment. Depending on the kind of information and the pace at which it is generated and utilized, several types of storage are used. Storage devices include things like hard drives in personal computers, memory in mobile phones and digital cameras, DVDs, and CD-ROMs. Data storage solutions for businesses include internal hard drives, external disc arrays, and cassettes.

Architecture and Storage Technology Evolution

In their data centers, corporations possessed centralized mainframe computers as well as tape reels and disc packs for storing information. It is now feasible for business units and departments to have their own servers and storage thanks to the development of open systems and the accessibility and simplicity of deployment they provide. Earlier open system implementations often used internal server storage.

An organization's operational costs soared due to the proliferation of departmental servers, which created unsecured, poorly managed, dispersed islands of information. For managing these computers and the data produced, there were initially relatively few regulations and procedures. Storage technology progressed from non-intelligent internal storage to intelligent networked storage to address these issues. The following are some notable developments in technology:

Redundant Array of Independent Disks (RAID): This technology was created to meet the needs of data in terms of cost, performance, and availability. It is still developing today and is a component of all storage designs, including DAS, SAN, and others.

Direct-attached storage (DAS): Direct-attached storage (DAS) is a form of storage that is linked directly to a host server or a cluster of servers. The server's storage options include internal and external options. The difficulties caused by insufficient internal storage capacity were eased by external DAS.

Storage area network (SAN): This is a specialized, high-performance Fiber Channel (FC) network that makes it easier for servers and storage to communicate at the block level. For the purpose of accessing its data, storage is partitioned and given to a server. Comparing SAN to DAS, you may see improvements in scalability, availability, performance, and cost.

Network-attached storage (NAS): Network-attached storage (NAS) is specialized storage for programs that serve files. In contrast to a SAN, it connects to an existing LAN and gives clients of all types' access to files. In comparison to general-purpose file servers, it provides better levels of

scalability, availability, performance, and cost advantages due to its design for delivering storage to file server applications.

IP-SAN (Internet Protocol SAN): IP-SAN, one of the most recent advancements in storage design, is a fusion of SAN and NAS technology. Data consolidation and availability are increased because to IP-provision SAN's of block-level communication across LANs and WANs.

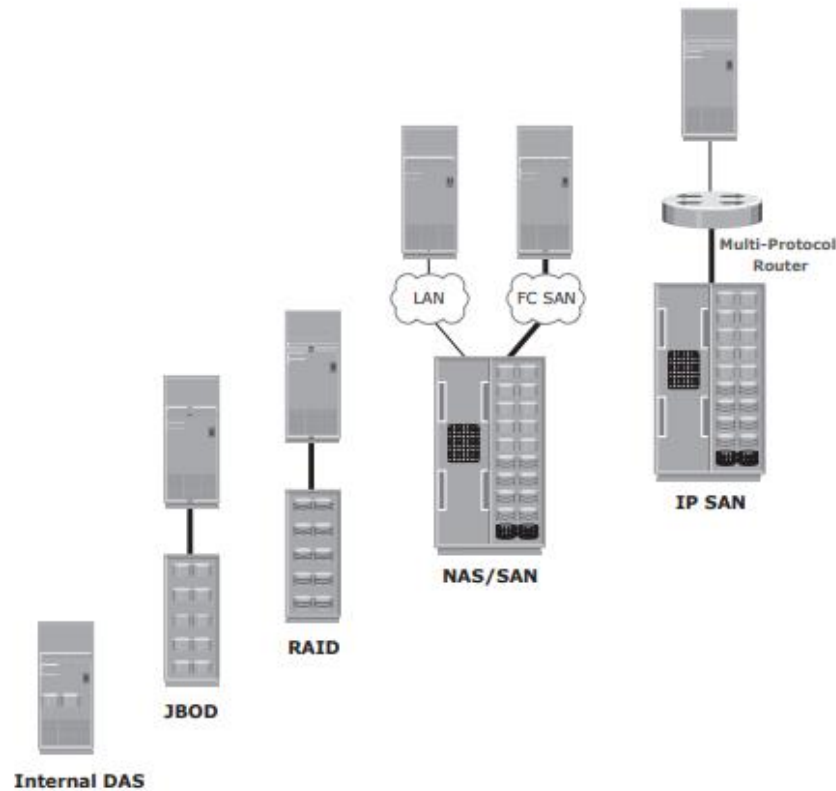


Figure 1.1: Represented the Storage Architecture Evolution.

Principal Features of a Data Center

The continued functioning of data centers is essential to a company's existence and prosperity. Organizations need a solid infrastructure to guarantee that data is always available. The emphasis is on storage systems even though the traits apply to every component of the data centre architecture. The many technologies and answers to achieve these needs are covered in this book.

Availability: A data center should make sure that information is accessible when needed. Businesses in the financial services, telecommunications, and e-commerce industries might lose millions of dollars every hour as a result of information not being available.

Security: To prevent unauthorized access to information, data centers must set up rules, processes, and core element integration.

Scalability: As a business grows, it often has to add extra servers, apps, and databases. Resources in the data center should grow according to needs without interfering with business activities.

Performance: Based on the necessary service levels, every component of the data center should operate at its peak efficiency.

Data integrity: The term "data integrity" refers to systems that guarantee that data is saved and retrieved precisely as it was originally received, such as correction codes for errors or parity bits.

Capacity: To store and handle massive volumes of data effectively, data center operations need enough resources. The data center must be able to expand its capacity in response to rising capacity demands while maintaining or improving availability. It is possible to manage capacity by either adding new resources or reallocating already-existing ones.

Manageability: A data center should make all of its components simple and integrated to administer. Automation and a decrease in the need for human (manual) involvement in routine chores may increase manageability.

Essential Components of a Data Center

A data center must have these five fundamental components in order to operate properly:

- A. **Application:** A computer software that supplies the reasoning behind calculations.
- B. **DBMS, or database management system:** gives a structured method for storing data in connected, logically arranged tables.
- C. **Host or compute:** A computing platform (hardware, firmware, and software) that serves as a host for databases and program.
- D. **Network:** A data route that enables communication between numerous networked devices is known as a network.
- E. **Storage:** A system for storing data persistently for future usage

Although these fundamental components are frequently treated and controlled separately, they all need to function together to meet data processing needs.

Responsibility a Data Center

Numerous responsibilities are involved in running a data center. The following are some of the important managerial activities:

- 1) **Monitoring:** It is the process of continuously acquiring data on numerous components and services that are active in a data center. Security, performance, availability, and capacity are among the elements of a data center that are observed.
- 2) **Reporting:** It is done on a regular basis regarding resource capacity, performance, and use. Establishing business justifications and payback of expenditures related to data center operations is made easier by reporting tasks.
- 3) **Provisioning:** The process of delivering the hardware, software, and other resources needed to run a data center is known as provisioning. Resources management is a key component of provisioning activities since it helps to meet capacity, availability, performance, and security needs.

The provisioning and management of data center infrastructure resources has undergone a significant transformation as a result of virtualization and cloud computing. To maximize their use, businesses are quickly implementing virtualization on different components of data centers.

The popularity of cloud computing is a result of ongoing cost pressure on IT as well as demands for on-demand data processing.

Information Lifecycle: The "changing in the value of information" through time is referred to as the information lifecycle. Data typically has the most value and is utilized most frequently when it is initially produced. Data becomes less valuable to the company and is accessed less frequently as it becomes older. Knowing the lifespan of information can aid with storage infrastructure deployment that is appropriate for the information's changing value. For instance, from the moment the purchase is placed until the warranty expires, information's value varies in a sales order application. The information is most valuable when a business executes a fresh sales order to ship the product. The client or order information need not be accessible in real-time after order fulfilment. The business can move this information to less expensive secondary storage with less stringent standards for accessibility and availability until a warranty claims or another circumstance necessitates its use. The business can archive or discard data once the warranty expires to make room for additional highly valuable information.

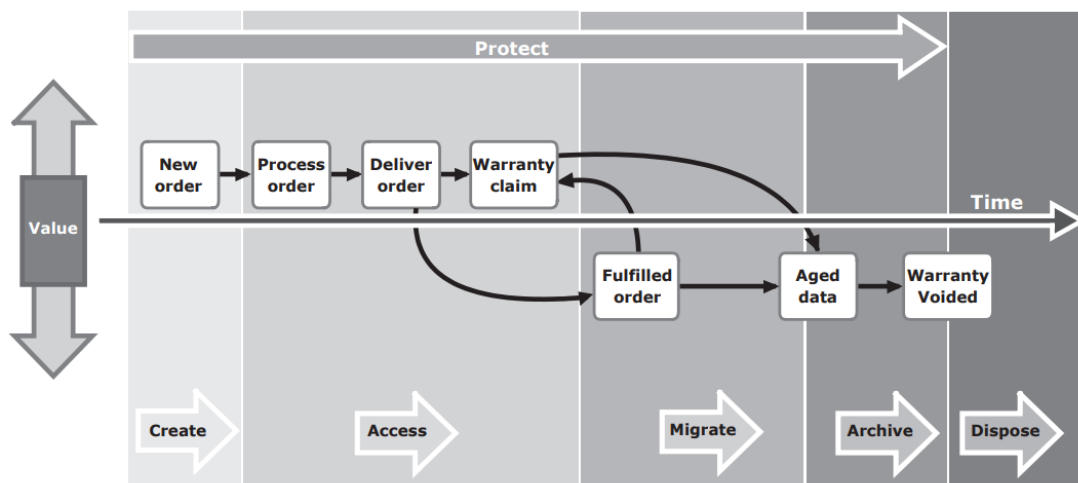


Figure 1.2: Information Lifecycle Management

Data centers may do this by making the best and most efficient use of their storage infrastructure. To support this infrastructure and take advantage of its advantages, an efficient information management strategy is needed. A proactive method called information lifecycle management (ILM) enables an IT company to efficiently manage data throughout its lifespan based on established business standards. This enables an IT company to maximize return on investment by optimizing the storage infrastructure. The following traits should be present in an ILM strategy:

Business-centric: To accommodate both the existing and anticipated expansion of information, it should be connected with critical business activities, processes, and applications.

Centrally managed: All of a company's information assets should be centrally managed and fall within the ILM strategy's scope.

Policy-based: ILM adoption shouldn't be limited to a few departments, according to policy. All business applications, procedures, and resources should be covered by ILM, which should be established as a policy.

Heterogeneous: An ILM approach should include all variations of operating systems and storage platforms.

Optimized: Because the worth of data varies, an ILM method should take into account the various storage needs and allot storage resources in accordance with the value of the information to the company.

Information Lifecycle Stages

The phases or stages that data undergoes as part of the information lifecycle are often used to explain the ILM process. Even while many of them are often conceptually similar, various resources frequently characterize these stages in different ways. The seven stages listed below provide a basic picture of what occurs with data during the course of its lifecycle:

Gather data: Companies constantly generate new data and get data from outside sources. Both manually and automatically created data are possible. Social media, the industrial internet of things (IoT), organizational collateral, user-generated content, customer feedback, sales records, and a variety of other sources may all be used as data sources.

Keep data: Data-generating and -collecting organizations need to develop efficient methods for storing their data. They could use file, block, or object storage technologies to keep the data. They could make use of various configurations and storage media, such storage area networks or network-attached storage. They may keep their data on-site, on the cloud, or in both places.

Control data: It is not sufficient to just store data. Additionally, organizations need to be able to handle such data well. They have to make sure that the data is secure, accessible, and complies with all applicable company, industry, and governmental standards. Additionally, they may categorize, de-duplicate, or compress the data, or they might develop a system for tracking their information and storage systems.

Modify the data: Few businesses just collect data and store it without changing it in some manner to make it more accessible and understandable. The data may be cleaned, filtered, aggregated, enhanced, merged, or modified in some other manner as part of this procedure to satisfy their business goals.

Use data: To guarantee that users and apps have the information they need to do business and complete their assigned responsibilities, data must be captured, stored, and transformed. Users may read, alter, share, or work together on data at this stage. They could also do data analysis or create reports using it.

Data archives: When data is no longer required on a regular basis, it is often stored in case it becomes necessary in the future for commercial purposes or to satisfy legal or regulatory obligations. Because data access needs are not very high, organizations often utilize slower and less expensive storage solutions. Not all data has to be preserved, even though this step in the ILM procedure might be crucial. For instance, IoT device data may only need to be kept until it has been collected and analyzed or until abnormalities have been found.

Eradicate data: The data is said to have reached the end of its usefulness and may be erased when an organization is certain that it is no longer required and is not constrained by legal or regulatory constraints. Because it lowers the quantity of data that has to be kept and the organization's potential liability, the destruction phase is crucial to the ILM process. Even if the information is

not required, maintaining and storing it comes at a cost, thus the sooner it can be securely erased, the better. Additionally, if data is not promptly removed, it may be more challenging for work with the new data and base business choices on that new data.

Workings of Information Lifecycle Management

By using a policy-based approach to data processing, information lifecycle management offers a centralized, uniform technique for managing the full data lifetime. Storage tiering and automation are also made possible by ILM. By doing this, data may be moved automatically depending on the appropriate regulations from one storage tier or format to another. Newer and more often accessed data is typically kept on faster, more costly storage media, whereas less important data is typically stored on slower, less expensive media. IT teams may establish various rules for various kinds of data throughout its lifecycle using the ILM method. ILM takes into consideration the fact that various forms of data depreciate in value at varying rates, with some maintaining worth far longer than others. Path management features, which monitor where data is in the storage cycle and make it simpler to retrieve stored data, may sometimes be included with ILM. ILM must be a company-wide initiative, incorporating policies and practices in addition to software and technological platforms, for it to be successful. When confronted with e-discovery demands, the ability to better manage and access information offers a crucial advantage for IT, the legal team, and the company, according to consulting firm Deloitte. ILM may "bring managerial rigor and controls" of data for the whole company, according to Deloitte.

Importance of information life cycle

Organizations value information lifecycle management because it ensures that information is successfully handled from creation to disposal throughout the course of its entire existence. ILM can assist businesses in increasing productivity, cutting expenses, and enhancing security. Benefits of information lifecycle management include:

Lessened dangers: One of the primary advantages of information lifecycle management is the reduction of information hazards. When information is handled effectively, your firm should be able to comply with a variety of laws and rules, reducing the chance of receiving fines and penalties that might end up costing your business a lot. Additionally, you enhance the decision-making process by protecting and providing only high-quality information to workers and managers, reducing the risks of making business choices without sufficient knowledge.

Savings in costs: The main advantages of information lifecycle management include quick information retrieval, backups, and storage, which may help your business save money that it can then spend to enhance services or launch new products.

Enhanced safety: Better safeguard your data from unwanted access and exploitation by being aware of where it is and how it is being used. You can better safeguard your data from unwanted access and exploitation by being aware of where it is and how it is being used.

More efficient government: ILM could increase organizationally advantageous management consistency and controls. ILM may provide the company as a whole the added advantage of improved information management. By ensuring that data is correctly handled throughout its lifespan, ILM may assist you in complying with rules more effectively.

Better performance: ILM may help you increase the efficiency of your IM system by transferring data to quicker storage alternatives as it gets busier.

Greater agility: One of the main advantages of information lifecycle management is increased agility, which may help you react more rapidly to changing business demands by making it simpler to access and utilize the appropriate data at the appropriate time.

Cloud Computing and Virtualization

Through the process of virtualization, physical resources like computation, storage, and networks are abstracted and given the appearance of being logical resources. In various incarnations, virtualization has been around in the IT industry for a while. Virtual memory utilized in computing systems and partitioning of raw drives are typical instances of virtualization. Pooling physical resources is made possible via virtualization, which also offers a consolidated picture of the capabilities of the physical resources. Storage virtualization, for instance, makes it possible for many shared storage devices to become a single huge storage unit. Similar to this, by implementing compute virtualization, the pooled physical servers' CPU capacity may be shown as the sum of all the CPUs' power (in megahertz). Additionally, virtualization makes it possible to manage shared resources centrally.

Physical resources that have been pooled may be used to develop and provide virtual resources. For instance, a virtual server with a certain amount of CPU power and memory may be constructed from a compute pool, just as a virtual disc with a certain capacity can be generated from a storage pool. By sharing pooled physical resources, these virtual resources increase the efficiency of using actual IT resources. The virtual resources' capacity may be increased or decreased in accordance with business needs without affecting users or applications. Organizations may reduce expenses connected with the acquisition and administration of new physical resources by using IT assets more effectively. Additionally, less physical resources need less energy and space, which improves economics and promotes green computing.

Organizations need to be flexible and agile in today's fast-paced and competitive climate in order to adapt to changing market demands. In order to satisfy diminishing or static IT expenditures, this causes fast growth and improvement of resources. Cloud computing effectively handles these problems. IT resources may be used by people or companies as a network-based service thanks to cloud computing. It offers very flexible and scalable computing that makes it possible to deploy resources as needed. The demand for computational resources, especially storage capacity, may be scaled up or down by users with little administration work or service provider involvement. Self-service requesting is made possible by cloud computing thanks to a completely automated request-fulfillment system. Because consumption-based metering is made possible by cloud computing, users only pay for the resources they really use, such as the number of CPU hours utilized, the volume of data transported, and the gigabytes of data saved.

Chapter 2

STORAGE SYSTEM ENVIRONMENT

Pradeepa P, Associate Professor

Department of Electrical and Electronics Engineering, Faculty of Engineering and Technology, JAIN (Deemed-to-be University), Karnataka – 562112

Email Id- p.pradeepa@jainuniversity.ac.in

Data is a collection of unprocessed facts from which inferences can be made. Examples of items that include data include handwritten letters, printed books, family photos, printed copies of mortgage documents that have been duly signed, bank ledgers, and airline tickets. The methods used for data creation and sharing before the invention of computers were restricted to fewer media, like paper and film. Similar information can now be transformed into more practical formats like an email message, an e-book, a digital image, or a digital video. Computers are capable of creating and storing this data as strings of binary digits (0s and 1s). This type of data is known as digital data, and the user can only view it later a computer processes it. The rate of data generation and exchange has drastically expanded with the development of computer and communication technology. A few of the elements that have fueled the expansion of digital data are listed below:

Enhanced data processing abilities: Processing and storage capacities are significantly increased by modern computers. This makes it possible to convert many kinds of content and media from analogue to digital formats. Less expensive digital storage Low-cost storage options are now available as a result of technological advancements and falling storage device prices. The rate at which digital data is generated and stored has risen as a result of this cost advantage.

Affordable and quick communication technology: Digital data sharing is now done considerably more quickly than in the past. . While an email usually only takes a few seconds to reach its target, a handwritten letter could take up to a week to get there. Smart apps and the proliferation of smart devices have both made major contributions to the creation of digital content. These devices include smartphones, tablets, and newer digital gadgets.

Application-based data storage and retrieval are used by users. Hosts are the machines that these programs are installed on. Simple laptops to intricate server clusters may function as hosts. Physical components (hardware devices) make up a host, while logical components are used to interact with one another (software and protocols). The logical and physical elements of a host affect both data access and the storage system environment's overall performance.

Storage System Environment's Components

The host, connection, and storage are the three primary elements of a storage system environment.

Host: Application-based data storage and retrieval are used by users. Hosts are the machines that these programs are installed on. Simple laptops to intricate server clusters may function as hosts.

There are physical components in a host logical components are used by (hardware devices) to interact with one another (software and protocols). The physical and logical components of a host affect both data access and the storage system environment's overall performance. , the logical parts of the host are described in more depth.

Physical elements

Three essential physical elements make up a host:

The physical parts connect with one another via a bus, which is a communication channel. The CPU is linked to other parts, including storage and I/O devices, through a bus.

CPU

There are four primary parts that make up the CPU:

Arithmetic Logic Unit (ALU): The CPU's basic building element is the arithmetic logic unit (ALU). It can execute addition, subtraction, and Boolean operations as well as other mathematical and logical operations (AND, OR, and NOT).

Control Unit: A digital circuit that coordinates the CPU's functioning and regulates CPU activities.

Register: A grouping of fast storage facilities. Due to their closeness to the ALU, the registers provide quick access to intermediate data that the CPU needs to process an instruction. CPUs generally only have a few registers.

Cache at Level 1 (L1): It may be found on current CPUs and stores information and code instructions that the CPU will probably require in the near future. Despite being slower than registers, the L1 cache has greater storage.

Storage

Data is stored in memory and storage devices, either permanently or momentarily. Storage devices employ either magnetic or optical media, whilst memory modules are built utilizing semiconductor chips. Data access is made faster by memory modules than by storage medium. On a host, there are often two kinds of memory:

Random Access Memory (RAM): This enables immediate access to any region in memory and allows for the writing and reading of data. RAM is a volatile sort of memory; to sustain memory cell content, it needs a steady source of electricity. When the system's power is cut off or otherwise stopped, data is deleted.

Read-Only Memory (ROM): Non-volatile Read-Only Memory (ROM) that only permits data reading. ROM stores information needed to execute internal processes, such system initialization. Less costly than semiconductor memory are storage devices. The following are some examples of storage devices: Hard drive (magnetic), CD-ROM and DVD-ROM (optical), Diskette floppy (magnetic), drive for tapes (magnetic)

I/O Devices

Data may be sent to and received from a host using I/O devices. This kind of communication might be any of the following:

Communications between the user and the host are handled by basic I/O devices such as the keyboard, mouse, and display. Users may input data and see the outcomes of operations using these devices.

Host-to-host communication is made possible by tools like a modem or Network Interface Card (NIC). A host bus adaptor handles communications from the host to the storage device (HBA). An ASIC card called an HBA performs I/O interface tasks between the host and the storage, alleviating the CPU of extra I/O processing effort. HBAs also provide ports, which are connection outlets, to link the host to the storage device. A host might have many HBAs.

Connectivity

Interconnection between hosts or between a host and any other peripheral devices, such as printers or storage devices, is referred to as connectivity. Here, the connection between the host and the storage device is the main topic of discussion. Physical and logical connection components may be distinguished in a storage system environment. The hardware components that link the host to the storage are known as the physical components, and the protocols that are used to communicate between the host and storage are known as the logical components of connection.

Physical Connectivity Components

Bus, Port, and Cable are the three physical elements that make up the connection between both the host and storage. Processor HBA Port Cable BUS Disk. The bus is a network of pathways that makes it easier for data to be transferred between computer components, such as from the CPU to the memory. A specific outlet known as a port makes it possible for the host and external devices to communicate. Cables use copper or fiber optic media to link hosts to internal or external devices.

Physical components interact with one another through a bus by exchanging bits of data (control, data, and address). One of the two methods listed below is used to transfer these bits across the bus: Bits are sent in a sequential manner via a single route. Both unidirectional and bidirectional transmissions are possible. Bits are sent concurrently through many pathways in parallel. Parallel might have two directions as well. The quantity of data that can be conveyed over a bus at once is determined by the width of the bus, which is its size. A bus's width may be compared to a highway's number of lanes. A 64-bit bus, for instance, may send 64 bits of data at once whereas a 32-bit bus can send 32 bits at a time. Each bus has a clock speed that is expressed in MHz (megahertz). These show how quickly data is transferred between the bus's end locations. Applications may run more quickly thanks to a fast bus that facilitates quicker data transmission. As means of data transport on a computer system, buses fall into one of the following categories: The bus that transfers data from the CPU to memory is known as the system bus. A fast data channel that connects directly to the CPU and transmits information between peripheral devices, such as storage devices, and the processor is known as a local or I/O bus.

Logical Connectivity Components

Peripheral component interconnect is a common interface protocol used to link the local bus to a peripheral device (PCI). Integrated Device Electronics/Advanced Technology Attachment (IDE/ATA) and Small Computer System Interface are the interface protocols used to connect to disc systems (SCSI).

PCI: PCI is a specification that harmonizes the information-exchange process between PCI expansion devices and the CPU, including network cards and modems. The link here between CPU

and associated devices is made possible through PCI. The PCI plug-and-play functionality makes it simple for the host to identify and set up additional cards and devices. A PCI bus may have a width of 32 bits or 64 bits. A 32-bit PCI bus is capable of 133 MB/s of throughput. The PCI Express bus is a much faster and more efficient variant of the PCI bus.

IDE/ATA: The most widely used interface protocol for contemporary drives is IDE/ATA. This protocol provides outstanding performance at a fair price.

SCSI: High-end computers now favor SCSI over other protocols. Due to its greater price, this interface is far less often utilized than IDE/ATA on desktop computers. Initially, SCSI was utilized as a parallel interface to link devices to a host. SCSI has been improved and currently incorporates a large number of related standards and technologies.

Storage

The most crucial element of the storage system environment is the storage device. Magnetic or solid state media are used in a storage device. Magnetic media are used in discs, cassettes, and diskettes. An example of an optical media storage device is a CD-ROM, while an illustration of solid state media is a disposable flash memory card. Because they are very inexpensive, tapes are a common storage medium used for backup. Data centers formerly housed a sizable number of tape drives and handled thousands of reels of tape. But tape has the following disadvantages: Along the tape's length, data is stored in a linear fashion. Data searches and retrievals are carried out progressively, and it always takes a few seconds to obtain the data. Random data access is thus time-consuming and sluggish. As a result, tapes aren't a good choice for applications that need quick, real-time access to data. Data saved on tape can only be used by one program at a time in a shared computer environment since it cannot be accessed by several apps at once. The read/write head of a tape drive makes contact with the tape surface, which causes the tape to deteriorate or get worn down over time. The costs involved with controlling tape medium and the space required for storing and retrieving data from tape are substantial.

Despite its drawbacks, tape is often used because to its affordability and portability. High capacity media and fast drives are products of ongoing tape technology progress. To boost data throughput, modern tape libraries are equipped with extra memory (cache) and/or disc drives. Today's tapes are a component of an end-to-end data management system with these and more intelligence, notably as a cheap option for long-term data storage and seldom accessed data storage. Single-user computer settings, optical disc storage is common. On desktop or laptop computers, people typically utilize it to store images or as a backup media. It may also be used to move tiny quantities of data from one standalone system to another or as a distribution medium for specific applications, such games. The use of optical media as a corporate data storage solution is constrained by the speed and capacity restrictions of optical discs. One benefit of optical disc storage is the capacity to write once and read many (WORM). A WORM gadget is an illustration of one. Optical discs may be used as low-cost options for long-term storage of relatively modest quantities of fixed material that won't change once it is generated since them, in part, ensure that the content hasn't been changed. Jukeboxes, an array of optical discs, are still used as a fixed-content storage option. Optical discs also come in CD-RW and other DVD iterations. Disk drives are the most common kind of storage that contemporary computers utilize to store and retrieve data for online applications that need a lot of processing power. Disks provide quick access to arbitrary data locations. As a result, data may be swiftly written or retrieved for several concurrent users or

applications. Disks also have a huge capacity. Multiple discs are built into disc storage arrays to increase capacity and improve performance.

Inside its casing, the hard disk, which typically offers storage for applications and data in a computer, contains four essential parts: a platter for storing information, a spindle for rotating the platters, a read/write arm for read and write data, and an actuator for controlling the read/write arm's actions. Only the best skilled IT specialists should try to operate on the parts within a hard disc.

Platters: The 1s and 0s that comprise up your data are kept on circular discs within the hard disc called platters. A magnetic surface is present on plates that are constructed of aluminium, glass, or ceramic in in order to permanently preserve data. To improve the total capacity of bigger hard drives, multiple platters are employed. To keep information structured and findable, data is stored on the platters in track, sectors, and cylinders.

A Spindle: The platters are maintained in place and rotated as needed by the spindle. The hard drive's capacity and speed are determined by its revolutions per minute rating. 7,200 RPM is the normal speed of an internal desktop drive, while higher and slower rates are possible. To provide the read/write arm access, the spindle maintains the platters at a constant distance from one another.

Reader/Writer Arm: The read/write heads actually read and write data on the disc platters by converting the magnetic surface into an electric current, while the read/write arm regulates the movement of the heads. The arm, sometimes referred to as the head arm or actuator arm, ensures that the heads are in the proper position depending on the data that has to be retrieved or written. Usually, each platter side has a single read/write head that floats 3 to 20 millionths of an inch above the platter surface.

Actuator: The actuator, also known as the head actuator, is a tiny motor that receives commands from the drive's circuit board to direct the movements of the read/write arm and manage the transport of data onto and off of the platters. It is in charge of making sure the read/write units are always located precisely where they should be.

Controller: At the base of a disc drive is a printed circuit board that serves as the controller. A CPU, internal memory, a circuit, and firmware make up the device. The firmware regulates the spindle motor's speed and power supply. Additionally, it controls communication between the host and the drive. Additionally, it optimizes data access while controlling R/W operations by rotating the actuator arm and switching between R/W heads.

Physical Disk Structure: Track, which are concentric spheres on the platter centered on the spindle, are where data is stored on a disc. Beginning at zero and working outward from the platter's edge, the tracks are numbered. How closely tracks are arranged on a platter is influenced by the number of tracks per inch (TPI) on the platter (also known as track density). Sectors are the smaller, numbered units that make up each track. The smallest unit of storage that may be individually addressed is a sector. The drive manufacturer performs a formatting process on the platter to write the track and sector structure. Depending on the drive, different tracks have different numbers of sectors. 17 sectors were on each track of the earliest discs for personal computers. Sectors on a single track are substantially more numerous on recent discs. Depending on the platter's physical size and recording density, there might be thousands of tracks on it.

Although certain discs may be formatted with bigger sector sizes, a sector typically carries 512 bytes of user information. A sector holds user data as well as additional data, like the sector number, head or platter amount, and track number. Although keeping this information takes up disc space, it aids the controller in finding the data on the drive. Consequently, the capacity of a formatted disc and an unformatted disc varies from one another. Manufacturers of hard drives often list the unformatted capacity; for instance, a 500GB disc will really only carry 465.7GB of user data and 34.3GB of metadata. The collection of identical tracks on each drive platter's two surfaces is known as a cylinder. Instead of using the track number, the position of drive heads is identified by the cylinder number.

Zoned Bit Recording: Since the tracks on the platters are concentric, the outside tracks may store more data than the inner tracks because they are physically longer. Data density was poor on the outer tracks of early disc drives since they had the same number of sectors as the inner tracks. The usage of the limited space here was ineffective. Zone bit recording makes optimal use of the disc. This process divides tracks into zones according to how far they are from the disk's center. The zones are designated, with zone 0 being the most outside. Each zone is given the proper amount of sectors per track, thus a zone close to the platter's centre has less sectors per circuit than a zone on the outside. However, the number of sectors on each track is the same within a zone.

Block Logical Addressing: Earlier drives required the host operating system to be aware of the geometry of each disc in use and employed physical addresses made up of the cylinder, head, and sector (CHS) number to refer to particular places on the disc. By employing a linear address to retrieve physical blocks of data, logical block addressing (LBA), simplifies addressing. The host just has to be aware of the disc drive's block size since the disc controller converts LBA to a CHS address. A 1:1 mapping between the logical blocks and physical sectors is used. The engine has four cylinders, eight heads, and eight sectors per track. The number of blocks spans from 0 to 255 since $8 \times 8 \times 4 = 256$ blocks are involved. Every block has a distinct address. A 500 GB disc with a formatted capacity of 465.7 GB will contain more than 976,000,000 blocks if the sector size is assumed to be 512 bytes.

Other Ingredients: The front-end circuit board, together with the ports at the end of the drive, handles input and output signals in addition to the case that keeps all of the hard disk's components together. Regardless of the drive's kind, it has a single connection for a power supply and a single port for sending and receiving data and commands to the rest of the system.

Disc Performance

Partitioning a disc: The way disc partitions are managed has an impact on the disk's overall performance. In situations when data is continuously being accessible across the different partitions, creating too many fragments may result in a loss in overall disc performance. When distinct partitions are used for various data kinds, disc partitioning improves overall performance. For instance, the OS, data, and games all have their own partitions. When loading a specific software, the magnetic head only travels inside the partition it is in it does not move across partitions.

Driving Form Factor: One factor that sets one hard disc apart from another is this one. It relates to the drive's real size. The overall storage capacity and associated number of internal platters are specified by the drive form factor, which also establishes the disk's transfer rate. The speed at which information is written to or retrieved from the disc is referred to as the transfer rate.

Spindle Speed: The hard drive's speed in RPMs is determined by the spindle's rotational speed. Faster data accessibility and retrieving times result from the drive's platters rotating at a rapid rate, or high spindle speed.

Plate Dimensions: Although spindle speed affects data transmission rates significantly, it is not the sole one. The drive's data transmission speeds are also impacted by the platter's diameter. This is due to the fact that although the spindle's speed remains constant, the discs' outside perimeters cover a larger area than their inner perimeters. This implies that the diameter will have an impact on how quickly the magnetic head reaches various regions of the platter. At the same spindle speed, platters with a big diameter have greater transfer rates than platters with a diameter.

Access Times: On a certain platter, data may be kept in any location. The magnetic head needs to reposition itself to the proper platter and its placement for every read or write request. Various criteria make up access times. Write seek time is the amount of time it takes the magnet head to find free space on a platter so that it may execute a write operation. Time spent by the magnetic head seeking the data to be read is known as the read seek time. Full stroke time: the time it takes for the electromagnetic head to fully rotate around each platter. The amount of time it takes the magnetic head to transfer successive tracks on a particular platter from one to the next.

Buffer Memory: When data has to be transferred quickly from one disc region to another, buffer memory is a special area of memory set aside. These days, this reserved area may have a 16MB capacity. The speed of data access increases with buffer size.

Flash drives

Users of storage continue to expect ever-increasing performance standards for their business applications due to the proliferation of information. High I/O demands were traditionally satisfied by adding extra discs. The situation has altered as a result of the availability of enterprise level flash drives (EFD). Flash drives, commonly known as solid state drives (SSDs), are modern drives that provide the ultra-high performance needed for applications with high performance requirements. Flash drives store and retrieve data using semiconductor-based solid state memory (flash memory). Because flash drives don't have any moving components, they don't have search or rotational latencies as traditional mechanical disc drives have. Flash drives provide a high IOPS with very fast reaction times. Moreover, because they are semiconductor-based, flash drives use less power than mechanical drives.

Flash drives are particularly well suited for workloads requiring random reads with tiny block sizes and constant low response times. Flash drives are beneficial for applications that require to handle enormous volumes of data fast, such as real-time data feed processing, electronic trading systems, and currency exchange. EFD offers up to 30 times the throughput and up to a tenth the reaction time of traditional mechanical disc drives. Additionally, flash drives may store data with up to 98 percent less power usage per I/O and up to 38 percent less energy per TB than conventional disc drives. Despite costing more on a per-GB basis, flash SSDs provide superior total cost of ownership (TCO). Businesses may achieve application performance standards by using flash drives with a lot fewer drivers. This decrease not only results in drive cost reductions, but also in power, temperature, and space usage savings. Less drives in the surroundings also implies lower management costs for the storage system.

Flash Drive Components and Architecture

To preserve compatibility, flash drives employ connections and a physical form factor comparable to mechanical disc drives. This makes switching a mechanical disc drive in a storage array container for a flash drive simple. The controller, I/O interface, mass storage (a group of memory chips), and cache are the essential parts of a flash drive. The I/O system provided power and data access, while the controller controls how the drive operates. Data is stored using a mass storage system, which consists of nonvolatile NAND (negated AND) memory chips. Cache acts as a temporary storage area or buffer for data processes.

Data access on a flash drive occurs via a number of parallel I/O channels that go from the drive controller to the flash memory chips. Generally speaking, the internal bandwidth of the drive increases along with the amount of flash memory chips and channels, which in turn increases drive efficiency. Typically, flash discs contain eight to 24 channels. Blocks and pages are the logical units through which memory chips in flash drives are structured. The smallest item that can be read or written on a flash disc is a page. Blocks are collections of pages. 32, 64, or 128 pages may make up a block. There is no set size for pages, however the most common sizes are 4 KB, 8 KB, and 16 KB. A page stretches over a succession of related data blocks because flash devices mimic mechanical drives that use logical block addresses (LBAs). For instance, a 4-KB page would consist of eight consecutive 512-byte data blocks. A read operation in flash SSDs may occur at the page level, while a write or erase action can only occur at the block level.

Characteristics of a USB flash drive

- 1) It doesn't have any delicate moving components and utilizes less electricity.
- 2) This device's data storage is resistant to dust, magnetic fields, and mechanical trauma. These characteristics of USB flash drives make them convenient for travelling.
- 3) In comparison to other gadgets, it can store more information.
- 4) Many USB flash drives are made specifically to have a waterproof capability and to be hard using rough rubber and metal. When immersed in water, this sort of drive does not lose its memory.

Logical Components of the Host

The software programs and protocols that enable data communication with the user as well as the physical components make up a host's logical components. The logical parts of a host are as follows:

- A. Operating system
- B. Device drivers
- C. Volume manager
- D. File system
- E. Application

Operating System: The entire computing environment is under the control of the operating system. It functions as a bridge between the application and the hardware of the computer system. Data access is one of the features it offers the application. Additionally, the operating system keeps an eye on human activity and the environment and reacts accordingly. It coordinates the distribution of hardware resources as well as the organization and control of hardware components. It offers fundamental security for using and gaining access to all managed resources. In addition

to managing the file system, volume manager, and device drivers, an operating system also carries out fundamental storage management operations.

Device Driver: A device driver is specialized software that enables the operating system to communicate with a particular device, like a hard disc, printer, or mouse. A device driver gives the operating system the ability to identify the device and to access and manage devices through a common interface (offered as an application programming interface, or API). Device drivers are operating system- and hardware-dependent.

Volume Manager: The operating system first saw an HDD as a collection of continuous disc chunks. The file system or other data entity utilized by the operating system or application would be allotted the entire HDD. Lack of flexibility had the drawback of making it difficult to expand the file system as an HDD ran out of room. Allocating the entire HDD to the file system frequently led to underutilization of storage as the HDD's storage capacity rose. In order to increase HDD flexibility and use, disc partitioning was introduced. An HDD is partitioned into logical containers known as logical volumes during partitioning (LVs) For instance, a sizable physical drive might be divided into several LVs to maintain data in accordance with the needs of the file system and applications. When the hard disc is first installed on the host, the partitions are made from collections of adjacent cylinders. Without being aware of partitioning or the physical makeup of the disc, the host's file system accesses the partitions.

File System: A file is an organized grouping of linked documents or data that has a name. An organized hierarchy of files is a file system. Access to data files stored on a disc drive, a disc partition, or a logical volume is made simple by file systems. Host-based logical structures and software procedures that manage file access are required by a file system. Users can create, change, remove, and access files using this functionality. The owner's permissions, which are also upheld by the file system, determine who has access to the files on the drives. A file system uses directories, which are containers for storing pointers to numerous files, to store data in a structured hierarchical fashion. Each file system keeps track of a pointer map to the files, directories, and subdirectories that make up the file system. Here are some examples of common file systems:

- A. FAT 32 (File Allocation Table) for Microsoft Windows
- B. NT File System (NTFS) for Microsoft Windows
- C. UNIX File System (UFS) for UNIX
- D. Extended File System (EXT2/3) for Linux

Along with the files and directories, the file system also contains a variety of additional related entries known as metadata. For instance, the superblock, inodes, and list of free and used data blocks make up the metadata in a UNIX system. In order for a file system to be regarded as healthy, its metadata must be consistent. The type of the file system, creation and update dates, size and layout, the amount of resources that are currently available (such as the number of free blocks and inodes, etc.), and a flag indicating the file system's mount status are all contained in a superblock. Every file and directory has an inode attached to it that contains details about the file's length, ownership, access privileges, date and time of its most recent update, the number of links, and the addresses for locating the spot on the physical disc where the data is kept.

Application: A computer program known as an application offers the reasoning behind computations. It offers a connection between the user and the host as well as between different hosts. Traditional business applications that use databases typically have a three-tiered

architecture, with the front-end tier being the application's user interface, the middle tier being the computing logic or the application itself, and the back-end tier being the underlying databases that house the data. The application asks read/write (R/W) operations be carried out on the storage devices from the underlying operating system. The database, which in turn leverages the OS services to carry out R/W operations to storage devices, can be stacked with applications. Between the front-end and back-end levels, these R/W operations (also known as I/O operations) allow for transactions.

- **Block-Level Access:** The fundamental method of disc access is block-level access. By setting the logical block address, data is stored and retrieved from discs in this sort of access. The geometric arrangement of the discs serves as the basis for deriving the block address. The fundamental unit of data storage and retrieval for an application is defined by the block size. When an I/O operation is carried out, databases like Oracle and SQL Server specify the block size for data access as well as the location of the data on the disc in terms of the logical block address.

File-Level Entry: Block-level access is an abstraction of access at the file level. The name and path of the file must be specified in order to grant file-level access to the data. It makes use of the underlying block-level storage access and shields the application and DBMS from the difficulties of logical block addressing (LBA).

Data Access from a Host

Applications use the underlying infrastructure to access and store data. Operating system, networking, and storage are the main elements of this architecture. The host may have an internal storage device or an external storage device. In either scenario, the host controller card uses predefined interfaces, such as IDE/ATA, SCSI, or Fibre Channel, to access the storage devices (FC). Small and personal computer settings often employ IDE/ATA and SCSI to access internal storage. In order to access information from an external storage device, FC and iSCSI protocols are utilised (or subsystems). The host may be directly linked to external storage devices or via a storage network. Direct-attached storage refers to storage that is directly linked to the host (DAS). Because it serves as the basis for storage networking technologies, understanding access to data across a network is crucial. One of the following methods may be used to access data across a network: block level, file level, or object level.

Typically, the application specifies the file name and location when requesting data from the file system (or operating system). The file system transmits the request to the storage device after mapping the file parameters to the logical block address of the data. The storage device gets the data after converting the logical block addressing (LBA) to a cylinder-head-sector (CHS) address. As shown in the Figure 1 Block-level access involves creating the file system on a host and accessing data at the block level via a network. In this scenario, the host is given access to raw discs or logical volumes in order to build the file system. As shown in the Figure 2 a file-level request is transmitted across a network, and the file system is built on a separate file server or at the storage side. This technique has more overhead than accessing data at the block level since data is accessible at the file level. Data is accessible through a network in the form of self-contained objects with a specific object identification under object-level access, which is an intelligent progression.

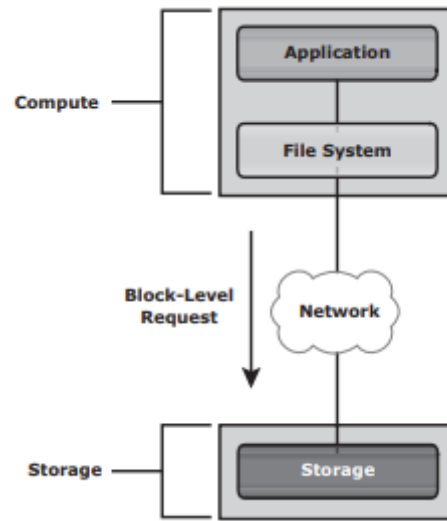


Figure 2.1: Represented the Block-Level Access




Figure 2.2: Represented the File-Level Access

Technologies for storage networks

InfiniBand, Fiber Channel, and Ethernet are the three network fabrics that are used most frequently in modern data centers. Although they have certain differences, all three offer a framework for connecting resources and sharing data among them in order to move data as quickly and securely as possible. These three major network topologies are still at war with one another over speed. And even when new storage networking technologies appear, this is unlikely to change.

Technologies for Storage

Traditional technologies like cassettes are unable to meet the demand for accessing and safeguarding this data as a strategic asset. Due to the tape medium's inherent problems, additional storage is now required, or the tape media may be completely replaced. The network is the

foundation of the present storage architecture. The most important issue for organizations nowadays is downtime. The high costs associated with downtime might result in service interruptions and unhappy customers. Federal rules have also produced strict compliance criteria for data security and high availability. There are now essentially three configurations for storage networks:

- A. **Direct Attached Storage (DAS):** Direct Attached Storage (DAS) is a time-honored technique for directly coupling storage devices and servers via a communication link between them known as buses. It is a specialized communication line that is independent from the network, and an intelligent access controller grants access (e.g. SCSI, SATA). Additional storage systems are installed when servers require more room. Additionally, using this technique, two servers may mirror one other.
- B. **Network Attached Storage (NAS):** Network Attached Storage (NAS) is a storage architecture with file-level access in which storage units are connected directly over a LAN. It enables diverse computer systems to access files on a file-level. The extra layer was added to deal with shared storage files. Typically, this technology makes use of the IP-based Network File System (NFS) or Common Internet File System (CIFS). This type of solution has the benefit that multiple servers can access the files kept on NAS storage devices. Although the servers may run on different operating systems (Linux or Windows), they all employ the same standard IP protocols.
- C. **Storage Area Network (SAN):** Storage Area Network (SAN) SAN is connected remotely from the servers, as opposed to locally to a server, like DAS. SAN offers block level access rather than file level access to a server. One file can give random access to any block on the storage device and contains many blocks. In critical data environments, SANs offer high availability and reliable business continuity. SANs are frequently switched fabric topologies connected by Fiber Channel (FC). The SAN Switch/Director is a crucial component in the development of the SAN-based storage architecture. One or more SAN switches/directors provide redundancy within the pathways to the storage units by connecting each SAN storage unit to each server. As a result, there are more channels for communication and one central switch is no longer a single point of failure.

High Speed connection to both open and private mainframe variations is provided by fiber channel communication (FCP). For supporting SANs, Ethernet has numerous benefits similar to those of Fiber Channel. Fiber Channel over Ethernet (FCoE) enables the convergence of storage and IP protocols by allowing FC traffic to travel across Ethernet-based networks.

Technologies for connectivity and storage networking

- A. **Fiber Channel (FC):** The networking technology known as Fiber Channel, also referred to as FC, is the industry standard that allows for the high-speed, low-latency, and lossless transport of block data to servers. It requires specialized cabling (which is where the "Fiber" as in fiber optical, in Fiber Channel comes in) and networking components in typical deployments. Gigabit per second (Gbps) and higher data transfer speeds are supported through fiber channel.
- B. **Ethernet over Fiber Channel (FCoE):** The FCoE protocol offers lossless Fiber Channel traffic transfer over 10 Gbps Ethernet networks, as its name implies. Instead of using Fiber Channel only for their SANs and Ethernet for the rest of their computer networking needs, this enables enterprises to standardize on Ethernet-based networking hardware.

- C. **Internet Small Computer Systems Interface (iSCSI):** Over TCP/IP networks, block-level storage operations are possible with iSCSI, another SAN protocol. Data is transferred between iSCSI storage devices and SCSI commands are packaged into packets that may be transmitted over a network by an iSCSI initiator, a piece of hardware or program that runs on a server.
 - D. **Independent Redundant Array of Disks (RAID):** RAID, which was formerly an abbreviation for redundant array of cheap discs, is not solely a storage networking term but has a significant impact on how networked storage systems function. The same data is typically stored using RAID across numerous hard drives to ensure that it is still available in the event of a storage drive failure. This is regarded as a crucial element of numerous enterprise storage system deployments that are trusted with important data.
 - E. **InfiniBand:** For enterprise organizations that need the quickest performance out of their networked storage systems, InfiniBand is a high-speed networking solution used to connect servers in high-performance computing (HPC) scenarios. Speeds measured in Gbps are supported by the switched fabric technology.
-

CHAPTER 3

REDUNDANT ARRAY OF INDEPENDENT DISKS

Dimple Bahri, Assistant Professor
Department of Civil Engineering, Faculty of Engineering and Technology, JAIN (Deemed-to-be University),
Karnataka – 562112
Email Id- dimple.bahri@jainuniversity.ac.in

For improved performance, data redundancy, or both, RAID, which stands for "Redundant Arrays of Independent Drives," uses a combination of many discs rather than a single disc. At the University of California, Berkeley, in 1987, David Patterson, Garth A. Gibson, and Randy Katz came up with the phrase.

Data Redundancy:

Although it consumes more capacity, data redundancy improves disc dependability. This means that if the same data is also backed up onto another disc, in the event of disc failure, we may retrieve the data and continue the procedure. In contrast, if the data is dispersed across just a few drives without the use of RAID, the loss of a single disc could have an impact on the entire set of data.

Important criteria for evaluating a RAID system

- A. **Reliability:** How many disc errors can the system withstand in terms of reliability?
- B. **Availability:** How much of the session time is spent in uptime mode, or in other words, how readily usable is the system?
- C. **Performance:** How quickly does it respond? How quickly is the job being processed, or throughput? Be aware that there are many more parameters involved in performance than simply the two. The usable capacity is accessible to the user given a set of N discs, each with B blocks.

Characteristics

- 1) Each RAID level has unique properties, including: Fault-tolerance, or the capacity to withstand one or more disc failures.
- 2) Performance, which displays how the read and write speeds of the complete array have changed as compared to those of a single disc.
- 3) The array's capacity, which is based on how much user data can be written to the array. The RAID level determines the array capacity, which is not always the same as the total size of the RAID member discs.

Level of RAID

Either a driver or an isolates produced card (a hardware RAID controller) may include the software necessary to carry out RAID functionality and manage the discs. Software RAID capability is included in several Windows versions, including Windows Server 2012 and Mac OS X. Although more expensive than pure software RAID controllers, hardware RAID controllers perform better,

particularly with RAID 5 and 6. SATA, SCSI, IDE, or FC are just a few of the interfaces that RAID systems may be employed with (fiber channel.) While some systems include a FireWire or SCSI connection for the host system, others employ SATA discs inside. JBOD, which means for Just a Bunch of Disks, is a designation for discs in a storage system that is sometimes used. This indicates that the drives function as stand-alone discs without using a particular RAID level. For drives that house swap files or spooling data, this is often done. The most prevalent RAID levels is provided below:

Striping at RAID level 0: In a RAID 0 system, data is divided into blocks and written across all of the array's devices. This provides better I/O speed by employing many drives (at least 2) concurrently. Using numerous controllers, preferably one per disc, may improve this speed even further. For non-critical data storage that must be read/written quickly, such as on an image restoration or video editing station, RAID 0 is the best option. Considering mounting one disk in the folder path of the second drive if you wish to utilize RAID 0 only to combine the storage space of two drives into a single volume. This has the benefit that a single power failure has no bearing on the data on the second disc or SSD drive and is supported by Linux, OS X, as well as Windows.




Figure 3.1: RAID 0

Benefits of RAID 0

- 1) Great performance is provided by RAID 0 for both read and write activities. The overhead of parity controls does not exist.
- 2) There is no expense since all storage space is being utilized.
- 3) The technology is simple to use.

Drawbacks of RAID 0: RAID 0 cannot tolerate errors. The RAID 0 array's data is lost if even one disc fails. For systems that are mission-critical, it shouldn't be utilized.

Mirroring at RAID level 1: By writing data to a mirror drive in addition to the data drive (or collection of data drives), data are saved twice (or set of drives). In the event of a drive failure, the controller continues to operate and retrieve data using either the data drive or the mirror drive. For

a RAID 1 array, need at least two discs. For example, mission-critical storage for accounting systems is best served by RAID-1. It is also appropriate for tiny systems with only two data discs.




Figure 3.2: RAID 1

Benefits of RAID 1

- 1) Excellent read and write speeds are provided by RAID 1, and the latter is equivalent to those of a single disc.
- 2) Data does not need to be rebuilt if a drive dies; it just has to be transferred to the backup drive.
- 3) A very basic technique is RAID 1.

Drawbacks of RAID 1

- A. The key drawback is that because every data is written twice, the effective storage capacity is just half of the overall drive capacity.
- B. A hot swap of a failing disc is not always possible with software RAID 1 solutions. Therefore, replacing the faulty disc requires shutting down the computer to which it is linked. This could not be appropriate for servers that are being accessed concurrently by multiple users. These systems often use hardware processors that enable hot switching.

RAID 4

Similar to RAID 0, RAID 4 strips data over many drives. In order to provide redundancy, it additionally saves parity information for each disc separately on a dedicated drive. The parity disc in the figure below is disc 4, which contains the parity blocks A_p , B_p , C_p , and D_p . Therefore, the data may be recreated using the symmetric cryptography of the failing disc if one of the discs is involved. Compared to RAID 1, space is handled more effectively here since parity information takes up far less room than replicating the drive. Because all the parity data is written to a single disc, which is a bottleneck, the write performance becomes poor. As we shall see in a moment, RAID 5 has a solution for this issue.




Figure 3.3: RAID 4

Advantages of Raid 4

- A. Data redundancy that is economical in terms of memory cost
- B. Data striping improves read operation performance

Drawbacks of Raid 4

- A. Slow write operation
- B. Redundancy in the data is lost if the specialized parity disc malfunctions.

RAID level 5: Parity-based striping: The most popular secure RAID level is RAID 5. At least three drives are needed, however up to sixteen may be used. A binary pattern of all the current block is copied to one device, and blocks are striping across the drives. As seen in the figure below, the integrity data are distributed over all drives rather than being written to a single fixed disc. If any of the other information blocks' data is no longer accessible, the computer may recalculate it using the parity data. In other words, a RAID 5 array may survive a single disc failure without removing information or the ability to retrieve data. Although software may be used to implement RAID 5, a hardware controller is advised. On these controllers, additional cache memory is often employed to increase write speed. A strong all-around solution, RAID 5 combines effective storage with top-notch security and passable speed. It is perfect for files and application servers with few data discs.




Figure 3.4: RAID 5

Benefits of RAID 5

- A. Write data transactions take a little longer than read data operations to complete
- B. Even when the damaged drive is being repaired and the system's functions refurbishes the data on the replacement drive, if a drive fails, may still access all of your data.

Drawbacks of RAID 5

- 1) Throughput is impacted by drive failures, however this is still adequate.
- 2) This technology is sophisticated. Depends on the load on the arrays and the performance of the controller, rebuilding the data when a 4TB disc fails and has to be replaced might take a day or more. Data are lost permanently if another disc fails at that time.

RAID level 6 – Striping with double parity:

Similar to RAID 5, RAID 6 writes the parity data to two discs. That implies it needs at least 4 drives and is capable of handling 2 drives failing at once. Of course, there are extremely little odds that two drives will fail simultaneously. It takes hours or even more a day to reconstruct the swapped drive in a RAID 5 system, however, if a drive dies and is replaced with a new drive. During that time, if another disc fails, still lose all of information. Even that second failure will not destroy the RAID array when using RAID 6. A strong all-around solution, RAID 6 combines effective storage with top-notch security and passable speed. In file and applications servers that employ several big discs for data storage, it is preferred over RAID 5.




Figure 3.5: RAID 6

Benefits of RAID

- 1) Read data operations happen extremely quickly, similar to RAID 5.
- 2) Even when the failed drives are being replaced, if two drives fail, you may still access all of your data. Therefore, RAID 6 is safer than RAID 5.

Drawbacks of RAID 6

- 1) Due to the extra parity data that must be computed, write data transactions take longer than RAID 5 operations. One study I saw showed a 20% decrease in writing performance.
- 2) Throughput is impacted by drive failures, however this is still acceptable.

- 3) This technology is sophisticated. It might take a while to rebuild an array once a drive fails.

Combining RAID 1 with RAID 0, or RAID level 10

The benefits (and drawbacks) of RAID 0 and RAID 1 may be combined into a single system. This RAID setup is layered or hybrid. By mirrored all data on secondary drives and applying striping across each pair of discs to speed up data transfers, it improves security.




Figure 3.6: RAID 10

Benefits of RAID 10

A RAID 10 configurations rebuild time is quite quick if one of the discs has a problem since all that is required is moving the data from the remaining mirror to a new drive. For drives with 1 TB of storage, this might finish in as little as 30 minutes.

Drawbacks of RAID 10: When compared to huge RAID 5 or RAID 6 arrays, mirroring uses half the storage space, making it more costly to provide redundancy in this method.

Performance as a Capacity Factor: While creating RAID performance formulas in terms of the amount of spindles, which is quite logical. This enables us to analyse the relative performance of several suggested choices, which is highly helpful in figuring out the performance of a planned array or even an existent one when measurement is not available. The performance of RAID in these words. This is not usually a wise strategy, however, since we often consider RAID as a generating capacity rather than a speed or spindle count factor. Someone choosing an eight-drive RAID 6 array over an eight-drive RAID 10 array would be very uncommon, but it is feasible. This may sometimes happen because of a chassis restriction or another related issue. But more often than not, rather than considering spindle count, speed, or any other element, we look at RAID arrays from the perspective of total array capacity. So it seems strange that we should start considering spindle count as a factor in RAID performance. If we shift our perspective and focus on capacity rather than performance (X), while maintaining the assumption that particular drive capacity and performance are consistent among comparators, we arrive at a very different performance environment. By doing this, we may see, for instance, that read performance no longer consistently fluctuates and RAID 0 is no longer the most standards compliant RAID level.

Although capacity may be unpredictable, we can boil it down to the number of spindles required to meet the specified capacity. This makes this conversation much simpler. Therefore, the first step is to calculate the number of spindles required for raw capacity. Ten spindles, for instance, would be required if we needed a 10TB capacity and were utilizing 1TB drives. Alternatively, six spindles would be required if we needed 3.2TB and were utilizing 600GB drives. Now refer to our thread count as "R." (R is used to indicate that this is the Raw Capacity Count instead of the total number of spindles; see below.) As previously, "X" is used to indicate how well each drive is doing. RAID 0 is still very basic. There are no extra drives, therefore performance is still RX. IOPS for read and write operations are both NX. RAID 10 has two more read IOPS than writes. Dramatic, isn't it? Suddenly, we discover that RAID 10 has double the read performance of RAID 0 when performance is seen as a factor of stable capacity. RAID 5 becomes a little bit difficult. The formula for write IOPS is $(R + 1) * X / 4$. $(R + 1) * X$ is the formula used to express read IOPS. As anticipated, RAID 6 exhibits the same pattern as RAID 5. For RAID 6, the write IOPS are $(R + 2) * X / 6$. And $(R + 2) * X$ is the formula for the Read IOPS. By just considering read performance, RAID 0 becomes the slowest RAID level rather than the quickest, and RAID 10 becomes the fastest for both read and write, regardless of the values for R and X. This new perspective alters the way we think about performance.

RAID Controller

Hard disc drives in a shared storage are managed by a RAID controller. By displaying collections of discs as logical units, it may serve as a layer of abstraction between the operating system and actual discs. Performance may be enhanced and data can be protected in case of a crash by using a RAID controller. A RAID controller might be based on software or hardware. In a hardware-based RAID device, the whole array is controlled by a physical controller. Along with supporting drive formats like Advanced Technology Attachment Serial and Small Computer System Interface, the controller may also be built to support those. The motherboard of a server may also have a hardware RAID controller. When using software-based RAID, the interacting as a system advantage of the physical system's memory and central CPU. Although a software-based RAID controller performs the same tasks as a hardware-based RAID controller, it may not provide as much of a speed gain and may have an impact on the performance of other server-based applications. Firmware, or driver-based RAID, is a possible solution if a software-based RAID implementation is incompatible with a system's boot-up procedure and hardware-based RAID controllers are too expensive. Similar to software-based RAID, firmware-based RAID is controlled by chips on the motherboards, and all operations are handled by the CPU. The RAID mechanism is only used with firmware, however, and only when the device boots up. The driver takes over RAID operations after the OS has loaded. Although a software RAID controller is more CPU-intensive than a hardware solution, it is less expensive. Device software RAID, composite model RAID, and fake RAID are other names for firmware-based RAID.

RAID implementation: Both computer hardware and software are capable of managing the data distribution over several drives.

Hardware-based: A separate controller has to be placed in the server for hardware-based RAID. Even before operating system is run, hardware RAID controllers may be set up using the card BIOS or Option ROM. Additionally, each controller's vendor offers unique setup tools once the operating system has booted. Separate hardware is used to generate hardware RAID. There are two possibilities. A cheap RAID chip could be integrated onto the motherboard. Option that is

more costly and has a sophisticated standalone RAID controller. These controllers often support hot swapping and may be outfitted with their CPU and battery-backed cache memory. A hardware-based RAID card manages the RAID array(s) entirely, giving the system logical drives without adding any additional load to the system itself. Hardware RAID may also provide the system with a wide variety of RAID configurations at once. For the huge storage array, a RAID-5 array is also provided in addition to a RAID 1 array for the boot and application drive. Other operating systems may have their own built-in generic frameworks for interacting with any RAID controller and tools for keeping track of the state of RAID volumes. There are various benefits of a hardware RAID over a software RAID, including:

- 1) It doesn't use the host computer's CPU.
- 2) Users are able to create boot sectors.
- 3) Due to direct communication with the devices, it manages faults better.
- 4) It allows for hot-swapping.

Software RAID

All of Stead fast's local servers come with the software RAID option as a standard feature. This indicates that software RAID 1 is FREE and is strongly suggested if you're utilising local storage on a machine. It is strongly advised that all of the discs in a RAID array be the same kind and size. One of the least expensive RAID technologies is software RAID. Software-based RAID will manage the RAID setup using a portion of the system's processing power. When utilizing ordinary HDDs and trying to improve system performance, such as with RAID 5 or 6, it is advisable to utilize a hardware-based RAID card. Software RAID implementations are available in many contemporary operating systems. Software RAID may be put into practice as:

- 1) A layer that unifies various hardware into a single virtual computer.
- 2) A layer that lies over any file system and shields user data from parity attacks.

The system must be clever enough to boot from the remaining drive or drives if a boot drive fails. The software RAID cannot be used to boot the system without certain restrictions. Boot partitions can only be included in RAID 1, and RAID 5 and RAID 0 software cannot allow for system boot.

RAID Works

RAID increases performance by distributing data across numerous discs and enabling input/output processes to overlap in a balanced way. Keeping data redundantly increases fault tolerance because using different drives lengthens the mean time among failures (MTBF). The operating system sees RAID arrays as a single logical drive. Disk striping or disc mirroring are methods used in RAID.

- 1) Disk mirroring will duplicate data across many drives.
- 2) Disk striping partitions aid in distributing data across several disc drives. The storage capacity of each disc is divided into blocks that range in size from 512 bytes to several megabytes. All of the discs' stripes are interspersed and addressed sequentially.
- 3) A RAID array can also combine disc striping and mirroring.

The stripes are normally set up to be modest (512 bytes) in a single-user system where substantial records are stored so that a single record covers all the discs and can be accessed rapidly by reading all the discs at once. Better performance in a multi-user system necessitates a stripe broad enough

to accommodate the typical or maximum size record, enabling overlapping disc I/O between drives.

Utilization of RAID

When a significant quantity of data has to be recovered, having a RAID arrangement might be helpful. Since the data is also kept on other devices, it may be swiftly recovered if a disc malfunctions and the data is lost. When availability and uptime are crucial business issues. If data restoration is required, it may be done swiftly and without a break in service. While handling huge files. Working with massive files requires dependability and speed, which RAID offers. When a company wants to improve overall performance while lessening the load on its hardware. A hardware RAID card could have extra memory on it that can be utilized as a cache, as an example. When experiencing disc I/O problems. RAID will increase performance by reading and writing data from several drives rather than delaying operations until one disc is ready. When price is a concern. A RAID array is now less expensive than it was in the past because to the widespread usage of inexpensive drives.

RAID data recovery: the possible difficulties

- A. Because availability and efficiency are crucial for business-critical tasks inside enterprise IT infrastructures, RAID arrays are extremely sophisticated, and the difficulties they pose are heightened.
- B. Although they can be helpful in some situations, add-on technologies like virtualization and database applications can make an already troubled system even more expensive to fix.
- C. From the standpoint of data recovery, it is typically necessary to rebuild the RAID file system, get around any physical issues, and evaluate any virtualized design. Although recovering from this can be a difficult and drawn-out process, it is doable with the appropriate knowledge.

Advantages of RAID

These are some advantages of RAID.

- A. A rise in cost-effectiveness due to the widespread use of discs at decreasing prices.
- B. RAID can increase the performance of a single hard drive by using numerous hard drives.
- C. Depending on the settings, increased computer speed and dependability following a crash.
- D. RAID 5 offers greater availability and resilience. RAID arrays with mirroring can have two discs with identical data. It guarantees that if one continues to work.

E. Disadvantages of RAID

The following are negative aspects or shortcomings of RAID:

- A. Nested RAID levels require more discs than conventional RAID levels, hence they are more expensive to implement.
- B. Because many of the drives are used for redundancy in nested RAID, the cost per gigabyte of storage devices is higher.
- C. When a drive in the array fails, the likelihood that another drive will soon follow suit rises, increasing the likelihood that data will be lost. This is due to the simultaneous installation of all the discs in a RAID array. As a result, the wear on all drives is the same.

- D. RAID levels like 1 and 5 can only withstand the failure of one drive.
- E. Until a failing drive is replaced and the new disc is filled with data, RAID arrays are exposed.
- F. Since RAID was introduced, rebuilding failing discs requires a lot more time due to the larger size of drives.

Hot Spares: A slice, not a volume, that is available and operational but not being used is called a hot spare. A hot spare is held in a sub mirror or RAID 5 volume when it is available to replace a failed slice. Hot spares provide hardware failure protection because when a slice from a RAID 1 or RAID 5 volume fails, it is instantly replaced and resynchronized. In the meanwhile, a failing sub mirror or RAID 5 volume slice may be replaced or rectified by using the hot spare. Hot spare pools include hot spares that you generate.

A hot spare pool may include one or more individual hot spares. For instance, you may have two hot spares in addition to two sub mirrors. Two hot spare pools may be created using the hot spares, with each pool containing the two spares in a differently preferred order. This tactic successful system implementation by having more hot replacements accessible and allows you to select which hot spare is utilized first. Only a hot spare whose length is equal to or larger than the size of the damaged slice in the sub mirror or RAID 5 volume may be used.

A hot spare for a sub mirror that uses 1 Gbyte discs, for instance, must also be 1 Gbyte or larger. The array's spare disc is known as a hot spare RAID. It is only required at a certain period and is not utilized at all. In the case that one of the array drives fails, this time will occur. The RAID controller initiates the RAID rebuilding process and switches out the faulty disc with a hot spare disc. As a result, there is a substantially shorter length of time during which the array information is unavailable. In RAID 1, RAID 5, or RAID 6, hot spare discs are possible. Any failing disc in the storage array may be replaced with a hot spare since it is not reserved for a particular volume group. The disc that it safeguards must be similar to the hot spare disc, among other requirements. Since the simultaneous destruction or failure of two discs rarely happens under unusual circumstances, the array achieves maximum data protection. Hot spare makes it simple to repair a single drive failure.

A hot spare is a spare HDD in a RAID array that momentarily takes the place of a RAID set's failing HDD. The array's failed HDD's identification is transferred to a hot spare. Depending on how RAID is implemented, one of the following data recovery techniques is used: If integrity RAID is used, the data is reconstructed onto the hot spare using the parity and the data from the other HDDs in the RAID set. If mirroring is used, the data is copied using the data from the remaining mirror. One of the following things happens when a fresh HDD is used to replace a failing HDD: The hot spare completely replaces the new HDD. As a result, it can no longer serve as a hot spare, and the array will need to setup a new hot spare. Data from the hot spare is transferred to the new HDD when it is inserted into the system. When the next failed drive fails, the hot spare resumes its idle position.

Data from a failing drive should fit in a hot spare's capacity. Multiple hot spares are used by certain systems to increase data availability. The configuration of a hot spare determines whether it will be utilized automatically or manually in the case of a disc failure. In an automated setup, the disc subsystem attempts to automatically transfer data from the failed disc to the hot spare when the recovered error rates for a disc surpass a set threshold. The subsystem changes to the hot spare and designates the failing disc as useless if this action is finished before the damaged disc fails. If not,

it recovers the data using parity or a mirrored drive. In the event of a user-initiated setting, the rebuild procedure is within the administrator's control. For instance, the rebuild might take place at night to avoid any system performance decrease. If a hot spare is not available, the system is susceptible to another failure.

CHAPTER 4

INTELLIGENT STORAGE SYSTEM

Dr. Uthama Kumar A, Assistant Professor

Department of Data Science & Analytics, School of Sciences, Jain (Deemed-to-be University), Bangalore-27, India

Email Id- uthamakumar.a@jainuniversity.ac.in

Intelligent Storage Arrays or Intelligent Storage Systems are terms used to describe feature-rich RAID arrays with highly efficient I/O processing capabilities. The needs of today's I/O-intensive next-generation applications may be satisfied by these intelligent storage solutions. High levels of performance, availability, security, and scalability are necessary for these applications. As a result, many manufacturers of intelligent storage systems now provide SSDs, encryption, compression, deduplication, and scale-out design in order to satisfy application needs. Massive numbers of IOPS can be handled because to the utilization of SSDs and scale-out architecture. These storage systems provide connection to several types of computing systems. Additionally, APIs are supported by the intelligent storage systems to facilitate interaction with cloud and Software-Defined Data Center (SDDC) environments.

Intelligent Storage Systems: These storage systems feature an operating system that manages the provisioning, management, and consumption of storage resources intelligently and effectively. To fulfil the needs of performance-sensitive programs, the storage systems are set up with a large quantity of memory called cache, several I/O pathways, and complex algorithms. The controller and storage are the two main parts of an intelligent storage system. A controller is a computing device that runs an operating system designed specifically to handle a number of crucial tasks for the storage device. Serving I/O requests from the application servers, managing storage, protecting RAID arrays, enabling remote and local replicating, providing storage, automating tiering, encrypting and compressing data, and managing caches intelligently are a few examples of these tasks.

For redundancy, an intelligent storage system often contains many controllers. To handle a huge number of I/O requests, each controller has one or more computers and a certain quantity of cache memory. These controllers are either directly or indirectly linked to the servers via a storage network. The servers send I/O requests to the controllers, and the controller reads or writes data from or to the storage on their behalf. A storage system may be categorised as a block-based storage system, file-based storage system, object-based storage system, or unified storage system depending on the kind of data access. Block-based, file-based, and object-based data access are all available in a single system with a unified storage system. The next posts detail them.

Intelligent Storage System Architecture

Either a scale-up architecture or a scale-out architecture may be used to build an intelligent storage system. According to the needs, a single storage system's capacity and performance may be scaled up using a scale-up storage architecture. Upgrading or adding controllers and storage are necessary for scaling up a storage system. These systems' fixed capacity limitation restricts their potential to scale, and performance also begins to suffer when the capacity limit is reached. By simply adding nodes to the cluster, a scale-out storage design offers the potential to maximize its capacity. When

the cluster needs more performance or capacity, nodes can be added quickly and without any downtime. This gives the option to combine several nodes with average performance and availability to create a system with greater aggregate performance and availability. The workload is split across all the nodes in a scale-out design, which pools the cluster's resources. This causes performance to increase linearly when nodes are added to the cluster.

Characteristics of Intelligent Storage Systems

Storing In Tiers: A method of creating a hierarchy of several storage types is called storage tiering (tiers). This makes it possible to efficiently store the appropriate data at the appropriate tier depending on service level needs. Different tiers provide varying degrees of performance, affordability, and protection. Using this method, we are able to store data on the best tier possible. It is beneficial to store active data on quick media and inactive information on slow media. By not having to load the storage array with fast discs when the majority of the data is only sometimes accessed, this may improve the efficiency of the shared storage and save expenses. Data is moved in accordance with established tiering procedures. The tiering approach may be based on variables like access frequency. For instance, tier 1 storage may be set up with high speed SSDs or FC drives to hold data that is accessed often, while tier 2 storage can be set up with inexpensive SATA drives to keep data that is accessed less frequently. Application performance is enhanced by storing frequently accessed information on SSD or FC. Data that isn't accessed as often may be moved to SATA to free up space on high-performance drives and cut the cost of storage.

Data is often moved automatically from one sort of layer to another. Automated storage tiering automatically moves active data to a higher performance tier and idle data to a greater capacity, reduced performance tier while continuously monitoring the application demand. Data is moved across the levels without causing any disruptions.

Redundancy: Intelligent storage systems with redundancy features make sure that malfunctioning parts do not stop the array from operating. Multiple pathways are often set between host and storage in multi - path I/O configurations, even at the host level, to ensure that the system won't go down in the event of a path or network connection failure between the host and storage array.

Replication: Storage system-based replication creates distant copies of production volumes that might be crucial in planning for business continuity and disaster recovery (DR). Remote replicas may either be zero-loss synchronized replicas or asynchronous replicas, depending on the application and business requirements. While synchronous replication needs a distance of no more than 100 miles between the source and destination volumes, asynchronous replication solutions allow for distances of thousands of miles.

Tight Provisioning: The capacity of the storage systems may be utilized more efficiently by using thin provisioning solutions. Running out of storage capacity due to excessive provisioning might ultimately occur.

In a hybrid cloud architecture, intelligent storage

The use of intelligence to storage management to optimize for workloads with apps and across hybrid cloud architecture may be its finest usage. Databases are primarily reliant on applications to handle structured data. Data-aware devices may leverage information about application-level processes to enhance the databases' storage performance. For instance, a data-aware device may forecast which data will be required in the future using data supplied by a database management

system. However, it's uncommon to utilize databases and apps separately. Rather, they are often a component of bigger operations.

Advantages of intelligent storage and DevOps

The ability to monitor operations and optimize storage across various storage devices utilizing data gathered from computing, networking, and storage devices is the highest degree of storage intelligence. This kind of intelligence may aid in reducing cloud storage expenses, which are sometimes difficult to monitor and frequently result in unforeseen expenditures. Hardware upgrades are no longer necessary to increase storage performance and cost-effectiveness. The most significant developments are currently being driven by intelligent software that covers everything from low-level I/O operations on a single device to workload monitoring and data placement optimization across hybrid cloud architecture.

Four essential parts make up an intelligent storage system: the front end, the cache, the back end, and the actual discs. These elements and their relationships are shown in Figure. To allow data storage and retrieval from the actual disc, an I/O request received from the host at the front-end port is handled via cache and the back end. If the desired data is located in cache, a read request may be fulfilled immediately.




Figure 4.1: Intelligent Storage System Architecture

Front End

The front end acts as the host's interface with the storage system. Front-end ports and front-end controllers make up its two parts. Hosts may connect to the intelligent storage system through the front-end ports. For storage connections, each front-end port includes processing logic that performs the proper transport protocol, such as SCSI, Fibre Channel, or iSCSI. On the front end, redundant ports are offered for high availability. The internal data bus is used by front-end controllers to transport data to and from cache. The controller notifies the host with an acknowledgement message when the cache receives write data. Controllers use command queuing methods to streamline I/O processing Commands.

Queuing Front-End: Front-end controllers use a method called command queuing. It chooses the sequence in which orders are executed after being received, which may cut down on pointless

drive head motions and enhance disc performance. The command queuing methods assign a tag that specifies the order in which commands should be performed when a command is received for execution. Regardless of the sequence in which the orders were received, command queuing enables several commands to be performed simultaneously depending on the arrangement of data on the disc. The following are the command queuing methods that are most often used:

First In First Out (FIFO): First In First Out (FIFO) is the standard algorithm, which executes orders in the order they are received. Requests for optimization are not reordered; as a result, performance is ineffective.

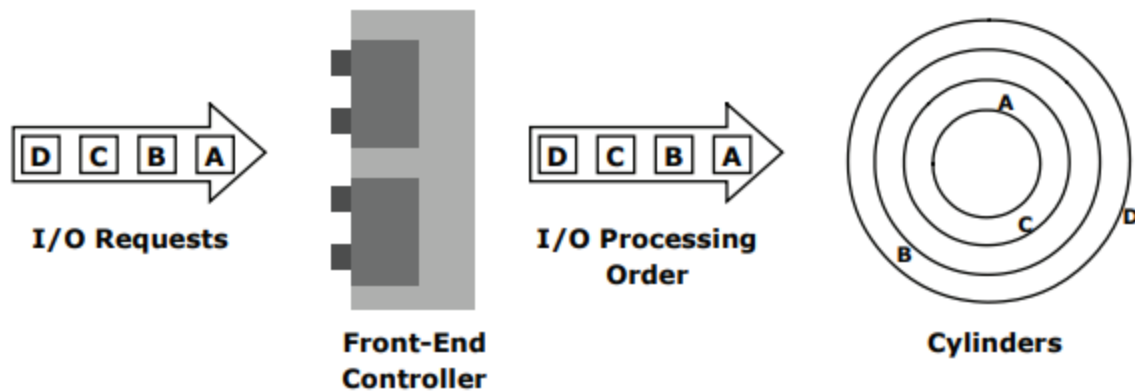


Figure 4.2: Without Optimization in Disk command queuing

Seek Time Optimization: Instructions are performed using read/write head movement optimization, which may require rearranging the commands. The instructions are carried out in the order they are received if search time optimization is not used. For instance, the orders are carried out in the sequence A, B, C, and D as illustrated in Figure. When C comes after A, there is a smaller radial movement needed by the head than there would be for B. Figure illustrates the command execution process with seek time optimization as A, C, B, and D.

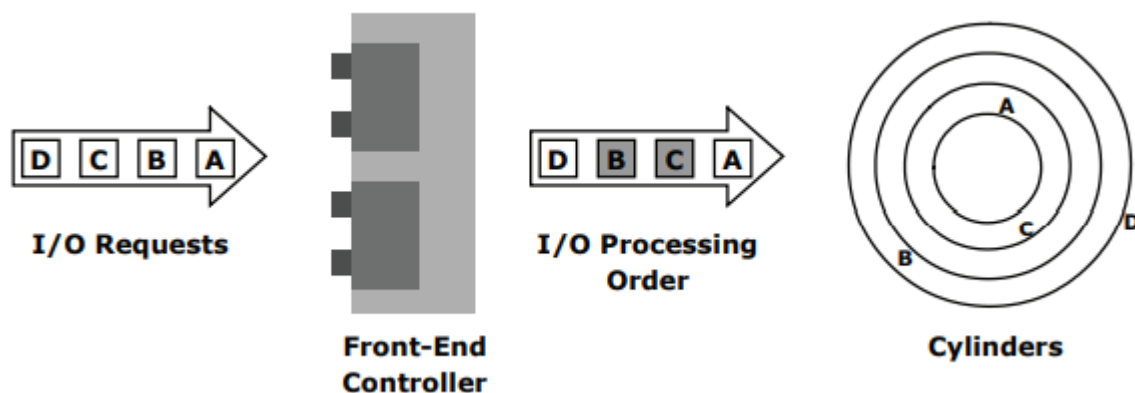


Figure 4.3: With Seek Time Optimization in Disk command queuing

Access Time Improvement: For optimum performance, commands are carried out using a combination of search time optimization and a rotational latency analysis. Additionally to the

command queuing performed on the front-end controllers, command queuing may also be done on disc controllers. Command queuing is a feature that certain SCSI and Fibre Channel drive types' controllers have adopted.

Cache: In an intelligent storage system, cache is a key element that improves I/O performance. To speed up the processing of I/O requests from the host, data is temporarily stored in cache, a kind of semiconductor memory. By insulating hosts from the mechanical delays associated with physical discs, which are the slowest parts of an intelligent storage system, cache increases the speed of storage systems. Due to search delays and rotational latency, accessing data from a physical disc often takes a few milliseconds. Requests are queued and responses are delayed if the host must access the disc for each I/O operation. Less than one millisecond is needed to access data that is in cache. After being cached, write data is then written to disc. The host is promptly recognized when the data is safely stored in cache.




Figure 4.4: Structure of Cache

Organization of Cache

The lowest unit of cache allocation, pages or slots, make up the structure of the cache. A cache page's size is set in accordance with the application's I/O size. The data store and tag RAM make up the cache. The data store keeps the data while tag RAM keeps track of where it is on disc and in the data store (see Figure in above figure).

Data discovered in cache and its proper location on the disc are both indicated by entries in tag RAM. A dirty bit flag in tag RAM lets you know whether or not the data in cache has been committed to the disc. It also includes time-based data, such as the time of the most recent access, which is used to locate cached data that hasn't been visited in a while and may be released.

Operation with Cache may be read:

The front-end controller checks the tag RAM when a host requests a read to see whether the requested data is already in cache. A read cache hit, also known as a read hit, occurs when the requested data is discovered in the cache and is given to the host immediately without the need for a disc operation (see Figure).

This gives the host a quick reaction time (about a millisecond). A cache miss occurs when the requested data cannot be located in cache and must instead be retrieved from the disc (see Figure). The required data is retrieved by the back-end controller by accessing the proper disc. The front-end controller then sends the data to the host once it has been cached. I/O response time increases due to cache misses.

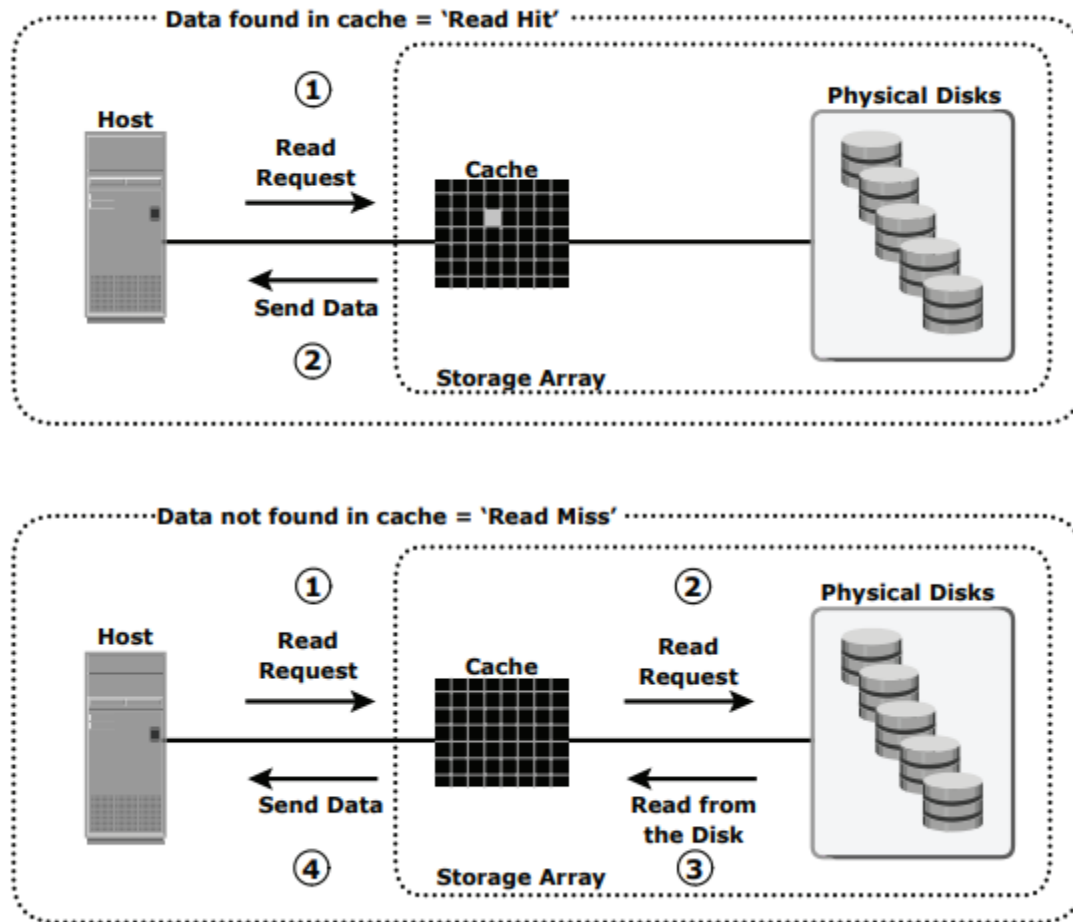


Figure 4.5: Read hit and read miss

When read requests are sequential, a pre-fetch, or read-ahead, method is used. An associated collection of contiguous blocks is returned in response to a sequential read request. Other blocks that the host hasn't yet asked for may be read from the disc and put into cache beforehand. The read operations are read hits when the host later requests these blocks. The host's reaction time is markedly improved by this procedure. Fixed and flexible pre-fetch sizes are offered by the intelligent storage system.

The intelligent storage system pre-fetches a predetermined quantity of data during fixed pre-fetch. It works best when the I/O sizes are consistent. The storage system pre-fetches a quantity of data that is multiples of the size of the host request when using variable pre-fetch. In order to avoid the discs from being overloaded with pre-fetch at the cost of other I/O, maximum pre-fetch sets a limit on the amount of data blocks that may be pre-fetched. Reading efficiency is determined by the

read hit ratio, also known as the hit rate, which is often given as a percentage. The number of read hits relative to the total number of read requests is represented by this ratio. The read performance is enhanced by a better read hit ratio.

Operation Write with Cache: Over writing to discs directly, write operations with cache provide better speed. From the host's viewpoint, an I/O is finished much faster when it is written to cache and acknowledged than it would be if it were written straight to disc. Sequential writes also provide chances for optimization since, with the help of cache, several smaller writes may be combined for larger transfers to disc drives. The following methods are used to accomplish a write operation using cache:

Write-back caching: When data is added to the cache, an instantaneous acknowledgement is delivered to the host. After multiple writes, data is later committed (de-staged) to the disc. Because write operations are not affected by the disk's mechanical delays, write response times are substantially quicker. However, in the case of cache failures, uncommitted data runs the risk of being lost.

Write-through caching: When data is added to the cache, it is instantly written to the disc, and the host receives an acknowledgement. Data loss risks are minimal since data is written to disc as it comes in, but write response times are greater because of disc operations. Under specific circumstances, such as extremely high size write I/O, the cache may be bypassed. In this implementation, writes are sent straight to the disc if the size of an I/O request surpasses the predefined amount, known as the write aside size, in order to lessen the effect of heavy writes occupying a large cache space. When cache resources are limited and need to be made accessible for little random I/O's, this is very helpful.

Implementation of Caches

Both dedicated cache and global cache may be used to implement caching. With a dedicated cache, distinct sets of memory space are set aside for reads and writes, respectively. Any memory address is usable for reads and writes in the global cache. One global set of addresses must be controlled, making cache management more effective in a global cache solution. Users may be able to select the percentages of cache that are accessible for reads and writes in cache management using global cache. The read cache is typically minimal, but if the program being used requires a lot of reading, it should be expanded. Other global cache solutions dynamically change the amount of cache that is available for reads vs writes depending on the workloads.

Management of Caches: The maintenance of cache, a limited and costly resource, is crucial. Intelligent storage systems may be designed with a lot of cache, but after all the pages are used up, some pages must be released to make room for new data and maintain performance. In intelligent storage systems, a variety of cache management techniques are used to proactively maintain a set of free webpages and a list of pages that could be able to be freed up at any time: The method known as Least Recently Used (LRU) constantly tracks data access in cache and finds cache pages that haven't been visited in a while. These pages are either made available for reuse or freed by LRU. This technique is predicated on the idea that a host won't request data that hasn't been accessed recently. However, data will initially be written to disc before the page is reused if it includes write data that hasn't yet been committed to disc.

Most Recently Used (MRU): The algorithm most recently used (MRU) is the opposite of LRU. The most recently visited pages in MRU are released or designated for reuse. This technique is predicated on the idea that recently accessed data may not be needed right away.

To maintain cache availability as it fills, the storage system must act to flush dirty pages (material stored into the cache but not yet written to the disc). Data from the cache is committed to the disc via the process of flushing. High and low levels known as watermarks are put in cache to regulate the flushing process based on the I/O access rate and pattern. The storage system begins high-speed flushing of cache data at the high watermark (HWM), which is the level of cache usage. When the storage system reaches the low watermark (LWM), it stops performing forced or high-speed flushes and switches back to idle flushing. The manner of flushing to be employed depends on the amount of cache use.

When the cache usage level is halfway between the high and low watermarks, idle flushing takes place continually at a moderate pace. When cache use reaches the high watermark, the high watermark flushing feature is activated. Flushing is given some more resources by the storage system. The host I/O processing is not significantly impacted by this kind of flushing. When the cache fills to 100% of its capacity during a heavy I/O burst, forced flushing takes place, which has a considerable impact on the I/O response time. Dirty pages are forcefully flushed to disc during forced flushing.

Data Cache Protection

Data that hasn't yet been committed to the disc will be lost in the event of a power outage or any other kind of cache failure since cache is volatile memory. Cache mirroring and cache vaulting may reduce the risk of losing uncommitted data stored in cache:

Mirroring cache: Each write to cache is stored on two separate memory cards in two distinct memory locations. The write data will be secure at the mirrored location in the case of a cache failure and may be committed to the disc. The data may still be retrieved from the disc in the case of a cache failure since reads are staged from the disc to the cache. Because only writes are replicated, this approach maximizes the use of the available cache. The issue of preserving cache coherency is encountered with cache mirroring methods. Data must always be same across two separate cache locations in order for there to be cache coherency. The array operational environment is in charge of ensuring coherency.

Cache vaulting: In the event of a power outage, cache is vulnerable to the possibility of uncommitted data loss. There are many methods to solve this issue, including employing battery power to write the cache's contents to disc while the memory is powered by a battery until the AC power is restored. Using batteries is not a realistic solution in the case of a prolonged power outage because intelligent storage systems may need enormous volumes of data to be committed to several discs, and batteries may not supply power for long enough to write each piece of data to its appropriate disc. As a result, storage providers utilize a collection of actual discs to dump the cache's contents in the event of a power outage. The practice is known as cache vaulting, and the drives are known as vault drives.

Data from these drives is written back to the write cache and then transferred to the proper discs when the power is restored.

Back End: The back end acts as an interface between the actual discs and the cache. Back-end ports and back-end controllers make up its two parts. Data transfers between the cache and the real discs are managed by the back end. Data is transported from cache to the back end and then directed to the final disc. On the back end, ports are linked to physical discs. When reading and writing to the discs, the back end controller interacts with them. It also offers a little amount of extra temporary data storage. The algorithms used by back-end controllers provide RAID capability as well as error detection and repair.

Storage systems are set up with twin controllers and numerous ports for the highest levels of data security and availability. In the case of a controller or port failure, such arrangements provide a backup route to physical drives. If the discs are additionally dual-ported, this dependability is increased even more. Then, each disc port may link to a different controller. Also made possible by several controllers is load balancing.

Hard Drive: Data is constantly stored on a physical disc. SCSI or a Fiber Channel interface is used to link discs to the back-end. Utilizing a combination of SCSI or Fiber Channel discs and IDE/ATA drives is possible with an intelligent storage solution. Logical Unit Numbers, or LUNs, are one of the most important aspects of storage devices. You could sometimes hear people discussing it or read about it online. The SCSI protocol uses this number to identify the logical units within a storage device.

Logical Unit Number

The region on which data is written in a storage device is determined by its logical units. To address each logical unit, the SCSI protocol needs a string of integers. Therefore, storage uses Storage Area Network, or SAN, protocols to encapsulate the SCSI while it is in use. As a result, the logical unit in use is addressed by a number, and SCSI uses them by their numbers, coining the name LUN or logical unit number. This numbered item may be applied in a variety of circumstances, including the defunct tape drive. Or practically anything that writes data to a real drive. In the Storage Area Network, it is a number that is used to identify a logical disc (SAN). The logical disc itself may sometimes also be referred to as a LUN. However, the majority of specialists think it is technically incorrect. The storage device is made up of a number of different units, each of which is identified by a unique serial number that we refer to as a logical unit number.

The Process of Logical Unit Number

Some find it challenging to comprehend how LUNs operate in storage devices. A hard disc drive, however, writes data to several actual discs. However, a serial number is assigned to every component (physical disc or memory chip). When necessary, SCSI targets this serial number. Therefore, the storage is divided into several volumes when the user formats the disc array or creates a partition. SCSI then configures a logical unit and assigns it a number in order to identify each volume. Therefore, the computer accesses the LUN given by SCSI whenever a straightforward activity, such as unlocking the disc in your PC, happens. Therefore, the CDB (Command Descriptor Block) is transferred to a physical storage unit when such actions are started by a computer administrator. The logical unit within the particular target is then identified by a 3-bit logical number unit. This enables the gadget to start up in accordance with the number and keep writing to that particular storage region. SCSI transmits CDB to a different physical storage unit when that space is full or is otherwise utilized. Then, until it too runs out, this logical number unit takes precedence.

Logical Unit Number Types LUN, or a logical unit number, exists in a variety of sorts even though it may seem like one single form. Each kind is intended for a certain storage device activity. As a result, the number types also vary when other logical unit types are needed. Here are the top four of these kinds:

Masked LUN: A fault-tolerant LUN type called a "mirrored LUN" creates two data copies on different physical discs. It enables better management of backup and redundancy. Programs that detect redundant or duplicate files also use this type.

LUN concatenated: The main function of a concatenated LUN is to combine many logical number units into one volume. When two logical number units are barred together, it is the most common kind of logical number unit used in RAID groups.

Patterned LUN: A single logical unit number is used in striped LUNs to write data across many physical devices. As a result, the computer's performance and the different physical discs' longevity are improved. Simply because it splits the input and output amongst many drives.

LUN with stripes and parity: Striped LUN and striped LUN with parity are essentially equivalent. This Logical Unit Number type's function is to parity data across physical discs, ideally three or more. This makes it possible to reuse data from other drives in the event that one disc malfunctions or fails.

Use Cases for Logical Unit Numbers: LUN is mostly used to identify storage devices. Although there may be variations based on the LUN type, the fundamentals are the same. For instance, if you partition a disc, the partition will be identified on the physical drive by a straightforward Logical Unit Number. In addition, LUNs are used to zone and mask SANs so they may virtualize each physical sector to map multiple physical LUNS. Additionally, as was already said, the word itself is employed to distinguish a logical unit as well as to identify the number.

LUN Uses

LUNs are primarily used to designate storage devices as an identifier. Each LUN type may, however, be used in a different way. An example would be the usage of a straightforward LUN to designate a portion or the full physical disc. A LUN that spans two or more physical discs is designated as a spanned LUN by the designator "spanned LUN". In the event that one disc fails, the mirrored LUN is used to command the copying of the data stored on the first disc to the second disc. LUNs may be virtualized to map numerous physical LUNs, or they can be used for zoning and masking in the SANs.

LUN Masking and Zoning: To guarantee predictable behavior, LUN zoning offers segregated pathways for I/O across the FC SAN hierarchy between end ports. The host is restricted to the zone in which it was assigned. The switch layer is often where LUN zoning is established. It may improve security and get rid of network hotspots. Access to certain SCSI targets and associated LUNs is restricted by LUN masking for hosts. LUN masking is often performed on the storage controller, although it may also be done on the switch layer or host bus adapter (HBA). Multiple hosts and zones may use the same storage device port thanks to LUN masking. They can only see certain SCSI targets and allotted LUNs, however.

Virtualization and LUNs: The LUN resembles virtualization in that it leverages the standard SCSI mechanism of identification and communication to hide the underlying hardware devices.

As long as the host's representation doesn't change, the storage item that the LUN represents may be set, compressed, or reduplicated. Within and across storage devices, LUNs may be moved, copied, replicated, snapshotted, and tiered. Establish a virtual LUN, which can be generated outside of the available physical space, to map to numerous physical LUNs and virtualize capacity. Due to the fact that physical storage is not reserved before data is written, creating a virtual LUN that is bigger than the actual capacity available may aid in optimizing storage consumption. A thin LUN is another name for it. The server operating system (OS), hypervisor, or storage controller may all be used to configure virtual LUNs. LUN zoning is not necessary since the virtual machine (VM) cannot view the storage system's actual LUN.

Intelligent Storage Array: There are two main groups that intelligent storage systems often belong to:

- A. High-end storage systems
- B. Midrange storage systems

Midrange storage systems, which are commonly utilized in small- and medium-sized businesses, have historically been implemented using active-passive arrays whereas high-end storage systems have traditionally been built with active-active arrays. Active-passive arrays provide superior storage options at more affordable prices. Businesses take advantage of this financial advantage and deploy active-passive arrays to satisfy certain application needs including performance, availability, and scalability. The differences between these two implementations are becoming less and less important.

High-end storage systems:

Active-active arrays, a kind of high-end storage system, are often designed for big businesses to centralize corporate data. These arrays have several controllers and cache memory built into them. If an array is active, it means the host may access its LUNs by any of the various pathways in the figure below.




Figure 4.6: Active-active configuration

The following functionalities are offered by these arrays to meet the demands of businesses for storage: Large quantities of cache are needed to provide excellent host I/O service. Large storage capacity Data availability will be improved through fault tolerance design. Accessibility to hosts running open systems and mainframe computers o Availability of different front-end ports and

interface protocols to accommodate many hosts. Scalability to handle growing connection, performance, and storage capacity needs. Access to numerous back-end Fiber Channel or SCSI RAID controllers to manage disc processing. Capability to manage a high number of servers and applications sending significant quantities of concurrent I/Os Support for array-based local and distant replication. High-end arrays also include certain special capabilities and functionality that are necessary for mission-critical applications in big businesses.

Midrange Storage System: Small and medium-sized businesses would benefit most from midrange storage systems, often known as active-passive arrays. A host may only do I/Os to a LUN in an active-passive array via the pathways to the LUN's owning controller. These routes are known as active routes. Regarding this LUN, the other routes have a passive relationship. As can be seen in Figure 8, the host can only read from or write to the LUN via the route to controller A since that controller is the LUN's owner. There is no I/O activity done via the channel leading to controller B, which remains inactive. Midrange storage systems are often built with two controllers, each of which has disc drive interfaces, host interfaces, cache, and RAID controllers.




Figure 4.6: Active-Passive configuration

Midrange arrays have less storage space and a global cache than active-active arrays since they are meant to fulfil the needs of small and medium-sized businesses. Less front-end ports are available for connecting to servers. They do, however, provide good performance and redundancy for applications with predictable workloads. Additionally, they offer local and distant replication based on arrays.

CHAPTER 5

STORAGE NETWORKING TECHNOLOGIES

Dr. Thirukumaran Subbiramani, Assistant Professor

Department of Data Science & Analytics, School of Sciences, Jain (Deemed-to-be University), Bangalore-27, India

Email Id- s.thirukumaran@jainuniversity.ac.in

Direct-Attached Storage

A type of storage known as direct-attached storage (DAS) is connected to a computer directly, bypassing a network. The storage could have an internal or external connection. The only direct access to the data is through the host computer. To interact with the data, other devices must pass via the host computer. Internal hard disc drives (HDDs) or solid-state drives are typically found in servers, desktop computers, and laptops (SSD). These items are all examples of direct-attached storage. External DAS devices are also used by some PCs. An enterprise server may occasionally connect directly to share discs used by other servers. Direct-attached storage is not part of a network. As opposed to network-attached storage (NAS) or a storage area network, there are no connections through Ethernet or Fiber Channel (FC) switches (SAN). Through an interface like Small Computer System Interface (SCSI), Serial Advanced Technology Attachment (SATA), Serial-Attached SCSI (SAS), FC, or Internet SCSI, an external DAS device is directly connected to a computer (iSCSI). The device plugs into a card that is inserted into the computer's internal bus.

Benefits of DAS include:

- A. High availability,
- B. High access rate due to the lack of a storage area network (san), elimination of network configuration issues
- C. Growth of storage capacity
- D. Data security
- E. Fault tolerance.

DAS drawbacks include:

- A. Only one user may access the data at a time
- B. Different user groups cannot access the data.
- C. Exorbitant administrative fees.

Type of Direct Attached Storage (DAS)

Internal DAS and External DAS are the two categories under which Direct Attached Storage (DAS) is categorized. This classification is made based on where the storage device is located.

Internal DAS: Internal serial or parallel buses are used to connect the storage device to the server in an internal DAS system. A physical bus is utilized to provide high-speed connectivity over shorter distances, and this is also one of its drawbacks. In addition, because most internal buses can only accommodate a small number of devices and take up a lot of room inside the server, maintaining other components is difficult.

External DAS: The server is directly connected to the external storage device in an external DAS architecture. Most often, the Small Computer System Interface (SCSI) or Fiber Channel protocol is used to connect the server and storage device (FCP).

DAS use: Internal storage in the form of a hard disc drive (HDD) or solid-state drive (SSD) that is physically linked to the motherboard is the standard form of DAS used in servers and personal PCs. Direct-attached storage devices include things like external hard drives and USB (Universal Serial Bus) storage devices. Small and medium-sized enterprises (SMBs) and data centers can employ DAS as private storage linked to dedicated servers or as file servers. DAS is sometimes used in conjunction with networked storage systems like SAN and NAS by larger businesses. It is a realistic storage option for SMBs that want straightforward storage systems and don't need to distribute data throughout the firm. DAS is utilised when high performance and significant amounts of storage space are required. Comparing external DAS to other storage options like SAN and NAS, the cost of expansion is lower.

DAS function: DAS doesn't need a network connection to connect to the host computer or server and can be internal or external. A high-speed host bus adapter can be used to connect an internal storage device to a server or personal computer (HBA). Each personal computer has at least one internal DAS drive, which can be either a slower SSD or a more modern HDD connected via the SATA interface. Additionally, servers include internal storage that may be accessed through SATA or other quick interfaces like SCSI or SAS. External hard drives and disc enclosures that can hold numerous drives are examples of external DAS. These are linked to workstations and servers through USB, eSATA, SAS, or SCSI. The computer it is connected to manages and regulates DAS. Instead of directly accessing the data, computers on the network must talk with the computer to which the DAS is connected.

Use of DAS for storage: Direct-attached storage has a number of benefits, including:

- **High Performance:** Because DAS is connected to the computer requesting and consuming the data, it provides quick access to data. Data read/write access is quick and requests are unaffected by network traffic or connectivity issues because it is not connected to the network.
- **Simple Setup:** DAS is easy to install, configure, and use. Internal direct-attached storage is frequently purchased with computers and servers, making it immediately usable without configuration. Most external DAS devices are "plug and play," connected to a computer via a USB port, and can be utilized right away.
- **Low Cost:** Comparing to NAS and SAN, which need hardware and software to operate and administer the storage system, DAS is far more cost-effective. Only the disc drives and drive enclosures are required to be purchased.

Disk Drive Interfaces:

The five different kinds of hard disc interfaces are parallel ATA (also known as IDE or EIDE), SATA, SCSI, Fibre Channel, and SAS. IDE is mostly used in home goods. Additionally, part of it is used in web servers. The server market is where SCSI is mostly employed. While pricey and exclusively utilized in high-end servers, fiber channel. The most common kind of hard drive is presently SATA. It is used by the majority of desktop and laptop computers as well as solid state hard drives. Servers often utilize SAS. Both its transmission speed and durability are good. There are many different kinds of particular interfaces that fall under the broad categories of IDE and SCSI that may be separated based on various technical specifications and transmission speeds.

Definition of Integrated Drive Electronics

"Integrated Drive Electronics" is the full name of IDE, and its original meaning referred to a hard disc drive that integrated a "hard disc controller" and a "disc body." By using this method, the hard disc interface requires fewer and shorter cables, which improves data transfer reliability and streamlines the hard disc manufacturing process. In the past, many hard drives used IDE ports, however now practically all hard disc interfaces are SATA standard. Although IDE refers to a category of hard disc connectors, in practical terms, it is also known as the first IDE-type hard disc, or ATA-1. Technology advancement has done away with this kind of interaction. As time goes on, new hard disc interface types are created, including ATA, Ultra ATA, DMA, Ultra DMA, and various IDE hard disc interfaces (Figure).



Figure 5.1: Integrated Drive Electronics

IDE Mode

IDE supports PIO (Programmed I/O), DMA (Direct Memory Access), and Ultra DMA as its three transmission modes (UDMA). PIO mode's primary flaw is that it uses a lot of CPU resources. Data transmission rates for the IDE interface in PIO mode range from 3.3 MB/s (PIO mode 0) to 16.6 MB/s (PIO mode 4). DMA modes come in two varieties: Single-Word DMA and Multi-Word DMA. The maximum transfer rate for Single-Word DMA mode is 8.33 MB/s, while the maximum transfer rate for Multi-Word DMA (Double Word) is 16.66 MB/s. The main distinction between the DMA and the PIO is that the DMA mode runs less dependently on CPU instructions, saving the operational code of the processor. PIO and DMA are promptly replaced by UDMA as a result of the appearance and quick rise in popularity of the UDMA mode. The Ultra ATA system, which is centered on the 16-bit Multi-Word DMA mode, uses UDMA as a standard protocol. One benefit of UDMA is that it uses CRC (Cyclic Redundancy Check) technology to improve the speed of

error detection and debugging during data transfer in addition to the benefits of DMA mode. Since the Ultra ATA standard's launch, its interface has used DDR (Double Data Rate) technology to boost transmission speed by double it, with a communication speed of up to 100MB/s.

IDE Benefits and Drawbacks

Advantages: Compatibility and affordability.

Cons: Limited number of connected devices, limited cable length, and sluggish data transmission speed, inability to allow hot swap, and poor interface speed upgradeability.

Simple SCSI

SCSI, which is entirely distinct from IDE and has the full term "Small Computer System Interface," is a kind of interface. SCSI is a high-speed data transfer system that is commonly utilized in minicomputers but is not especially designed for hard drives. It benefits from a broad variety of applications, multitasking, high bandwidth, minimal CPU utilization, and hot swapping. SCSI is therefore mostly used in high-end workstations and medium to high-end servers. But because of the increased cost, it is harder to become well known than IDE. SCSI might have some issues as well. It must be configured for every computer and only supports a small subset of system BIOS. Furthermore, there is no standard SCSI software interface.

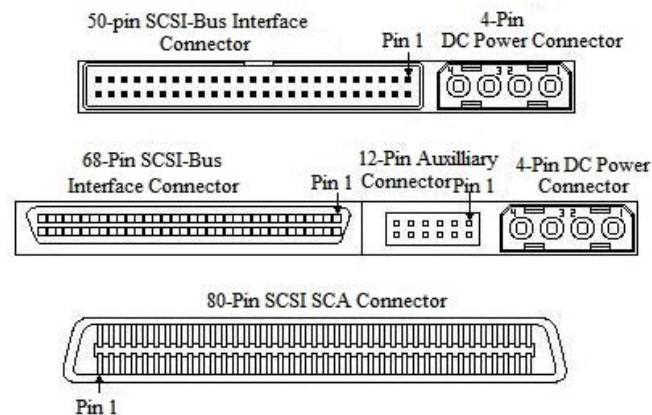


Figure 5.2: Simple SCSI

SCSI Version

- A. SCSI-1 is created in 1986, supported both synchronous and asynchronous SCSI peripherals; first released in 1979.
- B. SCSI-2 created in the year 1994. Fast SCSI, introduced in 1992, supported all SCSI devices.
- C. SCSI-3 is launched in 1995. It is presently accepted as the standard.

ATA Serial Definition: SATA, often known as serial ATA, is an acronym for serial advanced technology attachment. It is an interface for connecting ATA hard drives to the motherboard of a computer. Serial connection mode is used by SATA. The integrated clock signal used by the serial ATA bus has a better capacity for error correction. The main change from the previous is that it

can now examine transmission instructions (not just data). The automated correction of errors significantly improves the dependability of data transmission. SATA interfaces are now the standard interface. Because SATA offers much higher performance than IDE, it may take the role of IDE. SATA provides hot swapping and hot plugging and has a speed that is substantially greater than IDE (Figure).



Figure 5.3: ATA Serial

SATA Interface

The majority of the PCs we use have SATA connections as well. There are three versions of the SATA interface in use today: 1.0, 2.0, and 3.0. The greater the performance, which is mostly because of the quicker data transmission rate, the higher the version number, the later it appears. Although there have been four upgrades since its inception, namely 3.1 through 3.4, SATA 3.0 is still the most widely used interface today. The SATA interface version has backward compatibility, meaning that both the higher and lower versions are compatible. Jumpers are offered by several SATA hard drives. The SATA interface version number of the same hard drive is vary because of the various jumper settings. Additionally, a SATA motherboard must be supported for the interface's real transfer rate.

Greater interface needs, less size, and improved performance characterize SSD. Although the majority of high-performance SSDs now use M.2, U.2, and PCIe, SATA connectors won't be obsolete anytime soon. It remains popular, particularly in the HDD market. The SATA 3.3 standard, which has been improved by SATA-IO, adds a few additional capabilities as well as the ability to remotely turn off the device and optimise SMR compatibility. Storage density of HDDs may be increased by 25% using SMR (shingled magnetic recording) technology. In 2009, the SATA interface upgraded from the 3.0 standard to the 6Gbps era. Updates to SATA 3.1, 3.2, and 3.3 were made in 2011, 2013, and 2016, respectively. There aren't many brand-new features included in these subversion upgrades. Since speed is not the bottleneck of HDDs, significant interface improvements are difficult to achieve.

SATA vs. IDE Interface

SATA hard drives provide various benefits over IDE hard drives, including a novel design structure, quick data transfer, and space savings:

- A. SATA hard drives transmit data more quickly than IDE hard drives. SATA has a maximum transmission speed of 150MB/s. With the progress, it will reach 300 MB/s and 600 MB/s. At that point, the transfer rate will be over ten times quicker than that of IDE hard drives.
- B. The SATA cable is more compact and flexible than the PATA40-pin data cable used with IDE hard drives. Additionally, the transmission distance is considerable and can be increased to 1 metre, making it simpler to attach wiring and other components within the machine. This kind of cable significantly increases airflow within the computer and hastens the dissipation of heat in the case due to the tiny size of the connection.
- C. The power use has been decreased. Hard drives connected through SATA may run on 500 mA of electricity.
- D. By adopting multi-purpose chipsets or serial-parallel converters, SATA may be made backwards compatible with PATA devices. There is no need to update or modify the operating system since SATA and PATA may use the same disc.
- E. The master and slave disc jumpers do not have to be adjusted for SATA. It will be numbered by the BIOS in the sequence 1, 2, 3. While the IDE hard disc requires jumpers to configure the master and slave drives.
- F. SATA may be utilized as a U disc and also allows hot plugging. Hot swap is not supported by IDE hard drives.

Introduction to Parallel SCSI

A system interface known as the Shugart Associates System Interface was created by Shugart Associates and NCR in 1981. (SASI). SASI was created to provide a high-performance, proprietary standard mainly for usage by these two businesses. However, the standard was upgraded to a more robust interface and renamed SCSI in order to boost the industry's adoption of SASI. The new SCSI was recognized by the American National Standards Institution (ANSI) as an industry standard in 1986. SCSI, originally designed for hard drives, is sometimes contrasted with IDE/ATA. SCSI is appropriate for high-end systems because to its enhanced performance, expandability, and compatibility choices. SCSI's appeal among desktop users at home or in business is limited, nevertheless, by the high expense associated with it.

Development of SCSI

The interfaces used to interact with devices vary depending on the device before SCSI was developed. For instance, a hard disc drive was the only device that could utilise an HDD interface. SCSI was created to provide a method of connecting to and accessing host systems that is device agnostic. SCSI furthermore offered a multi-device, effective peer-to-peer I/O bus. SCSI is now a widely used hard disc interface. However, SCSI may be used to add hardware to the host computer, such as tape drives and optical media drives, without changing the hardware or software of the whole setup. SCSI has experienced significant revisions over time and developed into a reliable industry standard. This section goes into depth about several SCSI standards.

SCSI-1: The initial standard that the ANSI accepted is now known as SCSI-1, which was renamed to set it apart from later SCSI versions. The fundamentals of the first SCSI bus, such as cable length, signaling properties, instructions, and transfer modes, were established by SCSI-1. Only single-ended transmission and passive termination were supported by SCSI-1 devices. The highest data transmission rate for SCSI-1 was 5 MB/s and it made use of a constrained 8-bit bus.

Implementations of SCSI-1 led to numerous subsets of standards and incompatible devices. These problems led to the SCSI-1 standard being improved in 1985, a year before it was formally approved.

SCSI-2: A working paper was developed to provide a set of standard instructions for a SCSI device in order to manage the many issues brought on by the nonstandard implementation of the original SCSI. The common command set, or CCS, served as the foundation for the SCSI-2 standard. In addition to defining and formalizing the SCSI instructions, SCSI-2 was focused on boosting performance, increasing reliability, and introducing new capabilities to the SCSI-1 interface. In 1994, the ANSI discontinued the SCSI-1 standard and accepted SCSI-2 as a single, comprehensive document, X3.131-1994. SCSI-2's backward compatibility with SCSI-1 meant that concerns about the switch from SCSI-1 to SCSI-2 were minimal.

SCSI-3: SCSI-3, the next iteration of the SCSI standard, has been in development since 1993.

As opposed to SCSI-2, which is a single huge document, SCSI-3 is a collection of smaller but related standards.

Architecture SCSI-3

Different SCSI-3 standards and prerequisites for SCSI-3 implementations are defined and categorised by the SCSI-3 architecture. (See Technical Committee T10's "SCSI Architecture Model-3 (SAM-3)" paper at www.t10.org for further details.) The ANSI accepted the SCSI-3 design and published it as standard X.3.270-1996. This architecture makes SCSI easier to comprehend and use for consumers, hardware designers, and developers. The following are the three main parts of a SCSI architectural model: The SCSI-3 command protocol is made up of both general commands that may be used by all devices and device-specific commands that are specific to a particular type of devices. Transport layer protocols are a common set of guidelines used by devices to connect and exchange data. Physical layer interconnects: These are the specifics of the interface, such the data transmission and electrical signaling protocols. The ANSI software interfaces for SCSI devices are frequently used access techniques. The SCSI-3 standards architecture is shown in Figure 3 along with groupings of associated further SCSI-3 standards.

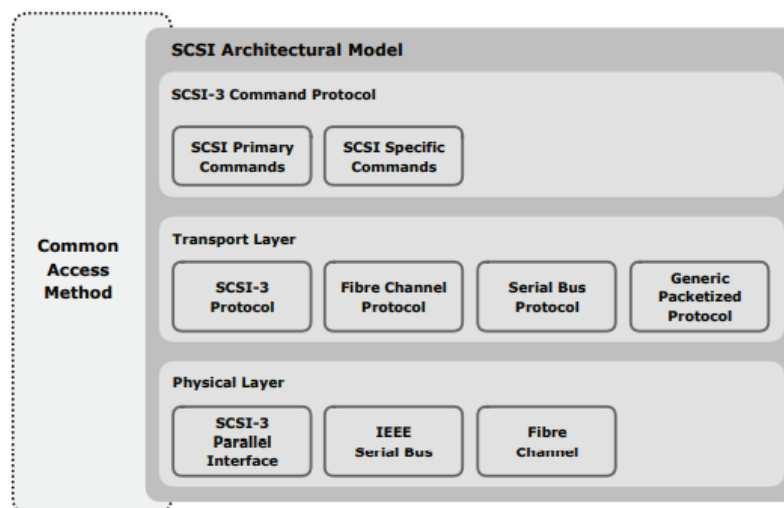


Figure 5.4: Architecture SCSI-3

Model for Client-Server

The client-server relationship, in which a client submits a service request to a server, which subsequently provides the requested service, is the foundation of the SCSI-3 architecture. An initiator-target paradigm stands in for the client-server model in a SCSI context. A specific SCSI device may function as a SCSI target device, a SCSI initiator device, or a SCSI target/initiator device in a SCSI-3 client-server paradigm. Each gadget carries out the following tasks:

A command is sent from the SCSI initiator device to the SCSI target device in order to accomplish a job. An instance of an initiator is a SCSI host adapter. **SCSI target device:** Carries out instructions for the job given by a SCSI initiator. Typically, a target device is a SCSI peripheral. The host adaptor, however, may sometimes also function as a target device. The SCSI-3 client-server concept is shown in Figure 4, where a SCSI initiator, also known as a client, requests something from a SCSI target, also known as a server. Using the protocol service interface, the target completes the tasks asked and provides the output to the initiator. One or more logical units may be found on a SCSI target device. According to the SCSI command standards, a logical unit is an entity that implements one of the device functional models. The logical unit executes the instructions that a SCSI initiator sends. As shown in Figure, a logical unit consists of two parts: a task manager and a device server. Requests from clients are handled by the device server, while administrative tasks are carried out by the task manager.

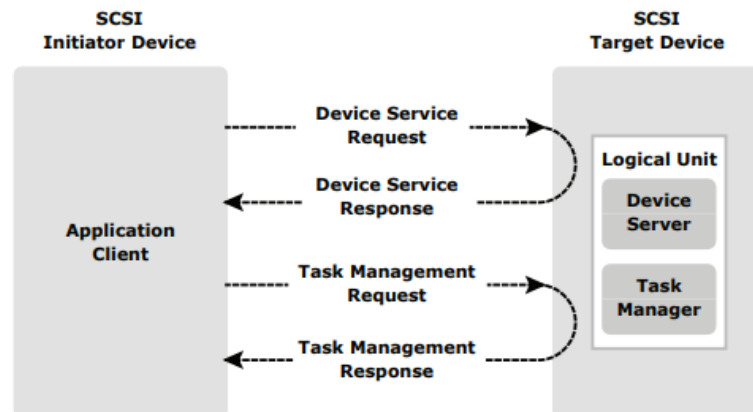


Figure 5.5: Model for Client-Server

Client-Server SCSI-3 Model:

The client-server relationship, in which a client submits a service request to a server, which subsequently provides the requested service, is the foundation of the SCSI-3 architecture. The client-server paradigm is represented by an initiator-target concept in a SCSI environment. A specific SCSI device may function as a SCSI target device, a SCSI initiator device, or a SCSI target/initiator device in a SCSI-3 client-server paradigm. Each gadget carries out the following tasks:

SCSI initiator device: Sends instructions to the SCSI target device so it may carry out a job. An instance of an initiator is a SCSI host adaptor.

SCSI target device: Puts instructions received from a SCSI initiator into action. Typically, a target device is a SCSI peripheral. The host adaptor, however, may sometimes also function as a target device.

The SCSI-3 client-server concept is shown in the following figure, in which a SCSI initiator, also known as a client, requests something from a SCSI target, also known as a server. Using the protocol service interface, the target completes the tasks asked and provides the output to the initiator. One or more logical units may be found on a SCSI target device. According to the SCSI command standards, a logical unit is an entity that implements one of the device functional models. The logical unit executes the instructions that a SCSI initiator sends. As shown in Figure, a logical unit consists of two parts: a task manager and a device server. Requests from clients are handled by the device server, while administrative tasks are carried out by the task manager.

Model for SCSI Communication: The SAM-3 defines three interconnected levels that make up a SCSI communication paradigm, which is comparable to the OSI seven-layer model. The lower levels provide services to the higher layers. By using the services that the low-level layer offers, a high-level layer may interact with it. The communication between entities at peer layers is defined by the protocol at each tier (Figure).

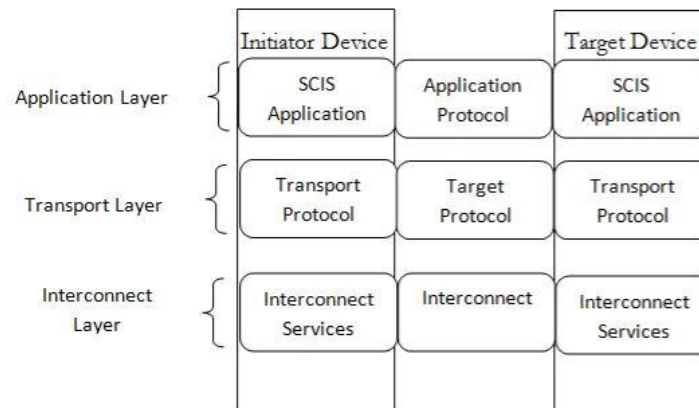


Figure 5.6: Model for SCSI Communication

The SCSI communication paradigm has three layers:

- A. SCSI application layer (SAL): This layer is home to client and server software that uses a SCSI application protocol to initiate and handle SCSI I/O activities.
- B. The SCSI transport protocol layer (STPL), which enables communication between an initiator and targets, is composed of the services and protocols.
- C. Data flow between the initiator and targets is facilitated by the interconnect layer. The service delivery subsystem, also known as the interconnect layer, includes the services, signaling methods, and interconnects for data transmission.

Addressing SCSI

For SCSI devices, Linux provides a four level hierarchical addressing scheme:

- A. SCSI adapter number
- B. Channel number
- C. Id number
- D. Lun

The standard SCSI shorthand for logical unit number is "LUN." The words in parentheses refer to the device pseudo file system's naming conventions (devfs). In the following description, "bus" is preferred over "channel." The SCSI adapter number is often generated randomly from the adapter cards on the computer's internal IO buses (such as PCI, PCMCIA, ISA, etc.). Such adapters are sometimes referred to as HBAs (host bus adapters). The kernel assigns SCSI adapter numbers in ascending order beginning with 0. A SCSI bus or buses may be controlled by a single HBA. In Appendix A, the different SCSI bus types are enumerated. Many SCSI devices may be linked to a single SCSI bus. The HBA occupies one SCSI id number and is referred to as the "initiator" in SCSI terminology. Targets, sometimes referred to as SCSI devices, and are communicated with by the initiator. The number of IDs on SCSI parallel buses is inversely proportional to the width. 8-bit buses, sometimes known as "narrow" buses, may contain 8 SCSI ids; 1 of them is utilized by the HBA, leaving the remaining 7 available for SCSI devices. A maximum of 15 SCSI devices (targets) may be connected to wide SCSI buses, which are 16 bits wide. A SCSI bus may have a large number of ids according to the SCSI 3 draught specification. Multiple Logical Unit Numbers may be present on a single SCSI device (LUNs). These are usually used by high-end tape and CDROM devices that handle various media.

CHAPTER 6

STORAGE AREA NETWORK

Dr. Santosh S Chowhan, Assistant Professor

Department of Data Science & Analytics, School of Sciences, Jain (Deemed-to-be University), Bangalore-27, India

Email Id- santosh.sc@jainuniversity.ac.in

A dedicated high-speed network or sub network known as a "storage area network" (SAN) links storage devices and provides shared storage pools to several servers. Storage accessibility and availability are crucial issues for business computing. For many business applications, traditional direct-attached disc installations inside of individual servers might be an easy and affordable choice. However, the disks and the crucial data they hold are connected to the physical server via a specialized interface, such as SAS. Modern workplace computing sometimes need a far greater degree of control, management, and adaptability. The storage area network developed as a result of these requirements (SAN). SAN technology provides a distinct, dedicated, and highly scalable, high-performance network intended to link a large number of servers to a variety of storage devices in order to meet sophisticated business storage needs. The storage may then be set up and controlled as tiers or pools that work together. Utilizing additional technologies, such as data deduplication and RAID, can optimize storage capacity and greatly improve storage resilience when compared to conventional direct-attached storage. A SAN enables an organization to treat storage as a single collective resource that can also be centrally replicated and protected (DAS). A fabric layer, host layer, and storage layer make up a storage area network.

Purposes Storage Area Networks Provide: A SAN is, to put it simply, a network of discs that a network of servers may access. SANs are widely used in a variety of corporate computing applications. Storage consolidation often involves the use of a SAN. For instance, it's typical for a computer system to include one or more local storage devices, such as a server. However, imagine a data centre with a large number of servers, each of which is running a virtual machine that can be launched and moved across the servers as needed. If the data for one workload is kept on that local storage, it could be necessary to relocate the data if the job is transferred to another server or to restore it if the server dies. A company may decide to migrate storage to a specialised storage subsystem, such as a storage array, where the storage may be collectively provided, controlled, and secured, rather than attempting to organise, track, and utilise the physical discs stored in individual servers around the data centre. Additionally, a SAN may increase storage availability. An interruption in one network connection may often be resolved by opening an alternate way across the SAN fabric since a SAN is basically a network fabric of linked computers and storage devices. Thus, business applications may still access storage even if a single cable or device fails. Additionally, treating storage as a single resource might increase storage usage by getting rid of "forgotten" drives on idle servers. Instead, a SAN gives managers a single place for all storage and allows them to pool and control the storage devices collectively. By enhancing its capacity to serve corporate workloads, each of these use cases may improve the organization's regulatory compliance, disaster recovery (DR), and business continuity (BC) postures. However, it's crucial to comprehend how SAN technology varies from conventional DAS in order to recognise its usefulness. Storage that is directly associated to a host and cannot be shared.

Through a specialised storage interface, such as SATA or SAS, one or more discs are directly attached to a particular computer with DAS. The drives are often used to store data and programmes tailored for that particular server. The DAS devices on a server may be accessed from other servers, however communication between the servers happens across the LAN's shared IP network along with other application traffic. Large data transfers via a typical IP network may be time-consuming, and the bandwidth requirements of these movements might have an impact on the server's application performance.

A SAN functions in a very different way. All of the drives are connected via the storage area network, or SAN. The dedicated network is a distinct entity from the shared LAN. This method basically treats storage as a single shared resource by allowing any server linked to the SAN to access any of the discs associated to the SAN. No SAN storage data must traverse the LAN, reducing LAN bandwidth requirements and maintaining LAN performance. The network may be built to prioritise performance and reliability, which are advantageous for corporate applications, since the SAN is a distinct dedicated network. Different storage designs and advantages come with SAN, DAS, and NAS.

Storage arrays, which are specialised storage subsystems that enable SANs, may expand to accommodate hundreds or even thousands of discs and can handle a large number of storage devices. Similar to this, a SAN may host several servers, and any server with an appropriate SAN interface can access the SAN and its enormous storage capacity. Fibre Channel and iSCSI are the two main networking technologies and interfaces used for SANs. Fibre Channel is one. When optical fibre cable and interfaces are employed, FC, a high-speed network that offers data speeds up to 128 Gbps over metropolitan area distances of up to around 6 miles or 10 km, is known for its high throughput and low latency. Block level storage may be centralised in one area with the use of this kind of dedicated network, while servers may be dispersed around a campus or a city. When storage and servers are nearby and distances don't exceed 100 feet, traditional copper cable and associated FC interfaces may also be employed (10 meters). The terms Gigabit FC and throughput FC have replaced the previous terms in more recent times, and the newest versions of the interface guarantee 128 and 256 GFC, respectively. As a network interface, FC offers a number of topologies, such as switched fabric, arbitrated loop, and point-to-point, similar to current Ethernet. Each server, storage device, FC network switch, or other network device has an FC host bus adapter (HBA) installed. Data is transferred over one or more ports on each HBA. HBAs and switches may work together to create a network fabric by connecting physical and virtual ports through cables.

iSCSI: Another network designed to link computers and shared storage is iSCSI. Although it can operate at rates of up to 100 Gbps, it offers data centre operators a number of conveniences. Unlike FC, which provides a unique and highly specialised network architecture, iSCSI combines Ethernet and TCP/IP networking with standard SCSI block data and command packets. In many circumstances, iSCSI may run on the same Ethernet LAN without a separate LAN and can share data over the LAN, WAN, and even the internet. This allows iSCSI storage networks to utilise the same cabling, network adapters, switches, and other network components used in any Ethernet network. The iSCSI data access is really another locally attached SCSI disc in the eyes of each server's operating system. The initiators and targets notions are used by iSCSI. A server taking part in the iSCSI SAN and sending SCSI instructions over an IP network is referred to as an initiator. Initiators may be hardware-based, like a storage array, or software-based, like an

operating system. A target may be either another computer or a storage resource, such a dedicated, network-connected hard drive storage device.

SANs operate:

In essence, a SAN is a network designed to link servers and storage. Any SAN's objective is to remove storage from individual servers and place it somewhere else where storage resources may be controlled and safeguarded centrally. Physically implementing such centralization involves installing drives into a specific storage subsystem like a storage array. However, centralised control may also be handled logically by software, such as VMware vSAN, which uses virtualization to identify and combine available storage. Storage traffic performance may be enhanced and expedited by connecting the collective storage to servers across a different network from the conventional LAN since the storage traffic no longer has to compete for LAN bandwidth required by servers and their workloads. As a result, corporate workloads may be able to access incredible storage volumes more quickly. The host layer, the fabric layer, and the storage layer are the three main layers that make up a SAN. Each layer is unique in its elements and traits.

The host layer: The servers that are connected to the SAN are represented by the host layer. Most of the time, the hosts, or servers, are running storage-dependent corporate workloads like databases. To allow the server and its workload to connect with other servers and users, hosts commonly utilise classic LAN Ethernet components. But in addition, SAN hosts have a unique network adapter built in for SAN connectivity. Host bus adapter is the name of the network adapter most FC SANs utilise (HBA). Like the majority of network adapters, the FC HBA uses device drivers to connect the HBA to the server's operating system and firmware to control its hardware. Through the operating system, this setup enables the workload to transmit storage orders and data to the SAN and its storage resources. One of the most well-liked and potent SAN technologies is FC, while iSCSI and InfiniBand are also widely used SAN technologies. When choosing a SAN technology, the business must carefully assess its workload and storage requirements since each system comes with a unique set of prices and tradeoffs. Ultimately, the same SAN technology must be used by the host, fabric, and storage layers.

Fabric covering: The network fabric between the SAN hosts and SAN storage is represented by the fabric layer, which is made up of the cabling and network equipment. SAN switches, gateways, routers, and protocol bridges are examples of SAN networking equipment that may be found at the fabric layer. For long-distance network communication, cabling and the accompanying ports of SAN fabric devices may use optical fibre connections, or conventional copper-based cables for close-quarters local network communication. Redundancy, or the availability of numerous alternative paths from hosts to storage throughout the fabric, is what distinguishes a fabric from a network. Numerous connections are often used when a SAN fabric is built to enable multiple pathways. SAN communication will utilise an alternate route in the event that one channel is broken or interrupted.

Layer of storage: The many storage devices gathered into different storage pools, layers, or kinds make up the storage layer. Optical media devices, such as CD and DVD drives, tape drives, SSDs, and classic magnetic HDDs are all examples of storage. Physical RAID groups, which may be used to enhance storage capacity, boost storage device reliability, or both, are how the majority of storage devices in a SAN are arranged. Each logical storage object, such as a RAID group or even a disc partition, is given a special LUN that functions fundamentally like the letter C or D on a disc drive. As a result, any SAN host may be able to access any SAN LUN across the SAN fabric.

An organisation may authorize which host to access certain LUNs by structuring storage resources and defining storage entities in this way, giving the company the ability to exercise fine-grained control over its storage assets. Zoning and LUN masking are the two fundamental techniques for managing SAN permissions. A list of LUNs that a SAN host shouldn't be able to access or that are inaccessible is called masking. Comparatively, zoning restricts host access to storage LUNs that are in an authorized permitted SAN zone by setting the fabric itself. Additionally, a SAN makes use of a number of protocols that allow applications to interact or organise data for storage. The Fibre Channel Protocol (FCP), which translates SCSI instructions across FC technology, is the most widely used protocol. An iSCSI protocol will be used by the iSCSI SANs to translate SCSI instructions across TCP/IP. There are additional protocol pairings, however, such as ATA over Ethernet, which maps ATA storage instructions over Ethernet, Fibre Channel over Ethernet (FCoE), and other less popular protocols, such as iFCP, which maps FCP over IP, and iSCSI Extensions for RDMA, which maps iSCSI over InfiniBand. In order to guarantee that all layers, operating systems, and applications can efficiently interact, SAN implementations often support numerous protocols.

Storage Area Network's Configuration

An organization must first satisfy the vendor's hardware and software compatibility criteria in order to integrate all SAN components:

Switch (firmware); host bus adapters (firmware version, driver version, and patch list); and storage (firmware, host personality firmware and patch list). Then, you must carry out the following in order to set up the SAN:

- Connect all the hardware parts using cables, and instal the necessary software.
- Examine the variations.
- Configure the HBA.
- Construct the storage system.
- Make any necessary configuration setting changes.
- Evaluate the integration: Test every aspect of the SAN environment's operating procedures, such as backup and regular production processing as well as testing for failure modes.
- Create a performance baseline for the whole SAN as well as each component.
- Publish the SAN installation and use instructions.

SAN fabric operation and architecture

A SAN's fabric, the scalable, high-performance network that links hosts (servers) and storage devices or subsystems, is its central component. The dependability and complexity of the SAN are strongly related to the fabric's design. Using optical cables for maximum speed and capability for networking over longer physical distances, an FC SAN may, at its most basic, simply connect HBA ports on servers to matching ports on SAN storage arrays. But such straightforward networking plans conceal the actual potential of a SAN. By removing single points of failure, the SAN fabric is really created to improve storage reliability and availability. Employing a minimum of two connections between any SAN parts is a key method in building a SAN. The objective is to guarantee that there is always at least one functional network link accessible between SAN hosts and SAN storage.

Host, fabric, and storage components are all included in SAN architecture.

In the figure above, two SAN hosts and two SAN storage subsystems must interact. This is a straightforward example. Because the HBA device itself is a single point of failure, each host uses a separate HBA rather than a multiport HBA. Each HBA's port is linked to a port on a various SAN switch, such as a Fibre Channel switch. Similar to this, different storage target systems or devices are connected to various ports on the SAN switch. Remove any one link from the diagram, and both servers may still interact with both storage systems to maintain storage access for the workloads on both servers. This is a straightforward redundant fabric. The fundamental operation of a SAN and its fabric. The host server will internally generate a request to access the storage device when it needs access to SAN storage. The conventional SCSI instructions used for storage access are enclosed inside network packets, in this case FC packets, and the FC protocol's regulations are followed while structuring the packets. The packets are delivered to the host's HBA, where they are loaded into copper or optical connections for the network. The request is sent by the HBA to the SAN, where it is received by the SAN switch (es). The request will be sent to the appropriate storage device via one of the switches after being received. In a storage array, the storage processor will take note of the request and communicate with the array's storage components to fulfil it.

Knowledge about SAN switches

The centre of any SAN is the SAN switch. The SAN switch, like the majority of network switches, receives a data packet, ascertains its source and destination, and then sends it to the designated destination device. The number of switches, the kind of switches (such as backbone switches, modular switches, or edge switches), and the manner the switches are linked ultimately determine the SAN fabric architecture. Modular switches with 16, 24, or even 32 ports may be used by smaller SANs, while backbone switches with 64 or 128 ports may be used by bigger SANs. Large and intricate SAN fabrics that link tens of thousands of servers and storage devices may be built by combining SAN switches. Storage resilience cannot be ensured by a fabric alone. In reality, the storage systems must include a variety of internal technologies, such as RAID, which groups discs for increased capacity and durability. RAID also has strong error management and self-healing capabilities. Thin provisioning, snapshots or storage clones, data deduplication, and data compression are just a few of the other technologies that the storage system will normally use for effective storage usage. Although a properly constructed SAN fabric enables any host to access any storage device, isolation methods, such as zoning and LUN masking, may be used to limit host access to specific LUNs for improved storage performance and SAN security.

Different SAN strategies

SAN technology has been around for a while, but recent upgrades and advancements have changed how SANs are deployed and designed. Virtual SAN, unified SAN, converged SAN, and hyper-converged infrastructure are some of these substitutes (HCI). The virtual SAN. The SAN was a perfect match for virtualization technology, which encompasses both storage and storage network resources to increase the physical SAN's scalability and flexibility. In contrast to standard SAN zoning, which effectively employs virtualization to establish one or more logical partitions or segments inside the physical SAN, a virtual SAN, marked with a capital V in VSAN, is a kind of isolation. Such isolation may be used by traditional VSANs to control SAN network traffic, improve performance, and increase security. As a result, VSAN isolation may stop possible issues on one SAN segment from influencing other SAN segments, and the segments can be altered

conceptually as necessary without affecting any physical SAN components. With the small v in vSAN, VMware offers virtual SAN technology that expands on basic VSAN approaches to provide advanced features like storage pooling or tiering, which detects and organises storage across hosts, as well as non-disruptive data migration, which moves storage from one platform to another without causing any downtime for the applications that depend on that data. Information lifecycle management, a function that enables vSAN to automatically shift data from one storage performance tier to another based on how the data is consumed, is another feature that VMware vSAN can provide. Data that is regularly accessed, for instance, may be started on a high-performance storage tier, relocated to a lower tier as it is used less often, and then eventually demoted to an archive storage tier once it is no longer needed.

Integrated SAN: Block storage, which is common for corporate applications, is supported by SANs. However, typical storage methods like network-attached storage would be required for file, object, and other sorts of storage (NAS). Multiple techniques, including file, block, and object-based storage, may be supported by a SAN that supports unified storage inside the same storage subsystem. Unified storage offers these features by supporting a variety of protocols, including block-based FC and iSCSI as well as file-based SMB and NFS. Users may benefit from potent capabilities, often reserved for conventional block-based SANs, by adopting a single storage platform for both block and file storage, including storage snapshots, data replication, storage tiering, data encryption, data compression, and data deduplication. But various storage protocols put different demands on the storage system, which may sometimes lead to inconsistent storage performance. For instance, compared to block-based data access, file-based data access may be slower and more erratic. Some business class applications may not be happy with the fluctuating requirements of unified storage systems and may nevertheless benefit from the dedicated performance features of block-based SAN.

SAN convergence: The expense and complexity of a separate network devoted to storage is a major drawback of a standard FC SAN. Using ordinary Ethernet networking components rather than FC components, iSCSI is one way to reduce the cost of a SAN. In order to combine the widely used IP and FC storage protocols into a single low-cost network, FCoE enables a converged SAN that can execute FC communication directly over Ethernet network components. To route and transfer FC data over an Ethernet network, FCoE operates by encapsulating FC frames into Ethernet frames. The choice of vendor is limited since FCoE depends on end-to-end support in network devices, which has proven challenging to provide on a large scale. Additionally, FCoE alters how networks are installed and maintained, particularly in terms of corporate data security and authentication, and businesses have been reluctant to accept such changes to their established procedures and policies.

Infrastructure that is hyper-converged: HCI use in data centres has increased significantly during the last several years. HCI integrates computation and storage resources into pre-packaged modules that can be added and controlled by a single utility as required. Modules, which are also known as nodes. To abstract and pool all the computational and storage resources, HCI uses virtualization. Following that, IT managers use the available resource pools to provision storage and virtual machines. HCI's primary objective is to streamline hardware deployment and administration while enabling rapid scaling. HCI 2.0 separates storage from computation resources so that they may be expanded independently while still achieving the same fundamental objectives. In essence, storage and compute are provided in distinct nodes. HCI is not a SAN, but depending on the needs of current company workloads, it may be utilised in lieu of SANs or even coexist

with conventional enterprise SANs. There are several factors to take into account when deciding whether to retain a SAN or switch to an HCI system.

SAN advantages

A SAN, whether physical or virtual, provides several compelling advantages that are essential for enterprise-class workloads.

Profound performance: The conventional SAN employs an exclusive network fabric for storage-related functions. For the best performance, the fabric is often FC, however iSCSI and converged networks are also options.

Excellent scalability: Extremely large installations including thousands of SAN host servers, storage devices, or even storage systems may be supported by the SAN. The SAN may be expanded as needed to fit the unique needs of the company by adding more hosts and storage as needed.

High accessibility: The concept of a network fabric, which should ideally link everything to everything else, is the foundation of a typical SAN. Because there is no single point of failure between a host and a storage device in a fully functional SAN deployment, communication across the fabric may always find an alternate channel to maintain storage availability to the workload.

Features for advanced management: Data encryption, data deduplication, storage replication, and self-healing technologies are just a few of the essential enterprise-class storage capabilities that a SAN will enable. These features are designed to increase storage capacity, security, and data resilience. Almost all features are consolidated and simply applied to all of the SAN's storage resources.

Downsides of SAN

SANs are far from flawless, however, and IT directors should take a variety of possible drawbacks into account before adopting or updating a SAN.

Complexity: Traditional SANs present the added complexity of a second network, complete with pricey, dedicated HBAs on the host servers, switches and cabling within a complex and redundant fabric, and storage processor ports at the storage arrays, even though more convergence options, such as FCoE and unified options, exist for SANs today. Such networks need to be carefully developed and managed, but the complexity is getting in the way for IT firms with smaller staffs and resources.

Scale: A SAN is often only viable in bigger, more sophisticated setups with several servers and huge storage due to the expense. A SAN can be implemented on a modest scale, but the expense and complexity are hard to justify. An iSCSI SAN, a converged SAN over a single common network like FCoE or an HCI deployment, which is skilled at pooling and provisioning resources, may often provide sufficient results for smaller installations.

Management: The concept of complexity is centred on hardware, hence managing SANs is a substantial difficulty. For busy businesses, configuring features like LUN mapping or zoning may be challenging. In order to maintain the organization's compliance, DR, and BC postures, setting up RAID and other self-healing technology, as well as accompanying monitoring and reporting — not to mention security — may be time-consuming.

Importance of SANs in enterprises:

SAN storage has existed for a long time. However, a number of variables are making this network-based storage option even more crucial for regular commercial usage. For instance: Organizations must always provide people dependable access to data. Both organized and unstructured data are expanding quickly. The technology used to manage, transfer, back up, and recover data is coming under more and more pressure from cloud computing, digitalization, and other IT developments. Organizations also want to minimize storage access problems that might impair the performance of their applications since application environments are becoming more diversified. Storage networks are also under strain from the expanding use of new technologies, such as high-performance storage systems with Non-Volatile Memory Express (NVMe), to enable newer app environments. Today, a lot of firms are also increasingly concerned with business continuity and disaster recovery, and they wish to consolidate their data into a SAN network to facilitate data replication. In other words, enterprises need to make sure the SAN they depend on is up-to-date and built to handle new difficulties in order to support rising application workload availability and performance needs and speed their access to data.

SAN fabric: SANs provide any-to-any connection for servers and storage devices by using FC-compliant Fibre Channel (FC) switches. To build the SAN fabric, devices inside a SAN are connected using FC SAN switches. The SAN fabric also includes the routers, gateways, and cables that the SAN uses to link servers and storage devices. The performance, scalability, and availability of the SAN for the company and its customers may be optimised by IT teams using fabric layer components. The SAN has two extra layers that are not a part of the SAN fabric: The servers that connect to the SAN to access data from storage devices to support application and database workloads make up the host layer. Lastly, there is the storage layer, which unifies all the storage resources, such as disc arrays, tape arrays, flash solid-state drives, RAID arrays, and more, into a single network.

Distinguishes SAN and NAS:

Network-based storage options include SANs and network-attached storage (NAS), however SANs are more complicated. A SAN is a network of several storage devices that runs independently of the operating system of an organization. It saves data at the block level and makes use of FC switches. The SAN often appears to a client operating system as a disc. To provide file-based data storage services to other devices on the network, a NAS device, also known as a NAS unit, connects to a local-area network (LAN) via an Ethernet connection. A NAS device presents itself as a file server to a client operating system. NAS storage connects directly to the LAN, however NAS also removes storage devices from the server to form a central pool of data. A dedicated network is used to offer capacity in SAN storage.

SAN's modernization disruptive: It's not necessary to be. If a company currently has the appropriate SAN infrastructure in place, it may collaborate with a top supplier of SAN technology to rapidly and painlessly update its storage area network without requiring significant equipment upgrades. The company may also invest in modern SAN technology, which can develop to keep up with the performance and availability demands of growing application workloads. This may assist the company in avoiding future forklift upgrades for its SAN infrastructure as demands evolve. The ESG white paper "Optimize, Accelerate, and Simplify SANs Non-Disruptively with Cisco MDS" provides further information on how enterprises may upgrade their SAN fabric and maximize current infrastructure assets.

Storage Area Network evolution (SAN)

High-speed data transfer: SAN employs fibre optic lines and the FCP (Fiber Channel) protocol. Because of this, it can transport data at rates that are around 20 times quicker than those of conventional copper lines.

Application Availability: SAN improves application availability thanks to block-level access.

Secured and Dynamic Failover Protection: SAN is much more secure than other storage network systems of a comparable kind. Additionally, it has top-notch data protection.

Backup: You will have several copies of your data to restore from even if you lose it. This is so that SAN can provide dependable disaster recovery using a centralised backup system.

The First Days of SAN: Ever since humans were first exposed to this technology, storage systems have been developing. A novel storage system known as Storage Area Network, or SAN, first appeared in 1993 or 1994. It wasn't all that well-liked at first, but that is to be expected with any new technology. The price of creating a brand-new storage system was also a significant consideration. As you are aware, users need a variety of parts to construct a SAN. SAN switches, servers, storage discs, tape libraries, JBODS, etc., are a few of the items you'll need. Therefore, this technology was highly costly overall at the time. As a result, only large businesses were able to benefit from this innovative storage solution. As time went on, several data storage system suppliers from all over the globe began to supply pre-built SAN storage to companies and organisations. As a consequence, those businesses became more competitive with one another and the cost of the components decreased as well. And SAN emerged as a well-liked alternative for managing and storing data. Additionally, when people realised that it transfers data using fibre channel technology and fibre channel protocol, it became popular. In more depth, we shall discuss SAN's popularity in the post's later part. Overall however, we can conclude that SAN entered the picture in the 1990s and ruled the storage system industry for more than 15 years.

SAN technology innovation and advancement

Like many other widely used technologies, SAN has been developed and enhanced in a variety of ways. Due of the intense rivalry among the providers, they began to experiment and personalise the technology. Let's investigate this further...

SSD SANs: To store data in a SAN system at initially, customers would utilise HDD (Hard Disk Drives) as the storage discs. They later began integrating SSDs (Solid State Drives) in the SAN storage system after its release. As you may already be aware, SSDs are between 5 and 25 times quicker than a regular HDD. Therefore, the technology is now quicker than ever thanks to the inclusion of SSD. Nearly all SAN suppliers in the globe have gradually begun to provide SSD in instead of HDD.

Virtualizing SAN: Back then, server virtualization was a huge deal. The storage system providers merged all of the common storage protocols, including SAN, DAS, and NAS, when virtualization technology first appeared in order to create a virtual storage pool that used all of these protocols. In this manner, a user may simultaneously access many storage systems without suffering from poor data transmission quality. However, at that time, virtual SAN was distinct from conventional SAN. However, the conventional SAN technology had to be replaced as a result of the

commercialization of the virtual server technology. The virtual SAN system is thus more common than the classic SAN in today's market.

SAN: Growing in Popularity: Prior to SAN, data was kept on servers that were directly linked to the storage discs. As a result, a server's storage space was constrained. Therefore, the user would have to upgrade their servers with new disc drives if they wanted extra storage. It cost money and took a lot of time. SAN entered the picture at that point. With the advent of access to block-level data storage, SAN provided a solution to that issue. The user may then freely build LUNs and allocate them to various servers and storage blocks. In addition, SAN was very quick in comparison to other data storage options at the time. Last but not least, SAN began to gain popularity as a result of the drop in price of its component parts. The year 2000 saw the popularity of SAN. Because of this, organizations and companies from all over the globe began using SAN as their main storage management technology. It has many of features and was dependable. As a result, SAN was developing into the next-best storage controlling network system.

The Situation with SAN: As one new technology becomes obsolete, another new one is introduced. In the IT sector, that is the absolute truth. And SAN falls under this as well. The popularity of SAN has begun to wane despite the fact that it was once the most widely used and popular storage technology. There are a few important causes behind this.

SANs' Popularity Diminishing:

The past ten years have seen significant advancements in software-defined storage systems. We don't absolutely need all the parts, such the switches, servers, etc., with the advent of hyper-converged storage architecture. With the aid of the latest cloud and hyper-converged technologies, all of these may now be merged into one. The majority of SAN providers now provide software-defined storage (SDS) solutions as an alternative to SAN, which is the reason for this. These systems are more cost-effective, always-on, high-performing, scalable, and dependable.

Finishing it off: Even while SAN isn't as popular as it once was as a storage option, many businesses are nonetheless using it right now. Additionally, the storage solution sector would have remained unchanged without the development of SAN. In fact, SAN transformed the management and storage solution industries. The method of networking data storage was streamlined, which also helped to address several important problems. Additionally, we currently employ software-defined storage systems as a result of SAN. So even if SAN may become obsolete soon, it will still exist in the shape of a better storage networking solution.

Fibre Channel: Fibre Channel is a high-speed data transport technology that delivers raw block data in-order and without any kind of compression. It has been created to link general-purpose computers, mainframes, and supercomputers to storage systems. The technology mostly provides switched fabrics (devices linked by Fibre Channel switches), but it also allows point-to-point (two devices connected directly to each other) connections. In order to link host servers to shared storage, usually shared arrays that provide block-level data storage, a storage area network (SAN) is employed. Fibre Channel SANs are often used for block-based databases used for high-speed online transaction - oriented applications (OLTP), such as those used in online ticketing, banking, and virtual environments. Fibre Channel SANs are ideally suited for low latency applications. Although it may also be used with copper cabling, fibre channel primarily uses optical fibre cables to connect data centres. Fibre Channel is a high-speed transmission technology that connects data storage to host servers by offering lossless, in-order delivery of raw block data. Most SANs are

often constructed with redundant fabrics since Fibre Channel fabrics may be extended across distance for Disaster Recovery and Economic Continuance.

A high-speed network technology called fibre channel is used to link servers to data storage area networks. On several corporate networks, it manages high performance disc storage for applications. Data replication and backup are supported. Fibre Channel is required since it is very versatile and allows for quicker data transport. The topologies that provide the fibre channel flexibility are as follows:

Point to point topology: In this architecture, a single connection links two ports. Although this architecture is less costly, a hub is not necessary. On every node, you may supply numerous "N" ports to construct point-to-point setup. The whole bandwidth provided by "N" ports is offered by each point-to-point connection. The distance between the two nodes may range from 500m (multi-mode fibre) to 10km (single-mode fibre), depending on the kind of connection being used (single-mode fiber).



Figure 6.1: Point To Point Topology

Fibre channel arbitrated loop: In this high-speed fibre channel [FC] architecture, fibre channel ports and hubs employ arbitration to create a point-to-point circuit and stop other ports and hubs from delivering frames simultaneously. Devices are arranged in a one-way ring here. In order to determine which port or hub in a loop topology may utilise the channel, ports or hubs must send an arbitration signal when they have information to convey.




Figure 6.2: Fibre channel arbitrated loop

The channel-controlling port then transmits its data to the destination port and sends an arbitrated signal marked "open." Since every port in the loop is linked, all ports will view and send along data, but only if it is specifically targeted to their port. On one controller, FC-AL may connect up to 126 ports. Many fibre channel switches still use it internally, but seldom link hosts to storage these days. If one device malfunctions or is removed, bypass circuits are provided via FC-hubs to keep the loop from breaking.

Switched fabric topology: This topology is the one that is now used extensively. The fabric in a fibre channel habitat is the collection of switches. Ports on different nodes connected to the same fabric may communicate with ports on one another. A large number of connections may be awake at once because to the fabric topology. The core of the fibre channel architecture is the any-to-any connection service and peer-to-peer communication service offered by a fabric. Both channel and network protocol may be supported by fibre channel at the same time (Figure).

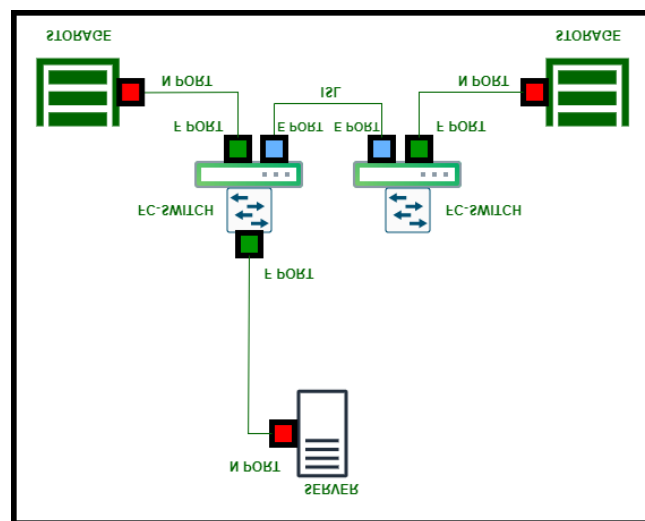


Figure 6.3: Switched fabric topology

Fibre Channel Ports: the many connections that nodes may make to one another in order to share information using the Fibre Channel topologies. Many ports that may be configured on a switch to allow these devices to properly interact with one another. Fiber channel ports essentially instruct the switch as to what kinds of devices are connected and how to handle these logged-in devices. For example, a loop-only capable device cannot connect to a fabric-only competent device. So to speak, they don't communicate in the same language. The fabric device employs port ids, but the loop-capable device uses AL PAs. The two devices use different addressing schemes.

Various ports:

N Port: N Port is a port for a node. A server or storage device can be the culprit. A single device, not more than one, is meant by a N Port. N Port doesn't exist on the switch itself. To set up N Port devices on a switch, utilise a F Port.

F Port: A fabric port is called F Port. N Ports are connected to the switch through this port. The 24 bit port address is used by this port.

L Port: This port has loop capability. You should attach the device as a L Port if it can operate in a loop.

FL Port: A fabric loop port is called FL Port. A port marked with a L often has many devices attached to it. So, even though there is just one port, there are several devices hidden behind it. Storage systems like Sun Microsystems' 3510 and 3511 might be set up as FL ports. Additionally, F-capable devices might utilise this port. If they are having trouble logging in with F, some folks may sometimes use FL instead of F.

E Port: AN extension port is called E Port. Switches are linked together using this port.

G Port: G Port This port could be present on certain Qlogic switches. Depending on whether the device is N Port or another switch, generic ports may reflect either F Ports or E Ports.

GL Port: Generic Loop Port (GL Port) A F Port, E Port, or FL Port may be represented by this port type. Everything relies on what is wired to the switch.

The legacy ports used by Sun Microsystems on the QLogic switches they market are listed below. Although they are somewhat outdated, I have included them for your convenience.

Segmented loop port: SL Port Pure fc-al devices may be connected to switches like the A5200 Sun Enterprise Network Array using these ports.

Translated Loop port (TL Port): These ports were unique ports that enabled the switch's F Ports to work with vintage T3A storage.

T Port: Trunk ports Sun connected switches using these ports prior to the E Ports. You could utilize as many of these unauthorized ports as you wanted.

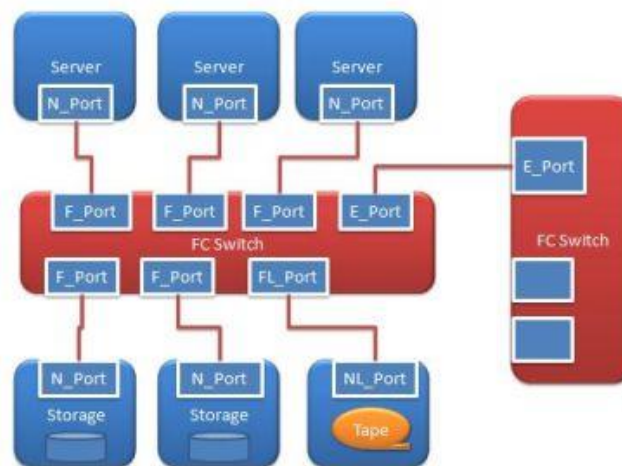


Figure 6.4: Fibre Channel Ports

Layers of Fibre Channel

To have a better understanding of a communications protocol, it is always a good idea to look at its layers. There are five layers in a fibre channel.

Physical layer, FC-0. Speed, cables, transmitters, etc.

Transmission protocol FC-1. Encoding from 8 to 10 bits.

Signaling protocol FC-2. Frames, flow control, and classes of service are ordered sets.

FC-3: Common services like multicast and striping.

Upper layer protocol, or FC-4. The application layer, in essence. IP, IPI, HIPPI, SCSI, etc.



Figure 6.5: Layers of Fibre Channel

Layers of fibre channels

Physical layer, FC-0: The physical interface and media are specified at this layer. things like speed, BER, and fibre optic cables (Bit Error Rate). From 100Mb/sec or 1Gbit/sec to 800Mb/sec or 8Gbit/sec, speeds are possible. There are already plans for faster speeds. It's incredible to consider that a serial communications system like fibre channel can achieve these speeds. Imagine a single cable clocking 8 billion bits per second. Amazingly, it isn't even a complete duplex! can accomplish 1.6Gbytes per second with full duplex. It is mind-boggling. A typical fibre optic cable may be seen in the figure below.



Figure 6.6: Physical layer, FC-0

Fiber Optic Cable: Light is used instead of electrical voltages like copper media in these fibre optic connections. Light has the benefit of being quick and able to cover large distances. Simply observe the stars at night. Some of that starlight has, in a sense, traversed billions of miles to get to us. The interference that occurs when light passes through these optic cables is similarly minimal. Because of this, we can go further. Many of the drawbacks of copper connections, such as signal interference, are eliminated by fibre optic lines. Unfortunately, copper wires cost less than fibre optic cables.

Multiple light wavelengths may pass through the core of a multimode fibre. The glass that lets light pass through is called the core. The core of multimode fibre is thicker and costs less to create. The size of the core is 50–100 microns. 62.5 micron is the most typical.

The signal may eventually become fuzzy or distorted due to the varied wavelengths of light in the core. At 100Mb per second, light may travel across multi-mode for around 500 metres. The core of a single mode fibre is substantially smaller. About 9 microns thick. The core allows just that wavelength of light to pass through it. With single mode fibre optic lines, 10 km at 100 MB/s is possible.

Using a strobe line, SCSI

Consider the two electrical cables running from one gadget to another in the following illustration. An electrical charge on the wire symbolizes the number one. No electrical charge exists at zero. The strobe line contains the solution. A short electrical current on the strobe line is sent together with each byte of data sent by the transmitting device. This instructs the receiving device to measure the data line's voltage. The voltage is a one if there is one. It is a zero if there is no voltage. Just like that. The transmitting device sends a byte or two of zeros in the diagram up top. Do not forget that a byte may have any value between 00 and FF in hex, 0 to 255 in decimal, or 0000 0000 to 1111 1111 in binary. The possibility of sending a string of zeros (0x00) or ones (0xFF) is thus extremely real. It doesn't cause an issue with SCSI. It only employs the strobe line. To demonstrate to you how SCSI would handle it, I redrew the diagram.

Eight-to-ten-bit encoding

Simply count the triangles on the strobe line in the figure to get the number of 0s that were sent. The example uses nine zeros (0 0 0 0 0 0 0 0 0). Easy enough. What about fibre channel, then? 8-to-10-bit encoding holds the key to the solution. Through the fibre optic connection, data is sent using fibre channel as a sequence of light signals. On an optical cable, a hex 00 would represent a persistent negative voltage or no light, while a hex FF would represent a sustained positive voltage or light. Unfortunately, the clock won't know if it has received more 1s or 0s. Keep in mind that the clock ticks more than 1 billion times each second. It will inevitably make a mistake. With 8bit to 10bit encoding, the standard 8 bits that are conveyed are changed to 10 bits by applying various criteria.

8-to-10-bit encoding technique

The original 8-bit byte is split into two blocks by the encoder: 3 most significant bits (MSB) y and 5 least significant bits (LSB) x. The 4 bits used to encode the 3 bit block and the 6 bits used to encode the 5 bit block. As illustrated in the above picture, they are then added together to create the new 10bit encoded value that is delivered across the fibre optic line to the other device.

Coded Characters:

Before encoding, the byte is transformed into a transmission character. If they are data characters, they will begin with a D, and if they are control characters, they will begin with a K. The K28.5 is the most popular. One of the rare characters with 5 bits of the same kind in succession is this one. This only informs the apparatus that the character is a control character and that a modification will be made. It could be an idle, a frame start, or a frame end. Signaling protocol FC-2. Frames, flow control, and classes of service are ordered sets. The most complex layer is this one. I'll attempt to convey this situation in a nutshell. Ordered sets are only groupings of encoded characters organised in 4 to inform the FC device of an impending change and its nature. The ordered sets often begin with a K28.5 and then continue with 3 further characters.

Frame made of fibre

A frame is signalled by its first 4 bytes. Normally, a collection of four bytes like this one is ordered. The first byte would be a K28.5, and the next three bytes would indicate what kind of control frame this is. The frame header is the following 24 bytes. Source id (S ID), destination id (D ID), sequence count, sequence id, exchange id, kind of frame, and other information are all included in the frame header. In essence, the header states what the frame is and where it is heading.

Standards for SAN technology

Several industry associations, notably the Storage Networking Industry Association, have established SAN technology standards, such as the Storage Management Initiative Specification. The SMI-S standard, as it is called, is designed to make storage device administration in storage area networks easier. The Fibre Channel Industry Association also promotes SAN standards, such as the Fibre Channel Physical Interface standard, which supports 64 GFC deployments, and Gen 7 solutions for the SAN market, which is the fastest industry standard networking protocol, enabling storage area networks of up to 128 GFC.

SAN administration

A SAN has significant management issues. The physical network might be complicated and must be constantly monitored. Furthermore, the logical network setup, such as LUN masking, zoning, and SAN-specific features like replication and deduplication, might vary and need frequent attention. SAN administrators should consider many management best practises to maintain the SAN running at top performance. Some of the most important procedures will rely on SAN monitoring and reporting. Administrators should monitor metrics or key performance indicators (KPIs) in the following areas of the SAN:

- A. Any kpis related to specific storage array subsystems, such as read/write throughput for each array; any kpis related to the SAN fabric or network, such as low or no buffer credits at a SAN switch or orphaned ports as zoning changes are implemented over time
- B. Any kpis related to host server I/O or workload performance, such as I/O throughput, for every virtual machine accessing the SAN;
- C. An administrator may guarantee a clear picture of the SAN's health and take proactive efforts to maintain the SAN working effectively by adopting a frequent review process and taking use of warnings and reporting tools inside the SAN.

Furthermore, features and capabilities meant to automate the SAN or reduce storage disturbances might help SAN administration. SANs that enable the use of rules for activities such as provisioning and data protection, for example, might assist administrators in avoiding oversights and errors that could waste storage or risk security. Similarly, features like native replication may help secure sensitive data while allowing ongoing access to it.

Remote SAN management is becoming more important in SAN administration. This allows SANs to be constructed in distant sites outside of the main data centre, or for a single SAN administrator to support one or more SANs from anywhere in the globe. Remote SAN administration necessitates a dependable network connection between the management tool (the administrator) and the SAN under control. The remote tool should be able to transmit full SAN health information, such as the aforementioned KPIs, as well as provide provisioning and launch diagnostics to assist in locating and eliminating any SAN faults. SolarWinds Storage Resource Monitor, IntelliMagic Vision for SAN, and EG Innovations Infrastructure Monitoring are examples of popular remote SAN products.

CHAPTER 7

NETWORK-ATTACHED STORAGE

Perna Vyas, Assistant Professor,
 School of Computer & Systems Sciences, Jaipur National University, Jaipur, India
 Email Id- Perna.vyas@jnujaipur.ac.in

The term "Network-attached Storage" (NAS) refers to a file storage system that is network-connected and allows several users to access data from a central disc capacity. LAN users utilize the Ethernet connection to access the shared storage. This storage is quick, inexpensive, and provides all the benefits of a public cloud at the location. It makes use of file access protocols including AFP, NFS, SMB, and NCP. On UNIX computers, NFS, a file-based protocol, is widely used. SMB, also known as Server Message Block, is a protocol that works with Microsoft Windows computers. Another file access protocol used with Apple systems is AFP. Those network systems, which may handle millions of operations per minute, are essentially the target audience for it. It supports storage devices for businesses that need dependable networking. Compared to file servers and external drives, it is more affordable and adaptable.




Figure 7.1: Network-Attached Storage

Network Attached Storage Architecture:

Network-attached storage (NAS) is a file-based storage architecture that makes stored data more available to networked devices. This implies that several users or client devices may get data from a single storage system. Because several clients or users linked through a Local Area Network access data from a centralized disc capacity via Ethernet, it is referred to as Network Attached Storage. Network Attached Storage (NAS) provides a single storage access point for all network-connected devices, known as a NAS Storage Server. This is a basic Network Attached Storage Architecture that indicates that there is a central storage system, i.e., NAS Server, and that various clients/users are connected to that NAS server and access data from it. NAS stores unstructured data like music, video, webpages, text files, and so on.

NAS Hardware: NAS hardware includes a NAS box, NAS unit, NAS server, or NAS head, which are basically simply servers with storage discs or drives, CPUs, and Random Access Memory (RAM). These are the fundamental hardware systems that comprise the NAS hardware framework.

NAS Software: Storage software is placed in the specific hardware of the NAS hardware system. The NAS software runs on a light operating system.

NAS Protocol: Data transfer protocols are used to deliver and receive data that is accessible by switches. The Internet Protocol (IP) and Transfer Control Protocol (TCP) are the most basic data transfer protocols used by most clients/users. The protocols' file formats are Network File Systems (NFS) and Server Message Blocks (SMB).

A network-attached storage (NAS) device is a data storage device that connects to and is accessible over a network rather than directly to a computer. NAS systems have a CPU and an operating system, allowing them to execute programmes and offer the intelligence required for data to be readily shared by authorised users. The convenience of a NAS device is that it may be accessed by several persons, computers, mobile devices, or even remotely (if set up properly).

Straight to your PC: Most of us do this at home when we need more storage than our computer or laptop can provide. A USB cable is often used to connect a hard drive or SSD to your computer's USB port. Thunderbolt cables and ports may be used by Mac users. There are methods to share access to the hard drive with others, but the hard drive is often utilised alone by the computer to which it is connected. This is sometimes referred to as direct-attached storage (DAS).

Through a network: local network at work or at home, which might be a hard-wired ethernet network or a WiFi-enabled network. As previously stated, storage devices linked to networks are referred to as network-attached storage (NAS) devices. NAS devices are often setup for access via permissions to users on an internal network, however you may normally set up access to your NAS devices through the internet. While cloud storage on a Wide Area Network (WAN) might be argued as the third technique to link computers and information, as noted below, for the purpose of this article, we'll stick to a local network.

Storage in the cloud: Azure, Amazon Web Services, iCloud, and many other services provide network-attached storage, but for the sake of this article, we'll stick to a local network. Most individuals choose to connect storage to their own networks when they want to keep expenses stable or predictable, when they want to access their data and files during internet disruptions, or when they are worried about privacy and data security.

The Future of Network-Attached Storage:

- A. Started supporting virtualization technologies.
- B. Support for Gigabit Ethernet (GigE) for quicker data transmission.
- C. Storage capacity may be increased as needed.
- D. Providing services to businesses of all sizes.

NAS device

NAS devices generally consist of a number of parts.

Drives for physical storage: NAS devices have a high-volume storage capacity and can house two to five hard drives. The redundant storage containers are organized logically using several

physical discs (RAID). Multiple physical storage components are combined into one or more logical units using the virtualization method known as RAID. Performance is enhanced and data is backed up as a result.

Central processing unit (CPU): NAS systems with central processing units (CPUs) feature a CPU that provides the computational ability and intelligence needed to maintain the file system. To process and serve files, manage many users, and, if needed, interface with the cloud, the CPU reads and writes data.

Operating system: An operating system serves as a software interface between the user and the storage device's hardware. Some basic network-attached storage devices may not have an operating system, despite the fact that more complicated ones do.

Interface for networks: Using the networking interface, the NAS device joins the network. Ethernet cables or Wi-Fi may be used for the network connection. Numerous NAS devices additionally provide USB connections for charging or connecting to other devices.

Fundamental NAS device storage concept

Network attached storage, or NAS, is used for file-based data. There are three primary methods of storage:

Storage of files: Store data in files, arrange files into folders, and put files in a hierarchy of directories and subdirectories when using file storage. It is a well-known and widely used storage method.

Block storage: A file is divided into smaller pieces (or blocks) via block storage, which saves each block independently under a different address. Blocks may be stored anywhere on the device by the computer. The blocks are put back together into the file by the server's operating system using the special address. Compared to searching across hierarchies to access a file, this is quicker.

Object holding: Discrete data pieces called objects are maintained in a database without a hierarchy or organization. Each item has the data, metadata—descriptive information about the data—and a special identification number. System software may locate and access the item using this data.

Understanding the Working of NAS

A network-attached storage device, sometimes known as NAS, is essentially a desktop personal server that customers may install. It could be connected directly to a computer via a USB connection, but it would undermine the point of the network. The private network that a NAS creates on its own is accessible by any device with the correct login credentials (username and password). A NAS is a better option than a simple external HDD and brings the customer one step closer to creating their own personal cloud storage. Any platform or operating system is compatible with NAS hardware. In its most basic form, it may be described as a group of hardware and software elements supported by an embedded operating system. A network interface card (NIC), a storage controller, a specified number of drive bays, and a power supply are often all that are required. A "box" or "head," the only piece of hardware that makes up NAS, is given an IP address and is capable of running on any platform or operating system. The box acts as the NAS's only interface with the PC clients. In order to connect to NAS, which might consist of several NAS devices linked to the same network, one can utilize an Ethernet network or LAN with an IP address.

Clients may connect to the NAS head, a single storage unit, by being given permission to do so, and users can attach numerous drives to the system to increase capacity. In NAS systems, there may be two to five hard drives, which provide redundancy and speedy file access. Although NAS is sometimes referred to as a mini-server, its controller does not function as a server; it solely works with discs for storage.

The computer can read data from and write (store) data to direct-attached storage (DAS) devices more rapidly than network devices. One could notice a variation depending on the file sizes and the task at hand. Using DAS is often beneficial while working with detailed design sheets, editing sizable photos or videos, or sending sizable data. The user may store data on a NAS configuration and access it from any other device. NAS devices are more sophisticated than DAS, despite the fact that read/write performance on them isn't as swift as on DAS external storage. A NAS device is just a piece of hardware with an Ethernet (RJ45) or Wi-Fi connection to the network. Instead of creating a Wide Area Network (WAN), this establishes a Local Area Network (LAN). Data is transferred across TCP/IP between users, servers, and a NAS after receiving an IP address. For remote file services and data sharing, NAS utilises the New Technology File System (NTFS) or the Network File System (NFS) as a standard file system. The devices supply shared storage capabilities as network-mounted volumes using a variety of protocols, including Network File System (NFS), Common Internet File System (CIFS), Server Message Block (SMB), etc. When used for shared storage, the NAS system links several servers to a single storage device. Since a cluster-shared volume enables the cluster nodes to access and retrieve the same data, these "clusters" are often employed for failover.

Key Features of NAS

NAS devices have a number of components built into their architecture to support a range of network-attached storage usage cases, including:

Storage: Network-attached storage (NAS) provides a range of storage devices that may function as file servers and act as storage simultaneously. With a hard disc, the main function of a NAS device is to store files. The most popular NAS devices for workplace workgroups, small enterprises, and home offices typically include two to five hard drives. Expanding the NAS's storage capacity is as easy as adding extra hard drives. Users don't have to replace or update their current servers in order to add extra storage; they may do this without shutting down the network. Multiple hard drives provide redundancy, quicker file access times, and more storage capabilities than a single hard disc. In NAS devices, 3.5-inch hard drives with particular NAS categories are often utilised since they can accommodate the needs of an always-running system. Secure, reliable central data storage for authorised network users and consumers. Businesses may store their data without depending on a third party by using a NAS, which gives them full control over who has access to it.

Safety: NAS systems provide built-in file system security capabilities or allow user databases to be utilised for authentication from a security standpoint. Another benefit of NAS is the reduction of network traffic due to devices being near to users. Data backup and recovery are made easy by granular security features. Another advantage is that NAS offers platform-independent access. For instance, enterprise NAS provides a method that allows users to access data independent of the OS from which they log in to permit NAS access. Given that many situations nowadays utilise more than one operating system, this is quite helpful. Obtaining data is usually challenging due to the network's pace. NAS is unlikely to outperform a server with flash memory, even with solid-state

drives (SSDs). The NAS device is made more secure by two-factor authentication by sending a pin code to a mobile device. Every time someone signs in, they must enter this pin code. This implies that even if hackers were able to guess the password, they would still be unable to access the NAS.

Reputation: To improve system speed and provide continuous user access, NAS employs specialized operating systems for network file access. By permitting common file access and compatible network protocols, these operating systems make it possible for NAS technology to meet specific requirements. Most NAS devices are equipped with an operating system. Because built-in data protection may encrypt data, users can secure their data in this way. It offers a fully functional operating system with access to additional applications. In order to expand the functionality of NAS systems, customers may do this by installing a variety of programmes, such as backup services, disaster recovery backups, corporate knowledge bases, and workplace surveillance cameras. Because they are connected to a specific computer network, users have access to data from anywhere. Simple installation and setup scripts are all that are required for NAS devices. The administration burden of a UNIX or NT file server is greater than that of NAS systems. The majority of these tools are free and come with the NAS software, which is the finest feature. With a higher usage rate and a longer MTBF, NAS drives are more durable and reliable.

RAID: The above-mentioned storage method known as RAID allows data to be duplicated from one disc to another in the event of a catastrophe or internal damage. Even while it's not the best option for a standalone backup solution, it's still a vital feature for NAS systems since businesses shouldn't retain their data in a single physical spot. A NAS system with four bays, for instance, may activate RAID and keep identical data on two of the drives. In the event that one fails, the data is still available on a backup disc. Although it doesn't totally prohibit physical data loss, it does provide some additional protection against frequent hardware failures. The RAID may be set up in a number of different ways. In cases when there is little data redundancy, RAID - 1 is advised. For more advanced data redundancy, RAID 5, 6, or 10 systems might be used. These layouts will spread the data over several drives, guaranteeing that it won't be lost even if several discs die. RAID configuration may be done manually or with software support, depending on the hardware components. The software for a RAID system is already present on modern NAS systems, it has been reported.

Examples of NAS Use Cases

Network-attached storage may be used in a variety of ways, depending on the IT architecture. The crucial ones are:

Virtualization: Users of Windows Powered Network Attached Storage (WPNAS) and Windows Server-based file servers now have access to NAS virtualization as a new tool. The utilisation is expanding, and their NAS datastores are supported by both VMware and Hyper-V. This is a well-liked choice for new or modest virtualized setups when a corporation doesn't already have a storage area network (SAN). Modern NAS systems now offer increased capability, storage, and performance. NAS systems are now being utilised in applications we never even dreamed of in the recent past thanks to increasing capabilities.

Home storage and file sharing: Many people now have NAS systems installed in their homes thanks to technological advancements, allowing them to store and share information with their loved ones. NAS, which is very elastic, scalable, and durable, may be used to build corporate

applications that communicate files and data. The primary use of network connected storage in medium-sized, SMB, and corporate distant offices is for this. When employing a hard disc, the main use of a NAS device is often to store files. The most popular NAS devices for home offices, small enterprises, and corporate workgroups typically include two to five hard drives.

Downloads of peer-to-peer files: Since NAS can do peer-to-peer file transfers discreetly in the background, a PC or laptop is no longer required to spend time on peer-to-peer file transfers through torrents. Equipment for Torrent storage NAS will be first set up by a field service specialist. It is then just necessary to set up the file transfer client, establish the connection, let the NAS manage downloads and uploads, and even request that the final file be dropped off in a central area. Controlling the device is possible using either the browser-based NAS user interface or a remote client on an Android tablet or phone.

Environments for business apps: There are open-source alternatives accessible, and many of them will serve an office from NAS even if a firm lacks the financial resources to run pricey server-based software. This is true for many business areas, including accounting, customer relationship management (CRM), enterprise resource planning (ERP), and HRM. For instance, if a CRM is required, OrangeHRM can fulfil HR requirements whereas SugarCRM or Vtiger are accessible on multiple NAS systems from various manufacturers. Users may also discover software tailored to certain industries or sectors, such as Moodle in the education sector.

Individual clouds: Given that NAS manufacturers specialise in storage and that it is not easy to make the NAS available over the internet, it is clear that they are aiming their marketing efforts at cloud storage providers like Dropbox, OneDrive, and Google Drive. A key feature of NAS-based commercial solutions for a long time has been the ability to upload and download data using a browser-based user interface. In addition, NAS storage is far less costly than conventional private cloud storage, and data ownership is maintained. However, if a user owns a company that maintains the data of other individuals, this has extra security obligations.

Multimedia in real-time: Users should transfer their huge archive of out-of-date videos to the NAS if they have a plan for them. From the DVDs, make H.264 video files and upload them as well. By installing a third-party programme like Twonky or Plex or setting the NAS's video server application, one may also create a Netflix account. Once operational, you may stream movies to linked smart TVs, media streaming devices, mobile devices, tablets, and other gadgets. No advertisements or unforeseen price rises will occur, and both file and streaming quality options are available.

Email server upkeep: It is unquestionably lot easier to set up and manage webmail or hosted email. Nevertheless, many NAS systems come with the tools needed to let an IT administrator manage their email system for either personal or commercial use. Some provide webmail so users may check their inboxes from any browser. One may access email servers with particular providers by using a POP3 or IMAP client.

Benefits

- 1) The NAS design is easy to instal and configure,
- 2) Network Attached Storage is easily accessible to any client or network user.
- 3) The reliability of NAS over conventional hard discs is a big benefit.

- 4) Consolidated storage space inside an organization's own network is another important benefit of NAS.
- 5) The performance is superb when it comes to serving files.
- 6) The NAS units are extendable and include remote access capabilities.
- 7) NAS is easy to use. Any computer linked to the LAN can save and recover data more quickly.
- 8) It further offers security.
- 9) It offers small businesses and private users a low-cost private cloud storage option.

Disadvantages of NAS:

- 1) The data transit speed is slower with NAS than with DAS, which is one of its drawbacks.
- 2) In order to use the NAS properly, users must also have a fundamental grasp of computer networks.
- 3) Clients or users who wish to back up their data can't do it right now. They can only do it with the operating system that is already in place.

Implementations of the NAS

As previously stated, there are two kinds of NAS implementations: integrated and gateway. All of the components and storage system of the integrated NAS device are housed in a single enclosure. The NAS head in a gateway configuration shares its storage with the SAN environment.

Integrated NAS (Network Attached Storage)

All of the NAS components, such as the NAS head and storage, are housed in a single enclosure, or frame, in an integrated NAS device. As a result, the integrated NAS is self-contained. The NAS head connects to the IP network in order to connect clients and serve file I/O requests. The storage comprises of a variety of discs ranging from low-cost ATA disc drives to high-throughput FC disc drives. The NAS head and storage settings are managed by management software. An integrated NAS system may vary from a low-cost single-enclosure device to a high-end solution with an externally attached storage array. A low-cost appliance-type NAS system is appropriate for applications that need Small departments may utilise it if their major requirement is storage consolidation rather than high performance or sophisticated capabilities like disaster recovery and business continuity. This solution has a limited capacity and may not be expandable beyond its initial design. To enhance capacity, the system must be expanded by deploying new units, which adds management overhead due to the various devices that must be managed. External and dedicated storage may be employed in a high-end NAS setup. This allows for independent capacity scalability in terms of NAS heads or storage. However, the scalability of this system is limited.

NAS Gateway: A gateway NAS device is made up of a separate NAS head and one or more storage arrays. While the storage is shared with other applications that need block-level I/O, the NAS head performs the same duties as it does in the integrated system. Because the NAS head and storage have distinct administrative duties, management activities in this sort of system are more difficult than in an integrated environment. A gateway solution, in addition to the components that are expressly linked to the NAS solution, may make use of the FC infrastructure, such as switches, directors, or direct-attached storage arrays. The gateway NAS is the most scalable since NAS heads and storage arrays may be ramped up individually as needed. Scaling is shown by increasing the processing capacity of the NAS gateway. When the storage limit is reached, it may scale up,

increasing SAN capacity independently of the NAS head. Administrators may improve their environments' performance and I/O processing capabilities without acquiring extra connection devices and storage. By sharing storage capacity with the SAN environment, gateway NAS provides maximum usage of storage space.

Built-in NAS Connectivity: A self-contained integrated system may connect to a regular IP network. The details of how devices are linked inside a NAS system vary depending on the manufacturer and model. Storage may be incorporated inside a NAS device and linked to the NAS head through internal connectors such as ATA or SCSI controllers in certain circumstances. In other cases, the storage may be external yet linked through SCSI controllers. External storage may be directly linked by FC HBAs or by specialised FC switches in a high-end integrated NAS solution. Backup traffic is shared on the same public IP network as ordinary client access traffic in the event of a low-end integrated NAS solution.

An isolated backup network may be utilised to separate traffic from limiting client access in the event of a high-end integrated NAS device. More advanced systems may incorporate an intelligent storage subsystem, which allows for quicker backup and greater capacity while also improving performance. Figure 2 depicts a case of integrated NAS connection.




Figure 7.2: Built-in NAS Connectivity

Connectivity to a NAS Gateway

Front-end connection in a gateway system is comparable to that in an integrated solution. Because an integrated system has a set number of NAS heads, determining IP networking needs is quite simple. In contrast, because of scalability choices, networking needs in a gateway context are difficult to calculate. Adding more NAS heads could need greater networking connection and capacity. In a gateway solution, communication between the NAS gateway and the storage system is accomplished through a standard FC SAN. Multiple data channels, redundant fabrics, and load distribution must all be addressed while deploying a robust NAS system. Figure depicts a gateway

NAS networking example. The implementation of a NAS gateway solution necessitates an examination of the present SAN infrastructure. This study is needed to establish the viability of adding a NAS workload to the existing SAN. Examine the SAN to see whether the workload is predominantly read or write, random or sequential. Determine the most common I/O size in use. Sequential workloads often have high I/O. NAS workloads are often random with modest I/O sizes. Introducing random workloads alongside sequential workloads might be detrimental to the sequential workload. As a result, it is advised to keep the NAS and SAN drives separate. Determine if the NAS workload is appropriately performing with the configured cache in the storage subsystem.




Figure 7.3: Connectivity to a NAS Gateway

The implementation of a NAS gateway solution necessitates an examination of the present SAN infrastructure. This study is needed to establish the viability of adding a NAS workload to the existing SAN. Examine the SAN to see whether the workload is predominantly read or write, random or sequential. Determine the most common I/O size in use. Sequential workloads often have high I/O. NAS workloads are often random with modest I/O sizes. Introducing random workloads alongside sequential workloads might be detrimental to the sequential workload. As a result, it is advised to keep the NAS and SAN drives separate. Determine if the NAS workload is appropriately performing with the configured cache in the storage subsystem.

NAS File Sharing Protocols:

Accessing files on a NAS system may be done via a variety of techniques or protocols. The techniques of accessing NAS systems or NAS file sharing protocols that are most often used include

- 1) Hadoop Distribution File System
- 2) Network File System (NFS)
- 3) Common Internet File System (CIFS), or Server Message Block (SMB) (HDFS)

Network File System (NFS)

NFS, a client-server protocol that functions natively across TCPIP networks and is mostly used for UNIX servers, was invented by Sun Microsystems. The implementation of it in Windows

servers, however, is less common than on UNIX systems. User Datagram Protocol served as the foundation for NFS at first (UDP). User data is represented via a machine-independent model. Additionally, it makes use of Remote Procedure Call (RPC) as a way for two computers to communicate amongst processes. For the following actions, the NFS protocol offers a set of RPCs to access a remote file system.

- 1) File and directory searches
- 2) Reading, writing to, closing, and opening a file
- 3) Modifying file characteristics
- 4) Changing directory and file links to transmit data, NFS establishes a link between the client and the distant system.

NFS is now used in three different versions. Version 2 of NFS (NFSv2), Versions 3 and 4 of the NFS protocol (NFSv4). The NFS server and NFS client are the two main elements of an NFS setup. Over the network, the NFS server exports certain folders to particular clients. To read and write from the NFS exports, the NFS client mounts them. NFS servers may be NAS devices that are executing an NFS service or general-purpose UNIX/Linux servers. The first version, NFSv2, is now almost ever used in data centres for IT.

Overview of NFSv3

Based on the Remote Procedure Call (RPC) protocol from Open Network Computing (ONC), NFSv3 is a stateless protocol that uses either UDP or TCP. Many businesses don't choose to utilise NFSv3 since it lacks strong security features. As a result, several businesses, particularly financial institutions, attempt to avoid NFSv3.

Additionally challenging to set up behind a firewall and access to through NFSv3, This is so that it may choose which network ports to connect to and listen on using the port mapper service. In essence, an NFS server informs the port mapper service which TCP/IP port numbers it is listening on for certain RPC numbers when it starts up. As a result, NFS clients must ask the NFS server's port mapper service which TCP or UDP port number to use to connect to the NFS server for a certain RPC number. If the NFS daemon is restarted, this mapping can change. As long as you can ignore the glaring security flaws, NFSv3 is still a highly functional and well-liked file-serving protocol.

Overview of NFSv4

The NFS protocol has advanced significantly with NFSv4, which puts it on par with other network file-sharing protocols like SMB/CIFS. NFSv4 is becoming more and more well-liked and has several useful features. Among the significant advancements are the following:

- 1) ACLs that are comparable to Windows ACLs are access control lists.
- 2) need tight security
- 3) Utilizing client-side caching, compound RPCs, and the well-known TCP port 2049
- 4) NFSv4.1 introduced parallel NFS as well (pNFS)

Security features in NFS4 are superior to those in NFS3. The principal security improvements include

- 1) Authentication
- 2) Integrity

3) Encryption

These security measures are supported by both Windows and UNIX servers.

Server Message Block (SMB) or Common Internet File System (CIFS)

SMB was once referred to as CIFS (Common Internet File System), a client-server protocol. A client-server application protocol called Common Internet File System (CIFS) allows client applications to send TCP/IP requests to distant machines for files and services. It is a version of the Server Message Block (SMB) protocol that is available to the public. Remote clients may access server files via the CIFS protocol. By employing unique locks, CIFS makes it possible for clients to share files. CIFS uses unicode characters to encrypt filenames. The following capabilities are offered by CIFS to guarantee data integrity.

- A. To prevent users from overwriting another user's work on a file or record, it employs file and record locking.
- B. It has fault tolerance and can automatically reopen files that were open before an interruption and restore connections.

CIFS utilizes CIFS shares, CIFS clients, and CIFS servers much like NFS shares. Any folder that is shared across a network utilizing the SMB/CIFS protocol and UNC path names is referred to as a CIFS share. For instance A Windows server that is sharing part of its discs or directories may function as a CIFS server, as can a file server instance running on a NAS array. Any device that connects to a CIFS share via the network is a CIFS client. Samba client must be installed in order to access these CIFS shares on Linux; no further software is required for Windows.

CIFS Benefits

One of the greatest file server protocols now in use is CIFS. CIFS utilises TCP port 445 and runs across TCPIP networks. CIFS provides the following beneficial characteristics.

- 1) Authentication
- 2) Encryption
- 3) Branch Caching using Quotas
- 4) advanced settings for permission
- 5) Ask for compounding
- 6) Demanding Pipelining

Microsoft released SMB 3.0 with Windows Server 2012 and included a cutting-edge capability called SMB Direct. SMB Direct enables the transmission of SMB data through RDMA (remote direct memory access) systems with high performance and low latency, such as InfiniBand. Additionally, CIFS uses Kerberos for cryptographic functions including encryption and authentication. Kerberos-based cryptographic services are not, however, supported by all NAS manufacturers.

Hadoop Distribution File System (HDFS)

User data may be saved in files using HDFS, a file system that spans several cluster nodes. In order to allow users or programmes to modify (create, rename, move, or delete) files and directories, it offers a conventional hierarchical file organisation. Additionally, it offers a streaming interface for leveraging the MapReduce architecture to execute any application of choice. Since HDFS cannot

be mounted, programmatic access is necessary. The TCP/IP protocol is placed on top of all HDFS connectivity. A master/slave architecture underlies HDFS.

Accessing and Hosting Files on NAS

The procedures necessary to host files and allow users to view the hosted files on a NAS system are as follows:

- 1) Make a storage array volume: Create volumes on the storage array and give them Logical Unit Numbers (LUNs). Show the NAS device the freshly created volumes.
- 2) Create NAS Volumes: Run a discovery operation on the NAS device to identify new array-volumes and create NAS Volumes (logical volumes). Multiple storage array volumes may be merged to make massive NAS volumes.
- 3) Create NAS file systems: On the NAS volumes, create NAS file systems.
- 4) Mount file systems: On the NAS device, mount the newly established NAS file system.
- 5) Gain access to the file systems: For client access, publish the mounted file systems on the network using NFS or CIFS.

I/O Operations at NAS: The NFS and CIFS protocols handle file I/O requests to a remote file system that the NAS device manages. The NAS I/O procedure is as follows:

- A. An I/O request is packaged into TCP/IP and sent via the network stack by the requestor. This network request is received by the NAS device.
- B. The NAS device translates the I/O request into a physical storage request, which is a block-level I/O, and then executes the operation against the physical storage pool.
- C. When data is received from the storage devices pool, the NAS device processes and repackages it into a file protocol response.
- D. The NAS device repackages this answer as TCP/IP and sends it to the client via the network.

Aspects Influencing NAS Performance and Availability

Because NAS operates on an IP network, bandwidth and latency concerns connected with IP have an impact on NAS performance. In a NAS context, network congestion is one of the most important drivers of delay (Figure 7-7). Other elements that influence NAS performance at various levels include:

Hop count: A high hop count may increase latency since IP processing is needed at each hop, adding to the delay produced by the router.

Authentication: using a directory service such as LDAP, Active Directory, or NIS: The authentication service must be accessible on the network, have acceptable bandwidth, and sufficient resources to handle the authentication demand. Otherwise, the servers get a huge number of authentication requests, increasing delay. Only when authentication happens does latency increase.

Retransmission: Retransmission might occur as a consequence of link faults, buffer overflows, or flow control mechanisms. This causes packets that have not yet arrived at their destination to be resent. When defining parameters for speed and duplex settings on network devices and NAS

heads, care must be given to ensure that they match. Improper setup may cause errors and retransmission, increasing delay.

Overutilized gateways and switches: The response time of an overutilized device in a network is always greater than the response time of an optimally or underused device. Network administrators may check vendor-specific information to identify how switches and routers are being used in a network. If the present devices are overworked, more ones should be added.

File/directory lookup and metadata requests: NAS clients use NAS devices to access files. Delays may occur due to the processing necessary before accessing the appropriate file or location. Deep directory structures may sometimes create delays, which can be rectified by flattening the directory structure. Inadequate file system structure and an overloaded disc system may also affect performance.

Overutilized NAS devices: Clients accessing many files on a NAS device might result in high usage levels, which can be identified by monitoring utilisation data. A bad file system structure or inadequate resources in a storage subsystem might create high usage levels.

Overutilized clients: Clients using CIFS or NFS data may be overutilized as well. Overloaded clients take more time to process answers from the server, increasing latency. Specific performance-monitoring tools for different operating systems are available to assist in determining client resource use. Setting the right Maximum Transmission Unit (MTU) and TCP window size, as well as configuring VLANs, will enhance NAS performance. High availability is ensured via link aggregation and redundant network setups.

CHAPTER 8

ADVANCED STORAGE NETWORKING

Perna Vyas, Assistant Professor,
School of Computer & Systems Sciences, Jaipur National University, Jaipur, India
Email Id- Perna.vyas@jnujaipur.ac.in

IP SAN transports storage traffic using Internet Protocol (IP) rather than Fibre Channel (FC) cables. It sends and receives block I/O across an IP network. Internet SCSI (iSCSI) and Fibre Channel over IP are two key technologies that use IP as a transport method for block-level data delivery (FCIP). iSCSI is a storage networking technique that enables the sharing of storage resources via an IP network. FCIP, on the other hand, is an IP-based protocol that allows remote FC SAN islands to be linked over an existing IP network. FCIP encapsulates FC frames onto the IP payload and transports them across an IP network. IP is an established technology, and employing IP as a storage networking solution offers a number of benefits.

Most businesses already have an IP-based network infrastructure that may be utilised for storage networking and may be a more cost-effective choice than establishing a new FC SAN system. Because IP networks have no distance limitations, they may be used to expand or link SANs across great distances. Organizations may use IP SAN to expand the geographical reach of their storage infrastructure and transport data that is spread across many locations. Many long-distance disaster recovery (DR) systems now use IP-based networks. Furthermore, there are several powerful and mature security alternatives accessible for IP networks. A storage system will often have both FC and iSCSI connections. This allows for both native iSCSI and FC connection in the same environment.

Overview of an iSCSI SAN

As previously stated, iSCSI is a storage networking technology that enables storage resources to be shared across an IP network, and the majority of storage resources shared on an iSCSI SAN are disc resources. iSCSI is a SCSI protocol mapping over TCP/IP, similar to how SCSI messages are mapped over Fibre Channel in FC SAN. The word iSCSI stands for Internet Small Computer System Interface, and it deals with block storage by mapping SCSI over standard TCP/IP. This protocol is generally used for sharing primary storage such as disc drives, although it is also utilised in disc backup environments in certain situations. At each stage of the network stack, SCSI instructions are encapsulated for eventual delivery over an IP network. The TCP layer ensures transmission dependability and in-order delivery, while the IP layer handles network routing.

SAN based on iSCSI

In the case of iSCSI SAN, initiators send read/write data requests to targets through an IP network. Over the same IP network, targets reply to initiators. This request response technique is used for all iSCSI interactions, and all requests and replies are sent over the IP network as iSCSI Protocol Data Units (PDUs). The iSCSI PDU is the basic unit of communication in an iSCSI SAN. Three major factors impact iSCSI performance. Dedicated iSCSI HBAs provide the best initiator

performance, while purpose-built iSCSI arrays provide the best target performance, and dedicated network switches provide the best network performance. As security is the most critical aspect of IT infrastructure, many levels of protection should be established on an iSCSI SAN. These include CHAP for authentication, discovery domains to limit device discovery, network isolation, and IPsec for in-flight data encryption.

Components of iSCSI

iSCSI is an IP-based protocol for establishing and managing links between hosts and storage devices. iSCSI is an IP-based encapsulation of SCSI I/O. iSCSI encapsulates SCSI instructions and data into IP packets that are then sent through TCP/IP. iSCSI is frequently used to transport SCSI data via IP between hosts and storage systems, as well as among storage systems. It is generally affordable and simple to instal, particularly in contexts where an FC SAN does not exist.

Key components for iSCSI communication

iSCSI initiators, such as an iSCSI HBA, are critical components for iSCSI communication. IP-based network such as a Gigabit Ethernet LAN iSCSI targets such as a storage device with an iSCSI port An iSCSI initiator transmits instructions and related data to a target, and the target responds with data and replies.

SAN IP Security

VLANS: The most prevalent way of safeguarding IP SANs is through Virtual Local Area Networks (VLANS). VLANS may be used to separate iSCSI nodes from other network devices.

CHAP: The Challenge Handshake Authentication Protocol (CHAP) is used for iSCSI target and iSCSI initiator authentication. While using Unidirectional CHAP, an iSCSI initiator authenticates itself with an iSCSI target using a secret key (i.e. a password) known as the CHAP secret; when using Bidirectional CHAP, the target additionally authenticates itself with the initiator using a second CHAP secret. When employing Bidirectional CHAP authentication, a RADIUS server may be used to facilitate CHAP secret key administration (a RADIUS server is a centralized authentication service). While the initiator's CHAP secret must still be specified, you are no longer needed to supply each target's CHAP secret on each initiator. IPsec is a collection of protocols established by the Internet Engineering Task Force (IETF) to facilitate secure packet exchange at the IP layer. IPsec is commonly used in the implementation of Virtual Private Networks (VPNs).

IPsec may operate in either Transport or Tunnel Mode:

- A. Protection is given all the way from the source to the destination in Transport Mode. For iSCSI, both the initiator and the destination must support IPsec.
- B. Tunnel mode secures gateway-to-gateway traffic. This needs no additional functionality in the iSCSI host or target driver. Until it reaches a network gateway, data in transit is unsecured. It is protected using IPsec at the gateway until it reaches the target gateway. Data packets are encrypted and confirmed at this stage.
- C. The data is subsequently delivered unprotected to the receiving host. When data must leave the safe bounds of a local LAN or WAN and travel between hosts across a public network such as the Internet, tunnel mode is often used.

FCIP: Fibre Channel over IP is abbreviated as FCIP. Storage transfer performance over IP networks is restricted, particularly over public networks, owing to ISP network latency, and this is

where FCIP comes into play. Fibre Channel over IP links FC SAN to an IP network in a transparent manner. It enables the tunnelling of FC SAN-to-SAN connections between geographically separated sites. FCIP is a system for facilitating company-wide storage access.

FCIP does not provide a built-in security feature. The IPSEC protocol suite is used to provide data secrecy and authentication, while IKE is used as the key management protocol. FCIP is a protocol that encapsulates the protocol for delivering Fibre Channel frames across IP networks. The whole FC stack, including the bottom layer, is loaded onto the TCP stack. To offer congestion management, error detection, and error recovery, FCIP takes advantage of the TCP mechanism. TCP connections are associated with an FCIP Link, which is used to establish Inter-Switch Links between Fibre Channel organisations. Fibre Channel over IP is a sophisticated mix of transport technologies that addresses the twin needs of storage networking and long-distance networking.

FCIP Terminology:

FC Gateway: A device that links multiple Ports and can route FC Frames utilising just the destination ID information in an FC Frame header.

FC Switch: An FC device that makes advantage of the FC Fabric's connecting services.

MDS 9000 IP services module SN 5428-2 Storage router FCIP port Adaptor for 7200/7400 CISCO FCIP Products Supported

Metrics for FCIP Design:

- A. GigE, OC48, or higher, with minimal latency. 60km is a short distance.
- B. OC3/OC12 has a low latency. 160km is a medium distance.
- C. Low speed, high latency on DS1/DS3. For distances more than 160 kilometres.

Considerations for High Availability - Use numerous FCIP tunnels for redundancy and enhanced performance.

FCIP Application: The following are the areas and domains where FCIP has been used because of its characteristics -

- 1) Scenario in which we have FC SANs in several locations
- 2) When there is a need for site connection
- 3) When it is necessary to transport FC traffic across an IP network as far as 500 kilometres.
- 4) Remote replication backup between two DCs.
- 5) Data replication through remote A-synchronization
- 6) Video creation that is distributed

FCIP addresses a common issue: latency. Reduces latency, allowing FCIP to work more effectively across long distances. Windows scaling, selective acknowledgements, and High Speed TCP are all examples of TCP Acceleration. Latency is also decreased by packet merging, which combines numerous smaller packets into a single bigger one, and by Dynamic Way Control, which chooses the shortest path to a distant site. When an FC error is identified, the time-out amount may need to be adjusted owing to IP Network delay, with the default value being 2 seconds.

Congestion: Improves FCIP performance over congested WANs. Lost packets are reassembled in real-time at the far end of a WAN network, reducing the delays associated with numerous round-trip retransmissions. Out-of-order packets are re-ordered, eliminating retransmission and processing delays caused by out-of-order packet arrival.

Content-addressable storage: Content-addressable storage (CAS), also known as content-addressed storage or fixed-content storage, is a method of storing data such that it may be accessed based on its content rather than its name or location. It has been used for high-speed storage and retrieval of fixed material, such as papers saved for regulatory compliance. The concept of content-addressable storage is analogous to that of content-addressable memory. CAS systems produce a unique key, the "content address," by putting the file's content through a cryptographic hash function. The directory of the file system holds these addresses as well as a reference to the actual storage of the information. Because attempting to save the same file generates the same key, CAS systems guarantee that the files contained inside them are unique, and because updating the file results in a new key, CAS systems ensure that the file is unmodified. The issue of storing information is becoming a major concern in companies as data and information volumes increase. Data that ages is less likely to be changed or updated and is referred to as fixed data. However, this fixed data has been viewed by numerous people throughout time and will be accessed further. As a result, this data cannot be erased or destroyed and must be kept for future use. The need for such fixed data or fixed content management and storage gave rise to the Content Addressed Storage System (CAS).

Among the qualities and benefits of CAS are:

Content Authenticity: It ensures the authenticity of stored items by establishing a unique or specific content address. The address allocated to each item of stored content is as unique as a fingerprint, ensuring content validity. CAS employs hashing methods, and any lost data is restored using mirrored data.

Content Integrity: This refers to the fact that the stored data has not been changed. The use of a hashing method assures integrity as well. In addition to mirroring, CAS enables parity RAID protection.

Location Independence: Location independence is achieved by the use of a unique identifier that programmes may use to get data. It makes the address visible by using a content address to access contents for apps.

Single Instance Storage: The unique address is used to ensure that only one instance of an item is stored. This address is obtained from the object's binary representation. If the item already exists on the system, it is not saved; instead, a reference to that object is produced.

Data Retention and Protection: Data retention and protection are critical. CAS produces two immutable components: There are two types of objects: data objects and meta-objects. All saved objects have access to data and meta objects. The meta-object maintains the properties of the object as well as the data processing rules. Retention rules are imposed on systems that offer object-retention capabilities until the policies expire.

Record Level Protection and Disposal: All fixed data is saved in CAS only once and is protected using a backup system. One or more storage clusters make up the array. Some CAS architectures

or designs offer further security by repeating the content into arrays in a separate place. Disposition assists with data backup and security.

Technology Independence: CAS operates on practically all technological platforms. The CAS system interface is unaffected by technological advancements. As long as the application server can map the original content address, the data is still available and usable.

Rapid Record Retrieval: CAS keeps all data on CDs. In CAS, random disc access allows for quick recovery.

Process of content-addressed storage

A CAS system assigns a content address to each object while storing data. The content address is a one-of-a-kind identifier computed from the content itself, serving as a digital fingerprint to confirm the data's legitimacy and uniqueness. Applications that need data access in a CAS system must employ content addresses to locate and obtain the required items. Data is kept on disc rather than tape in CAS, which simplifies the process of looking for archived data. Because the address of an item is dependent on its content, it may be used to verify that each stored object is unique, preventing data duplication. When an application tries to enter duplicate data, the system instead produces a reference to the original object rather than a second, identical object with the same address. (Identical items are assigned the same address.) Some CAS solutions, however, retain a backup copy of each object to improve dependability and reduce the chance of catastrophic data loss, but this data is kept distinct from the primary storage platform. Content-based naming also prevents data from being modified. When an object is updated, it obtains a new content address and the data is saved as a new object, leaving the previous object unaltered. Furthermore, once an item is saved, it cannot be removed until the chosen retention time expires.

CAS Architecture

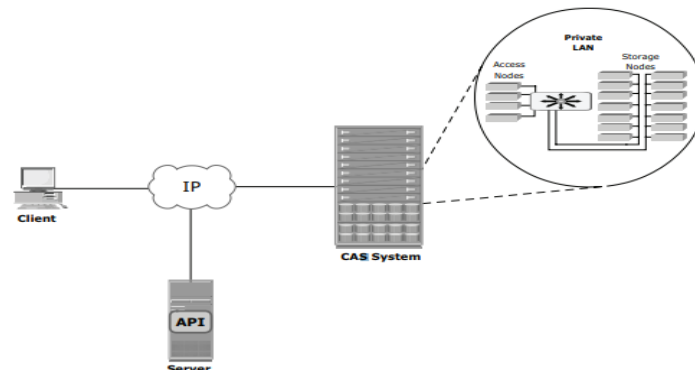


Figure 8.1: CAS Architecture

The CAS architecture is seen in Figure. A client connects to the CAS-Based storage through a LAN via the server that hosts the CAS API (application programming interface). The CAS API is in charge of carrying out functions that allow an application to save and retrieve data. The CAS design is a Redundant Array of Independent Nodes (RAIN). It consists of storage nodes and access nodes that are networked as a cluster over an internal private LAN. The internal LAN may be changed automatically to detect configuration settings such as the installation of storage or access nodes. Clients connect to the CAS through a separate LAN, which connects clients and servers to the CAS. The nodes are outfitted with low-cost, high-capacity ATA HDDs. These nodes run an

operating system with specialized software that provides the features and capabilities necessary in a CAS system.

When the cluster is installed, the nodes are assigned a "role" that defines the functionality they give to the cluster. A node may be set as a storage node, an access nucleus, or a dual-role node. Storage nodes store and safeguard data items. They are also known as back-end nodes. Access nodes link to application servers through the customer's LAN. They link to the cluster's storage nodes through a private LAN. The number of access nodes is determined by the amount of user throughput needed from the cluster. When a node is set exclusively as a "access node," its disc space cannot be utilised to store data objects. This arrangement is often encountered in older CAS implementations. Storage and retrieval requests are sent to the access node through the customer's LAN. Dual-role nodes may serve as both storage and access nodes. This node design is more common than a pure access node arrangement. Almost all CAS product provides the same features and settings. Some may be done differently, but the following characteristics are required for every CAS solution:

Integrity checking: It confirms that the file's content matches the digital signature (hashed output or CA). Integrity checks may be performed on each read or in the background. If any of the items has an issue, the nodes will automatically repair or recreate it.

Data protection and node resilience: This guarantees that the material saved on the CAS system is accessible in the case of disc or node failure. Some CAS systems have local replication or mirrors, which duplicate a data item to another node in the same cluster. This reduces overall available capacity by 50%. Another method for safeguarding CAS data is parity protection.

It requires less storage space but takes longer to regenerate faulty data. Remote replication replicates data items to a backup storage device in a different location. Remote replication is used as a disaster recovery option or for backup.

Load balancing: Distributes data items across many nodes to maximise throughput, availability, and capacity usage.

Scalability: Adding extra nodes to the cluster without disrupting data access and with little administrative cost.

Self-diagnosis and repair: Automatically discovers and fixes corrupted items, notifying the administrator of any possible issues. Failures may occur at the object or node level. They are not visible to people who view the archive. CAS systems may be set up to notify remote support teams, who can then diagnose and rectify problems remotely.

Report generation and event notification: Provides on-demand reporting and event notification. A command-line interface (CLI) or graphical user interface (GUI) may generate many sorts of reports. Any event notification may be sent to the administrator by syslog, SNMP, SMTP, or e-mail.

Fault tolerance: Ensures data availability if a component of the CAS system fails by using redundant components and data protection mechanisms. If CAS replication is used, failover to the remote CAS system happens when the main CAS system is unavailable.

Audit trails: Allow for the recording of management action as well as any data access and disposal. Compliance mandates audit traces.

Object Storage and Retrieval in CAS: When an item is saved in CAS, it is given a name that both uniquely identifies it and defines the storage location. This is referred to as a "content address." It removes the requirement for a centralised index, therefore tracking the location of stored data is no longer essential. Once an item has been saved, it cannot be removed until the retention time has passed. Data is kept on disc rather than tape in CAS. This simplifies the process of looking for stored things. Every item has a backup copy kept to improve dependability and reduce the chance of catastrophic data loss. A remote monitoring system alerts the system administrator in the case of a hardware breakdown. One key benefit of CAS is that it reduces the amount of storage space required for data backups and archives, eliminating what some engineers refer to as a "data tsunami" (the overwhelming buildup of information, much of which is obsolete, redundant, or unnecessary). Another benefit is authentication. Because an item has just one copy, establishing its authenticity is a straightforward task. The process of saving and retrieving items in CAS. This procedure requires knowledge of the following CAS terminologies:

Application programming interface (API): A high-level implementation of an interface that describes how customers may request services. The CAS API is located on the application server and is in charge of saving and retrieving items in a CAS system.

Access profile: A profile used by access apps to authenticate to a CAS cluster and by CAS clusters to authenticate to each other for replication.

Virtual pools: Allow a single logical cluster to be divided into several logical data groups.

Binary large object (BLOB): The data itself, without the descriptive information (metadata). The discrete bit sequence of user data reflects the real content of a file, regardless of its name or physical location.

Content address (CA): The address of an item generated by a hash algorithm ran over the binary representation of the object. The hash algorithm evaluates all elements of the content while creating a CA, delivering a unique content address to the user's application.

The sequence of bits that forms file content is used to generate a unique number. If even one character in the file changes, the resultant CA is different. A digest, sometimes known as a hash output, is a sort of fingerprint for a variable-length data file. This output reflects the contents of the file and is used to find it in a CAS system. The digest may be used to determine if the data is genuine or has been altered due to equipment failure or human tampering. When a user attempts to access or retrieve a file, the server sends the CA to the CAS system with the proper function to read the file. The CAS system locates the file and returns it to the application server using the CA.

C-Clip: A virtual package containing data (BLOB) and its corresponding CDF. The CA that the system returns to the client application is the C-Clip ID. It's also known as a C-Clip handle or a C-Clip reference.

C-Clip Descriptor File (CDF): An XML file created by the system when creating a C-Clip. This file contains CAs and information for all BLOBs that have been referenced. Metadata provides CAS object information such as size, format, and expiry date.

Content addressable memory (CAM): A data storage device that stores memory in cells is known as content addressable memory (CAM). When any memory aspect is entered, the CAM

compares the input to all stored data. It is a fast-paced technology. Memories in CAM are not organized chronologically and are not packed in independent modules.

CAM is utilised in very fast searching applications. Associative memory, associative storage, and associative array are other names for it. CAM characteristics include:

- 1) It is used in database management systems.
- 2) Associative memory is another name for it.
- 3) RAM is less costly than CAM.
- 4) CAM is appropriate for parallel search.
- 5) It produces a list of the data word addresses that were found.

CAM Workings:

- 1) Content-addressable memory (CAM) is a silicon device that allows for very fast yet unmistakable types of memory requests.
- 2) Queries using a CAM are conceptually similar to cooperating show reasoning in data structures, but the output is much more concise.
- 3) When a key is handed to a CAM sub-framework, the associated incentive is restored to that key. Because a "key -> esteem" pair is formed that may be used in the future.
- 4) The most important aspect is that a query of a section of a CAM may be conducted in silicon in a single clock cycle.
- 5) A RAM module that takes many clock cycles to create a single memory delivers a CAM cell in the chip, which includes two SRAM cells.
- 6) SRAM necessitates vast silicon entryways that demand a lot of power per door for rapid swapping.
- 7) Control usage in a chip generates heat and causes limits on warm dispersion due to the chip's limited impression.
- 8) When an address translation is required, utilize content addressable memory.
- 9) Encodes with high priority will be replaced with the assistance of CAMs.
- 10) It enables switching to forward traffic without overwhelming all ports.
- 11) Despite the fact that the CAM timeout is 5 minutes, a frame observed from a host is updated each time.
- 12) It may function as a search engine.

Advantages:

- A. CAM is precise.
- B. The input is connected with their memory contents in one clock cycle.
- C. The CAM is cascaded to enhance the size of lookup tables.
- D. Tables may be expanded with new entries.
- E. It is one of the higher-speed options.

Disadvantages:

- A. The cost is substantial.
- B. It has a big footprint.
- C. It uses more energy.

- D. Tables are being updated at the same time.
- E. Look up the requests on a regular basis.

EMC Centera Concepts in Action

EMC Centera is a simple, inexpensive, and secure information archiving repository. EMC Centera is the first platform that has been particularly built and optimised to cope with the retrieval and storage of fixed content while fulfilling performance, compliance, and regulatory standards. EMC Centera outperforms conventional archive systems in terms of record retrieval speed, SiS, assured content authenticity, self-healing, and support for several industry regulatory requirements.

EMC Centera Models:

EMC Centera is available in three versions to address a variety of user needs:

EMC Centera Basic

EMC Centera Governance Edition

EMC Centera Compliance Edition Plus (CE+):

EMC Centera Basic: Includes full features but does not impose retention periods.

EMC Centera Governance Edition: Adds to the features of EMC Centera Basic the retention capabilities needed by enterprises to appropriately handle electronic documents. Deploying Governance Edition enforces organisational and application information retention and disposal rules.

CE+: Offers comprehensive compliance capabilities. CE+ is intended to fulfil the criteria of the most demanding regulated business settings for electronic storage media, as defined by Securities and Exchange Commission (SEC) laws or other national and international regulatory bodies.

EMC Centera Architecture

The RAIN-based EMC Centera system's architecture is intended to be extremely scalable and capable of storing petabytes of data. While the EMC Centera cabinet can hold 32 nodes, the entry level configuration begins with as few as four nodes and two internal switches and may be expanded in increments of four nodes. To achieve all CAS features, the nodes use a Linux operating system and CentraStar software. A cube is an arrangement of nodes and internal switches. A cube may have up to 16 nodes and two internal switches.

The EMC Centera node has four SATA HDDs and a dual-source power supply, allowing EMC Centera nodes to connect to two different power sources. Each power outlet on a node is linked to a different power rail. The EMC Centera architecture is seen in Figure. Each node has over 1 TB of useable capacity and may be configured as an access or storage node. EMC Centera contains two 24-port 2 gigabit internal switches that allow up to 16 nodes in the private LAN to communicate with one another. An EMC Centera cluster may be formed by connecting several cabinets of these nodes and switches. One or more clusters comprise an EMC Centera domain. Within an EMC Centera domain, applications may support numerous clusters. Each cabinet may hold up to 23 TB of data. The useable protected capacity differs depending on whether content protection parity (CPP) or content protection mirrored is used (CPM).

CPP divides the data into segments, plus an added parity segment.

Each segment, like a file-type RAID, is on a distinct node. If a node or disc fails, the remaining nodes regenerate the missing segment on another node. Each data item in CPM is duplicated, and each mirror is located on a separate node (as seen in below mention Figure). If a node or a disc fails, the EMC Centera software immediately broadcasts to the node with the mirrored copy to regenerate another copy to a different node, ensuring that there are always two copies accessible. Using EMC Centera's unique self-healing features, both CPP and CPM offer ultimate failure protection. If any component in the node or the whole node fails, data is regenerated to another section of the cluster, ensuring that data is always secured. In addition to this "organic regeneration" process, additional processes operate in the background continually, checking items by cleaning and guaranteeing that objects are not corrupted. EMC Centera's self-management and configuration capabilities allow for quick installation and deployment. By enabling different generations to cohabit in a single CAS cluster, EMC Centera safeguards consumers against technological changes.

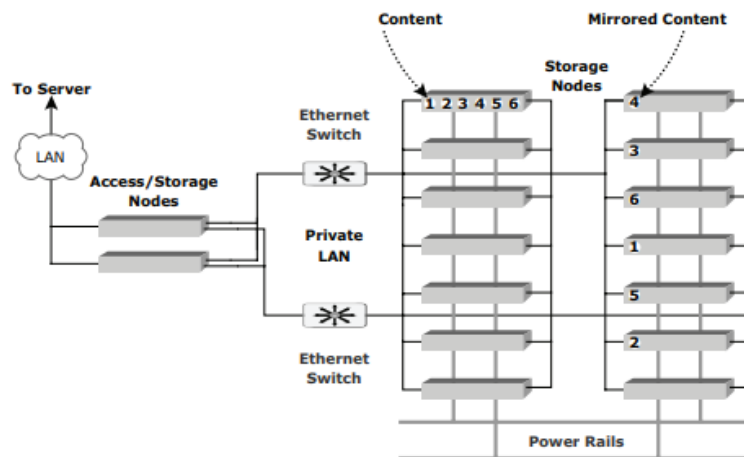


Figure 8.2: EMC Centera Architecture

Centera Instruments

Users and service workers may administer the functionality of EMC Centera using a set of tools. EMC Centera Viewer, EMC Centera Monitor, EMC Centera Console, and EMC Centera Health Reporting are among them.

EMC Centera Viewer: This is a graphical user interface (GUI) that is installed onto a client that has network access to EMC Centera. The tool makes it straightforward to see EMC Centera's capacity usage and operational performance. It also allows the system administrator to modify any site-specific information, such as public network and end-user contact information. Service professionals often utilise EMC Centera Viewer to do maintenance and update the CentraStar code.

EMC Centera Monitor: This is a utility that allows users to monitor a single EMC Centera cube by presenting system information like as configuration, capacity, and software version. EMC Centera Console is a web-based administration tool that allows system administrators to access specific information about alerts, settings, performance, and connections across various EMC Centera clusters.

Health Reporting at EMC Centera: It is accomplished by an automated e-mail message that an EMC Centera cluster delivers to the EMC Customer Support Center or a list of predetermined recipients on a regular basis. The notification provides information about the current state of the EMC Centera cluster. Remote monitoring, troubleshooting, and support for EMC Centera hardware and software are now possible.

Universal Access at EMC Centera:

When opposed to typical archive systems, an essential characteristic of EMC Centera is that the archive is available online. Furthermore, the EMC Centera archive may be accessed from any application or platform. EMC Centera Universal Access functions as a fast store and forward protocol translator. It interfaces with application servers using network file protocols (NFS, CIFS, and HTTP), as well as with an EMC Centera cluster via the Centera API (see Figure).




Figure 8.3: Universal Access at EMC Centera

EMC Centera Universal Access allows any corporate application that can mount a network disc or utilise FTP and HTTP to benefit from EMC Centera. EMC Centera Universal Access enables the use of EMC Centera in client settings with no changes to current applications, from home-grown apps to non-integrated versions of applications. This simplifies and speeds deployment significantly.

CHAPTER 9

STORAGE VIRTUALIZATION

Amita Kashyap, Assistant Professor,
School of Computer & System Sciences, Jaipur National University, Jaipur, India
Email Id-amita@jnujaipur.ac.in

The process of combining physical storage from several storage devices into what seems to be a single storage device, or pool of accessible storage capacity that is controlled from a central console is known as storage virtualization. The technique uses software to determine the amount of storage capacity that is available from physical devices, then to pool that storage into a pool that may be used by classic architecture servers or in a virtual environment by virtual machines (VMs). Input/output (I/O) requests from physical or virtual machines are intercepted by the virtual storage software, which then transmits the requests to the proper physical location of the storage devices that are a part of the total pool of storage in the virtualized environment. The user sees only the single physical drive, share, or logical unit number (LUN) that can handle normal reads and writes because the pool's different storage resources are hidden from view. A very fundamental form of storage virtualization consists of a software virtualization layer that sits between the hardware of a storage resource and a host, such as a personal computer (PC), a server, or any other device accessing the storage, and enables operating systems (OSes) and applications to access and use the storage. Sometimes, even a RAID array can be viewed as a form of storage virtualization. The user sees the array's numerous physical drives as a single storage device, but in reality, the drives are striping and replicating data to additional discs in the background to enhance I/O performance and safeguard data in the event of a drive failure.

Storage virtualization types include: File vs. block

Virtualizing storage may be done in two ways: file-based or block-based. File-based storage virtualization is a subset of network-attached storage (NAS) systems. File-based storage virtualization eliminates the dependency in a traditional NAS array between the data being accessed and the location of physical memory by using the Server Message Block (SMB) or Common Internet File System (CIFS) protocols in Windows server environments, or the Network File System (NFS) protocols in Linux systems. The pooling of NAS resources makes handling file migrations in the background simpler, which improves speed. NAS systems are typically not difficult to operate, however storage virtualization considerably simplifies the work of controlling many NAS devices via a single administration panel.

Storage resources commonly accessible through a Fibre Channel (FC) or Internet Small Computer System Interface (iSCSI) storage area network (SAN) are more often virtualized than file-based storage systems. Block-based systems decouple logical storage, such as a drive partition, from physical memory blocks in a storage device, such as a hard disc drive (HDD) or solid-state memory device. Because it works in the same way as native drive software, there is minimal overhead for read and write operations, hence block storage solutions outperform file-based systems.

The virtualization management software may aggregate the capacity of available blocks of storage space across all virtualized arrays and pool them into a common resource that can be given to any

number of VMs, bare-metal servers, or containers using the block-based operation. Storage virtualization is very useful for block storage. Unlike NAS systems, administering SANs may be time-consuming; combining a number of block storage systems under a single administration interface, for example, can be a considerable timesaver. IBM's SAN Volume Controller (SVC), currently known as IBM Spectrum Virtualize, was an early implementation of block-based virtualization. The software operates on an appliance or storage array and virtualizes LUNs linked to servers connected to storage controllers to form a single pool of storage. Customers may additionally layer block data to public cloud storage using Spectrum Virtualize.

Hitachi Data Systems' TagmaStore Universal Storage Platform, currently known as Hitachi Virtual Storage Platform, was another early storage virtualization device (VSP). Customers were able to construct a single pool of storage across several arrays, including those from other prominent storage providers, thanks to Hitachi's array-based storage virtualization.

Storage virtualization in action

To provide users access to data stored on physical storage devices, virtualization software must either generate a map using metadata or, in certain situations, utilise an algorithm to dynamically find the data on the fly. The virtualization software intercepts read and write requests from programmes and, using the map it has produced, locates and saves the data to the appropriate physical device. This procedure is comparable to how PC operating systems retrieve and save programme data. Storage virtualization masks the true complexity of a storage system, such as a SAN, allowing a storage administrator to conduct backup, archiving, and recovery chores more simply and quickly.

Virtualization in-band vs. virtualization out-of-band

There are two forms of virtualization that may be used on a storage infrastructure: in-band and out-of-band virtualization.

- A. In-band virtualization, also known as symmetric virtualization, manages data read or stored as well as control information (e.g., I/O commands, metadata) in the same channel or layer. This configuration enables storage virtualization to deliver more sophisticated operational and administrative tasks including data caching and replication.
- B. Out-of-band virtualization, also known as asymmetric virtualization, divides data and control routes. Advanced storage functionalities are often inaccessible since the virtualization facility only sees control instructions.

Methods of virtualization

Today, storage virtualization often refers to capacity aggregated from several physical devices and then made accessible for reallocation in a virtualized environment. Virtual storage is used in modern IT approaches such as hyper-converged infrastructure (HCI) and containerization, in addition to virtual computing power and, in certain cases, virtual network capacity.

Tape storage is still commonly used for preserving seldom accessed data, notwithstanding its decline as a backup destination medium. Archival data is often vast; storage virtualization may be used for tape media to make massive data repositories simpler to handle. Linear tape file system (LTFS) is a kind of tape virtualization that makes a tape appear like a conventional NAS file storage device and makes finding and restoring data from tape much simpler by employing a file-

level directory of the tape's contents. Storage may be used in a variety of ways in a virtualized environment, including:

Host-based storage virtualization is a kind of software-based storage virtualization that is often seen in HCI systems and cloud storage. The host, or a hyper-converged system made up of numerous hosts, offers virtual drives of different size to the guest machines, which may be VMs in an enterprise context, physical servers or PCs accessing file shares or cloud storage. Virtualization and administration are handled entirely via software at the host level, and physical storage may be nearly any device or array. Virtualization features are included into certain server operating systems, such as Windows Server Storage Spaces.

Array-based storage virtualization is a system in which a storage array functions as the main storage controller and runs virtualization software, allowing it to pool storage resources from other arrays and expose multiple kinds of physical storage for usage as storage tiers. A storage tier may include solid-state discs (SSDs) or hard disc drives (HDDs) on different virtualized storage arrays; the actual location and particular array are masked from servers or users using the storage.

The most popular kind of storage virtualization utilized in businesses today is network-based storage virtualization. A network device, such as a smart switch or purpose-built server, connects to all storage devices in an FC or iSCSI SAN and displays the storage as a single, virtual pool in the storage network.

Storage virtualization's advantages and applications

Storage virtualization was initially difficult to install and had limited application in terms of which brands and models of storage arrays the existing technology could operate with when it was originally presented more than two decades ago. Storage virtualization software had to be deployed and maintained on all servers that needed access to the pooled storage resources since it was initially host-based. As the technology progressed, it could be deployed in a number of methods (since explained above), making it simpler to deploy in a range of scenarios, as customers could choose the virtualization approach that best suited their companies' existing infrastructure.

Further development of virtualization software, as well as standards such as the Storage Management Initiative Specification (SMI-S), enabled virtualization products to work with a broader range of storage systems, making virtualization a much more appealing option for enterprises dealing with spiraling storage capacities.

Some of the advantages and applications of storage virtualization include:

- A. Easier administration. A single management console for numerous virtualized storage arrays reduces the time and effort required to manage the physical systems. This is especially useful when various manufacturers' storage systems are included in the virtualization pool.
- B. Increased storage usage. Pooling storage capacity across several systems simplifies allocation, allowing capacity to be assigned and utilized more effectively. With disconnected, heterogeneous systems, it is conceivable that some may operate at or near capacity, while others would be hardly utilized.
- C. Increase the lifespan of older storage systems. Virtualization allows you to prolong the usability of outdated storage equipment by incorporating it into the pool as a tier to handle archive or less essential data.

- D. Universally add sophisticated features. Tiering, caching, and replication are some advanced storage capabilities that may be implemented at the virtualization level. This contributes to the standardization of these practices across all member systems and allows these advanced features to be delivered to systems that may be missing them.

Storage virtualization types:

List every type of storage virtualization available in cloud computing, including: list every type of storage virtualization available in cloud computing, including:

- A. Hardware assisted virtualization
- B. Kernel level virtualization
- C. Hypervisor level virtualization
- D. Para-virtualization
- E. Full virtualization

Hardware-Assisted Virtualization (HAV):

Virtualization of this kind needs hardware backing. It resembles complete para virtualization. The unmodified OS can be used to manage hardware access requests, secure operations, and function as hardware support for virtualization.

Virtualization at the kernel level: It utilizes a different Linux Kernel version. Running several servers on a single host is possible at the kernel level. The main Linux Kernel and the virtual machine communicate with each other through a device driver. A unique type of server virtualization is this virtualization.

Hypervisor Virtualization: Between the operating system and hardware, there is a layer called a hypervisor. A hypervisor can support multiple operating systems in operation. Additionally, it offers the tools and services required for OS to function correctly.

Para-Virtualization: It is built on a hypervisor, which manages software emulation and trapping. Before being installed on any other machines, the guest operating system is updated. The upgraded system boosts performance by speaking directly with the hypervisor.

Totally virtualized: Similar to para-virtualization, this virtualization. The hypervisor does this by capturing the machine actions that the operating system uses to carry out the operations. Once the operations are captured, a certain piece of software is emulated, and the status codes are then returned.

Approaches to Virtualization: The term "virtualization" often refers to the pooling of various storage resources and maintaining them in a single storage in a virtual environment. More contemporary technologies, such as hyper-converged infrastructure, utilize virtual power and network resources in addition to storage resources. The various applications for these storages in a virtual environment:

Host-Based Virtualization Approach

The virtualization is done at the host level, providing the user with virtual storage with various capacity sets when the hosts are many, regardless of whether the end-user is using a virtual machine or a personal computer that contacts the cloud storage. With the aid of software, virtualization is accomplished, and any device can be used as physical storage. Let's look at some of the benefits and drawbacks of this strategy. The main benefits include its ease of designing and coding, ability to support any form of storage, and contribution to better storage use. Concerns include the fact that each OS has specific software, that synchronizing the host is challenging, and that optimization can only be done on a cost-based basis.

Array-Based Virtualization Approach

Basically, think of storage as a collection of objects that represent physical storage; often, these objects are HDDs and SDDs (Solid-State Drives). We manage these storage arrays using various software programs, and we conceal them from users and visitors. This technique has the benefits of not requiring any new hardware or infrastructure and having zero latency when attending a specific I/O. One drawback is that all storage types, including primary, secondary, and others, would demand the same amount of bandwidth, necessitating infrastructure. Another is that it is specific to the vendor's matrix and that overall storage utilization optimization is not done.

Network-Based Virtualization Approach

This strategy is well-known for being applied in many of today's large corporations. This method makes use of fiber channel, and any network device, such as a smart switch or a server that was developed specifically for the purpose, can connect to a storage area network (SAN) and appear to a guest user as a virtual storage pool. This method's main benefits include enabling real heterogeneous virtualization, boosting performance, requiring just one management device for all associated storage, and making it simple to replicate services across all connected devices. The main drawbacks are the fact that it is highly challenging to comprehend the matrices involved, that it adds some latency to I/O, that it is challenging to design and code, and that it is challenging to implement when working with rapid metadata.

Storage virtualization isolates storage hardware resources by software measures that merge the services or capabilities of one or more storage target devices with other additional functions, giving users with consistent data storage service through the abstraction layer. There are three forms of storage virtualization: host-based storage virtualization, storage device-based storage virtualization, and network-based storage virtualization.

Virtualization on the basis of a host

The host-based storage virtualization system will install storage virtualization management software in the server's host OS, allowing the server's storage space to span many heterogeneous disc arrays for data mirroring protection. In general, it is finished by logical volume management (LVM) software in the server OS, which differs depending on the OS type. The implementation idea is to place two LUNs from separate storage devices in the same volume group (VG) of the server OS and mirror the data between the two physical LUNs. Host-based storage virtualization, which requires no extra hardware, becomes the most simple-to-implement, cost-effective, and

well-developed option. However, because the need to instal storage virtualization software in the host OS adds overhead and consumes processing time on the host CPU, this solution has poor scalability and actual running performance, which may affect the system's stability and security levels, resulting in inadvertent unauthorised access to protected data

Storage virtualization based on storage devices

The storage virtualization system based on storage devices primarily operates by adding resource virtualization function modules to the controller host of the enterprise storage array, allowing the storage array to be virtualized for the management of physical storage resources other than its own storage space. Various storage suppliers have adapted this kind of virtualization technology to their enterprise-level storage controllers in recent years, resulting in a storage virtualization platform based on storage device controllers.

This controller for storage arrays allows you to connect to various brands of heterogeneous storage arrays. As a result, both the device's internal storage medium and external devices will appear in storage controllers with virtualization features that may be handled in the same. By adopting storage device-based storage virtualization, the original storage system is no longer required to manage storage resources, decreasing the difficulties of heterogeneous storage management and the storage network's complexity. It can manage the whole pool of internal and external storage resources using the virtualization system's unified LUN mapping. However, the performance of the storage controller will restrict the efficiency of the LUN mapping function.

Storage virtualization over the network

With the advent of NAS and SAN storage designs, the separation of storage from storage media, storage controllers, and servers has been accomplished. Servers and storage controllers are linked by network (TCP/IP or FC), allowing for flexible and efficient storage resource sharing, while the network also serves as the optimal place for storage virtualization. Network-based storage virtualization embeds an intelligent storage resource management device at the network layer, abstracting physical storage resources between servers and storage arrays. It links heterogeneous storage arrays to the storage layer's SAN switch through an FC interface, bringing them all together to an in-band storage virtualization device.

All LUN mapping and I/O requests, as well as data, will be transferred to the server through the device, while the server can only view the LUN given by the storage virtualization device without direct involvement. Network-based storage virtualization, with a separate storage virtualization management device, adds a number of additional features that the other two do not. It is capable of not only local heterogeneous storage data replication, but also distant data replication for disaster recovery. Building storage virtualization and disaster recovery inside data centers is a relatively low-cost option for large companies. However, this strategy is far more difficult to implement than the other two methods of storage virtualization. It must halt all LUN mappings from the server to the storage device and restart them on the storage virtualization device. And, after the operation is complete, all of the newly added storage devices cannot be de-virtualized since all of the storage LUN information has been preserved on the storage virtualization device.

Storage virtualization advantages

The following are some advantages of switching to storage virtualization now that we have seen what it is, its varieties, and how to implement them:

- A. Because keep data in a different, more convenient location, even if the host fails, our data won't be easily compromised.
- B. As build some level of abstraction in our storage, it is simple for us to protect, provide, and use our data.
- C. Additional services such as recovery, duplication, replication, etc. may be done easily.

Storage Virtualization Configurations

Storage virtualization in the network is accomplished using either in-band or out-of-band methods. The virtualized environment configuration is saved outside of the data flow in an out-of-band solution. As illustrated in Figure, the configuration is saved on a virtualization appliance that is configured to be independent of the storage network that transports the data. Because the control and data routes are separated (the control path travels via the appliance, but the data path does not), this configuration is also known as split-path. This design allows the environment to handle data at network speed while adding only minimum delay for conversion of the virtual arrangement to physical storage. The information is not cached at the virtualization appliance any farther than it would be in a regular SAN arrangement. The virtualization appliance can be considerably scaled since it is hardware-based and tuned for Fibre Channel connection. Furthermore, since the data is not changed in an out-of-band implementation, many current array features and functions, in addition to the advantages afforded by virtualization, may be used.




Figure 9.1: Out of band Storage Virtualization Configurations

The virtualization function is placed in the data route by the in-band implementation, as illustrated in Figure. General-purpose servers or appliances manage virtualization and act as a translation engine from virtual to physical storage. Data packets are often stored by the appliance during processing and subsequently delivered to the appropriate destination. In-band implementations are software-based, and data storage and forwarding via the appliance adds delay. Because the data lingers in the network for some time before being committed to disc, it causes a delay in application response time.




Figure 9.2: In band Storage Virtualization Configurations

In terms of infrastructure, the in-band design adds complexity and a new layer of virtualization (the appliance), while restricting the storage system's capacity to grow. In-band deployments are appropriate for static settings with predictable demands.

Challenges in Storage Virtualization

The capacity to control resource contention in the underlying storage infrastructure is one of the most difficult difficulties in virtualized infrastructure affecting application performance. To mitigate the effect of this difficulty, it is standard practise to over-provision storage systems with IO and disc capacity and to manage mission-critical and other IO-intensive applications in silos. This, however, is a trade-off that results in a greater cost per gigabyte of storage provided to each virtual machine, as well as a considerable influence on operational expenses associated with service assurance. Because all-flash arrays provide 10x+ the IOPS capacity of Hard Disk Drive (HDD) arrays, the adoption of so-called Solid State Drive (SSD) all-flash arrays tends to alleviate the IO bottleneck. When adding SSD arrays into virtual and cloud settings, however, the challenges of workload alignment, heterogeneous storage management, and effective capacity planning persist. Turbonomic's autonomic approach to virtualized storage management tackles these issues

by finding the whole stack of virtualized and cloud systems and empowering workloads to self-manage. Enterprises and service providers may use Turbonomic to:

- A. Reduce the cost of continuous storage infrastructure by 20-30%.
- B. Reduce operating expenses significantly by eliminating complicated storage issues and their effect on workloads, and hence on end users.
- C. Ensure application performance and risk reduction
- D. Allow for the reliable virtualization of IO-intensive applications, lowering the total cost of delivering computing services to the company and customers.
- E. Implement heterogeneous SSD/HDD storage settings in a seamless manner by recognising which applications need SSD access and when.

Storage virtualization is not a novel concept. It has been on the market for many years to address business storage issues. As data grows at an exponential rate, the necessity for storage virtualization is becoming increasingly apparent. However, it is not without its difficulties. The following issues may arise as a result of storage virtualization:

Agility and scalability: Storage virtualization may not always be a straightforward installation due to its agility and scalability. It has a few technological challenges, such as scalability. Companies have many hardware and software components provided by several suppliers. In such a case, managing software and hardware may become difficult. Due to the demanding nature of applications and the growing data, the storage system installed in such circumstances may also need quick improvements. Furthermore, difficulties such as lack of agility, scalability, better data analysis, and quicker data access are some of the obstacles that businesses must consider before selecting the best storage solution.

Data security: Data security is another issue that has to be addressed. While some may claim that virtual computers and servers are more secure than physical machines and servers, virtual environments might attract new types of cyber-attacks. Furthermore, data security and governance are becoming an issue as a result of storage virtualization.

Manageability and integration: Virtualisation disrupts your data's end-to-end perspective. The virtualized storage solution must be adaptable to current tools and systems.

Server virtualization presents significant issues to storage systems and the administrators who manage them. This basic truth is arguably one of the most major reasons why, despite the advantages of server virtualization being considerable, well-known, and genuine, only around half of the servers in data centres across the globe have already been virtualized.

At the most fundamental level, everything boils down to money. Server virtualization may result in significant cost reductions, but only if the storage systems that support it are adequate. One cause for rising storage prices is that, similar to how server virtualization decouples a virtual machine (VM) from the real hardware on which it operates, storage virtualization decouples the VM from the underlying storage, which is often placed on a SAN. Server virtualization manufacturers tout the ease with which new VMs can be created, but this may lead to VM sprawl and hundreds of ghost VMs – VMs that are no longer required or utilised yet use storage resources. This is exacerbated when VMs are created from standard images that provide significantly more storage resources than are required. In general, it is true that server virtualization, by definition, consumes a lot of storage resources. It may put a strain on storage systems because it can make

sequential accesses random precisely the kind of storage behaviour that storage systems struggle with the most." A scalable single system is simpler to manage, administer, and maintain; it takes up less precious data centre floor space; and it costs less to power and cool. When you consider that the purchase cost of a storage system might be as low as 20% of the overall cost of operating and maintaining it during its lifespan, the cost savings can be considerable.

Virtualization of the File System

The notion of a networked remote file system (also known as network attached storage, or NAS) such as NFS or CIFS is the most basic kind of file system virtualization. Dedicated file servers handle shared network access to files in the file system in this kind of virtualization. That file system is shared by multiple network hosts running various operating systems. Through an NFS share, for example, Windows and Unix hosts may access the same file system. The files on the shared file system are accessible using the same high-level procedures as are used to access local file systems, regardless of operating system. This implies that programmes and users may access files using the same interfaces regardless of where the file is physically located. This abstraction, or "hiding," of data location is a perfect illustration of one of storage virtualization's key features: location transparency. Another kind of file system virtualization is utilised to make database maintenance easier. In order to enhance efficiency, database table spaces and transaction logs are nevertheless often stored on raw storage devices. Other administrators choose to use a file system to handle these entities since raw disc devices are notoriously difficult to maintain. The increased overhead of the file system, on the other hand, reduces performance. In database settings, file system virtualization combines the benefits of raw partitions and file systems. The file system is exposed to the administrator, allowing for better control of database elements. However, from the perspective of the database, its entities are physically placed on raw disc drives; the file system is concealed, and so its buffered I/O is avoided, providing maximum performance.

Virtualization of Files and Records

Hierarchical Storage Management (HSM), the most extensively used form of file virtualization, automates the movement of seldom used data to low-cost secondary storage medium such as optical discs or tape drives, or low-cost high-density disc storage such as Serial ATA (SATA) arrays. This migration is invisible to users and programmes, who continue to access the data as if it were still on the main storage media. Again, virtualization produces location transparency. A file system pointer paired with information from the HSM programme guarantees that the migrated file may be quickly recovered and made accessible to the requestor without the requestor knowing the precise physical location of the file.

Virtualization of Blocks

Most current storage virtualization research has been on layer II of the SNIA shared storage model, which deals with block-level disc services, and it is this block virtualization that most manufacturers refer to when discussing storage virtualization. Beyond the fundamental CHS-LBA disc virtualization discussed previously, block virtualization is the next natural step. Whereas disc virtualization "manipulates" a single magnetic disc to portray it as logical block addresses, block virtualization takes it a step further by virtualizing many physical discs to display a single logical device. This is often known as block aggregation. The concept of block virtualization is straightforward: Overcome individual device physical constraints without needing any extra intelligence in programmes, so that the latter just perceive a "bigger" disc (a new virtual disc with

a greater logical block address range). However, basic block aggregation is just one aspect of block virtualization; additional services, such as performance, availability, and other critical storage qualities, may also be added into the block virtualization layer. In an ideal world, all storage management activities would be handled from the standpoint of the storage user (the application), rather than the storage provider (the array). The storage consumer's high-level needs are quite simple:

Capacity: Is there enough room to store all of the application's data.

Performance: Can the storage offered to the application match the program's reaction time requirements.

Availability: Is the storage offered to the application capable of meeting the application's availability requirements.

Physical features of storage, such as disc size and array count, are unimportant; storage customers do not want to deal with technical specifics; they just want to identify the storage services they want.

Storage administrators are responsible for addressing the needs of storage customers. To do so effectively in a dynamic, diverse environment, they need sophisticated tools that are flexible and scalable. Storage managers aim to be able to administer as many different servers and storage systems as possible using the same tools; their objective is easier administration. Storage virtualization is an important tool in the storage administrator's toolkit that may assist the administrator in meeting these objectives.

The purpose of block virtualization is to manage physical storage assets and combine them to create logical volumes with enough capacity, performance, and reliability to suit the demands of storage consumers without burdening them with unneeded low-level information. Instead of a difficult-to-manage collection of physical storage devices, users see one or more logical volumes that are indistinguishable from traditional physical discs. Consumers are unaware that logical or virtual units are involved; they only see capacity that fulfills the application's requirements. The virtualization layer is in charge of translating I/O requests to logical volumes onto physical storage. Simply said, block virtualization transforms physical discs into virtual storage devices that are as big, fast, and available (resilient) as storage customers want.

Storage customers' capacity, performance, and reliability needs are addressed by combining one or more of block virtualization's several capabilities:

- If storage users need more disc space, new volumes are created or existing logical volumes are expanded. Other free physical disc resources are added in the background without interfering with data access. Of course, the opposite is conceivable. Several smaller logical volumes may be produced from a single big physical disc; this is known as "slicing," and it can be highly beneficial when the number of LUNs that an underlying array can provide is limited.
- If storage users want more performance, the easiest solution is to stripe the data over many drives or even numerous arrays. This results in an improvement in throughput that is roughly proportionate to the number of physical discs over which the data is striped. Striping isn't the only option for enhancing I/O performance provided by storage virtualization, as we'll learn later in this paper.

- If storage users want increased availability, they may use clustering, RAID, synchronous mirrored to another array(s), and/or asynchronous data replication across long distances. All of these functions may be implemented as block virtualization layer functions.

Configuration updates may be made online with numerous copies of data in various places. Storage users may continue to work since their data is still accessible even if an entire storage array or cluster of arrays is replaced.

Virtualization Based on a Host

This sort of virtualization is usually connected with logical volume managers, which may be found on almost any computer, from the desktop to the data centre, in varied degrees of complexity. Logical volume managers, like storage-based virtualization, are not often connected with SANs. Despite this, they remain the most common type of virtualization due to their history and the fact that direct attached storage (DAS) is still widely used.

The logical volume manager (LVM) is increasingly becoming a regular feature of the operating system, although more powerful third-party implementations are also extremely prevalent. The following are the most typical applications for host-based LVMs:

- A. Combining physical storage from many LUNs to create a single "superLUN" that the host OS perceives as a single disc drive
- B. Introducing software RAID and other sophisticated features like as snapshots and remote replication.
- C. Maintaining the health of disc resources managed by the operating system

The durability of host-based virtualization after years of usage in practise, as well as its openness to diverse storage systems, are significant benefits. Because of the close proximity to the file system, which is also on the host, these two components may be closely coupled for effective capacity management. Many LVMs enable you to expand or shrink volumes and file systems without having to restart programmes.

The disadvantage of using a host-based solution is that the LVM is server-centric, which means that storage provisioning must be done on each host, which is a time-consuming process in a big, complicated system. As a result, several manufacturers provide cluster volume managers in a homogenous server environment to simplify storage administration by enabling numerous servers that share access to common volumes to be controlled as one.

Virtualization of Storage (Subsystems)

Your storage arrays may have been doing storage virtualization for years without your knowledge. Block-level storage virtualization includes features such as RAID, snapshots, LUN masking, and mapping.

Storage-based virtualization approaches may be used in both SAN and DAS systems.

Storage-based virtualization is often not reliant on a certain kind of host, enabling the array to accommodate heterogeneous hosts without having to worry about host operating systems or applications that differ. Furthermore, since features such as caching may be tailored to the individual hardware, storage-based RAID systems provide the best performance in proportion to their hardware. The disadvantage of this technique is that storage virtualization operations are

often constrained to a single array; for example, the source volume used for a snapshot and the snapshot itself are kept on the same array, rendering the snapshot worthless in the event of hardware failure. Virtualization functionalities may extend over many arrays or a cluster of arrays or controllers in certain circumstances; however, these solutions are often limited to a single vendor implementation.

Host- and storage-based virtualization are often coupled, combining the flexibility of host-based LVMs with the performance of hardware-assisted RAID.

For example, the host-based LVM may build virtual volumes that span several disc arrays by using multiple RAID-5 LUNs. LVMs in the host may mirror (with striped mirrors or mirrored stripes) volumes across several arrays in addition to just striping across LUNs from various arrays. In most contexts, host-based LVMs are also utilised to offer alternative route fail-over since the host is the only device in the I/O chain that knows for certain if its I/O has finished. Similarly, LVMs may employ load balancing on storage access channels to improve performance.

Virtualization based on a network

The benefits of both host- and storage-based virtualization may be merged in a storage management layer that already exists inside the SAN fabric. This network-based virtualization is the most recent advancement in the area. Network-based virtualization has the ability to provide the groundwork for the automated storage management required to restrict and manage storage capacity development.

A network-based virtualization strategy enables data center-wide storage management and can support a completely heterogeneous SAN with a broad set of host systems and storage resources. Network-based virtualization is often done using "black-box" appliances in the SAN fabric and maybe some agent software installed on the host (the requirement for and extent of such host-based agents dependent on how virtualization is performed in the appliance). The appliance itself might range from a commercially available server platform to a specialised, customised hardware architecture. Typical network-based virtualization functions include:

- A. Combining many LUNs from one or more arrays into a single LUN before providing it to a host; and
- B. Slicing a single LUN from an array into smaller virtual LUNs to deliver to the hosts.
- C. Synchronous and asynchronous replication inside the SAN and over WAN connections
- D. Device security to guarantee that only specified hosts have access to a LUN.

Other services include caching, sophisticated volume management, storage on demand, and QoS operations, however their availability varies from vendor to vendor.

SAN appliances are available as proprietary solutions, although they are more usually found as ordinary Windows, Unix, and Linux servers with accompanying virtualization software. All of the manufacturers with products on the market today understand the necessity of avoiding single points of failure inside the SAN and have solutions that enable redundancy and failover.

Switch makers have lately launched intelligent switches with inbuilt virtualization intelligence. In many aspects, these virtualizing switches are similar to SAN appliances; the primary distinction is

that the intelligence is built into the switch rather than being housed in one or more specialised appliances.

Future of Storage Virtualization

Management Integration

Unified management, which integrates discovery, reporting, storage virtualization, and storage automation. In a heterogeneous world, intelligent storage provisioning via active management of all levels, from the application to the physical storage devices, is only possible in the long run with virtualization and open standards like SNIA SMI-S. Storage as a utility, which was predicted years ago when the first SANs arrived, might soon become a reality.

Services for Automatic Data Migration

Storage virtualization will develop as well, and automated data movement services will become more widespread. The reasons for this are both technological and economic. Technically, such services provide fault-tolerant, high-performance data access. If the virtualization intelligence detects that storage systems are no longer capable of meeting demands, the data is moved or duplicated to other, faster storage resources without the administrator or users knowing anything. SLAs for performance and availability may therefore be met.

It should be recognised from an economic standpoint that not all data is equally important to a corporation. That is, all data does not need to be kept and handled on costly online storage systems. Statistically, the older the data, the less often access is required. This is where data lifecycle management enters the picture. Intelligent data migration services, well ahead of existing HSM ideas, may be used to guarantee that less often used data resources are relocated to and maintained in less expensive storage systems with several hierarchical levels, such as SATA, tape robots, or long-term archives.

Volumes and File Systems for Data Centers

This pamphlet is largely on block virtualization. Block virtualization builds and allocates volumes with defined capacity, performance, and availability to individual servers. If the same volume is allocated to many servers that access it at the same time, certain steps must be taken to prevent data integrity issues. Implementing a cluster file system that restricts write access by the multiple servers and manages the file system's metadata is a good example of such a precaution. However, volume sharing and cluster file systems are now mostly seen in homogenous server settings. At the time this article was written, there was no market solution for allocating identical volumes to several heterogeneous servers while also managing access using a heterogeneous cluster file system. The various operating system features and their footprints in volumes and file systems must still be "virtualized" for this to be viable. Several businesses, however, are working hard to merge block and file system virtualization. Once this goal is met, there will be no obstacles to transparent data access, regardless of the application, server operating system, network, or storage system.

CHAPTER 10

BUSINESS CONTINUITY

Amita Kashyap, Assistant Professor,
School of Computer & System Sciences, Jaipur National University, Jaipur, India
Email Id-amita@jnujaipur.ac.in

Business continuity refers to an organization's capacity to keep vital operations running during and after a crisis. Business continuity planning sets risk management methods and procedures with the goal of preventing disruptions to mission-critical services and restoring full organisation operation as fast and easily as feasible. The most fundamental need for business continuity is to maintain critical operations operational during a crisis and to recover with as little downtime as possible. Natural catastrophes, fires, disease outbreaks, cyberattacks, and other external hazards are all included in a business continuity strategy. Business continuity is critical for firms of all sizes, but it may not be feasible for any but the biggest enterprises to sustain all services during a crisis. Many experts believe that the first stage in business continuity planning is determining which operations are critical and allocating the available funds appropriately. Administrators may implement failover solutions after critical components have been identified. Disk mirroring, for example, enables an organisation to keep up-to-date copies of data at geographically distributed sites other than the core data centre. This allows data access to remain uninterrupted even if one location is deactivated and prevents data loss.

Imperative of business continuity:

Business continuity is crucial at a time when downtime is unacceptable. Downtime may come from a multitude of places. Some hazards, such as cyberattacks and harsh weather, seem to be worsening. It is critical to have a business continuity strategy in place that accounts for any possible operational interruptions. During a crisis, the strategy should allow the organisation to function at a bare minimum. Business continuity aids an organization's resilience by allowing it to react rapidly to an interruption. Strong business continuity saves money, time, and the reputation of the organisation. A prolonged outage poses a financial, personal, and reputational danger. Business continuity necessitates a business taking a look at itself, analysing possible areas of vulnerability, and gathering essential information such as contact lists and system technical diagrams – that may be beneficial outside of catastrophe scenarios. A company may strengthen its communication, technology, and resilience by implementing business continuity planning. Business continuity may also be required for legal or regulatory reasons. It's critical to understand which rules influence a certain organization's business continuity, especially in an age of rising regulation. Business continuity is a proactive approach to ensuring mission-critical activities continue in the event of an interruption. A thorough plan contains contact information, procedures for dealing with a range of problems, and instructions for when to utilise the document. Business continuity includes specific rules for what an organisation must do to keep operations running. When the time comes for a response, there should be no doubt on how to proceed with business procedures. Customers, workers, and the firm are all possibly at risk.

Different degrees of responsiveness are required for proper company continuity. Because not everything is mission-essential, it's crucial to choose what must remain operational and what can wait till later. It is critical to be open and honest about recovery time and recovery point goals. The whole company, from high management on down, is involved in the process. Although IT may be in charge of business continuity, it is critical to get management support and transmit critical information to the whole company. Another crucial area of cooperation is with the security team; although the two groups often function independently, a business may benefit greatly from exchanging information between both departments. At the absolute least, everyone should understand the fundamental processes for how the business intends to react.

Three essential elements of a business continuity plan

A business continuity strategy consists of three major components: resilience, recovery, and contingency.

- An organization's resilience may be increased by designing vital services and infrastructures with multiple catastrophe scenarios in mind, such as personnel rotations, data redundancy, and keeping a surplus of capacity. Assuring resilience against various situations may also assist firms in maintaining key services on and off site without interruption.
- It is critical to recover quickly after a catastrophe in order to resume company functioning. Setting recovery time targets for various systems, networks, or applications might assist in prioritising which pieces must be restored first. Other recovery options include resource inventories, agreements with third parties to take over firm operations, and the use of adapted premises for mission-critical services.
- A contingency plan includes procedures for a number of external eventualities as well as a chain of command that distributes duties inside the business. These obligations may include replacing hardware, leasing emergency office space, assessing damage, and hiring third-party providers for help.

Disaster recovery vs. business continuity

Disaster recovery planning, like business continuity planning, defines an organization's planned tactics for post-failure processes. A disaster recovery plan, on the other hand, is only a subset of business continuity planning. Catastrophe recovery plans are mostly data driven, focusing on storing data in a form that allows for easier access after a disaster. Business continuity considers this, but also focuses on the risk management, monitoring, and planning required for a business to remain functioning during an interruption.

Development of business continuity

Starting the planning project for business continuity is the first step. The stages of business impact analysis (BIA) and risk assessment are critical in acquiring information for the strategy. A BIA may expose any potential flaws as well as the effects of a calamity on multiple departments. The BIA report advises a company on which operations and systems should be prioritised in a business continuity strategy. A risk assessment detects possible threats to a company, such as natural catastrophes, cyberattacks, or technological failures. Risks may have an impact on employees, customers, building operations, and the company's brand. The evaluation also specifies who or what a risk might hurt, as well as the likelihood of the dangers. The BIA and risk assessment

complement each other. The BIA explains the probable consequences of the interruptions described in the risk assessment.

Management of business continuity

It is critical to identify who will be in charge of business continuity. It might be one individual for a small firm or a whole team for a bigger enterprise. Software for business continuity management is another alternative. Software, whether on-premises or in the cloud, aids in conducting BIAs, creating and updating strategies, and identifying areas of risk. Business continuity is a constantly changing process. As a result, a company's business continuity plan should not be left on the shelf. The organisation should reach out to as many individuals as possible. Implementing business continuity isn't only for times of crisis; the corporation should conduct training exercises so staff know what to do if there's a real interruption. Testing for business continuity is vital to its success. It's impossible to determine whether a strategy will work until it's been tried. A business continuity test may be as basic as a tabletop exercise in which employees debate what would happen in the event of an emergency. A comprehensive emergency scenario is included in more stringent testing. To properly simulate a crisis, a company may organise the test ahead of time or do it on short notice. After completing a test, the company should examine the results and change the strategy appropriately. Some portions of the plan are expected to work successfully, while other measures may need to be adjusted. A regular testing schedule is beneficial, particularly if the company's activities and personnel change often. Comprehensive business continuity is tested, reviewed, and updated on a regular basis.

The Institute for Business Continuity

The Business Continuity Institute (BCI) is a worldwide professional organisation that offers business continuity and organisational resilience education, research, professional accreditation, certification, networking opportunities, leadership, and direction. The BCI, located in the United Kingdom, was founded in 1994 and has over 8,000 members in more than 100 countries, both public and private. The BCI's products and services are offered to business continuity specialists and individuals interested in the subject. The BCI's aims and efforts include establishing business continuity standards, disseminating best practices in business continuity, educating and certifying BC professionals, increasing the value of the BC profession, and building the business case for business continuity. Among the numerous published tools available from the institution is its Good Practice Guidelines, which provide assistance for identifying business continuity initiatives that might support strategic planning. Membership in the BCI is an internationally recognized status; certification confirms a member's expertise in business continuity management.

Access to Information

The capacity of the infrastructure to function according to business expectations within its stipulated period of operation is referred to as information availability (IA). People (workers, consumers, suppliers, and partners) may access information whenever they need it if it is available. The dependability, accessibility, and timeliness of information may be used to describe its availability.

Reliability: Reliability is a component's capacity to work without failure under given circumstances for a set period of time.

Accessibility: The condition in which the needed information is available at the correct time and place to the right user. The length of time that the system is available is referred to as system uptime; when it is not accessible, it is referred to as system downtime.

Timeliness: Specifies the precise instant or time frame (a certain time of day, week, month, and/or year) during which information must be available. For example, if an application requires internet access between 8:00 a.m. and 10:00 p.m. every day, any delays to data availability outside of this time period are not deemed to influence timeliness.

Factors Contributing to Information Unavailability

Data is unavailable due to a variety of planned and unforeseen occurrences. Installation/integration/maintenance of new hardware, software updates or patches, backups, application and data restorations, facility operations (renovation and construction), and refresh/migration of testing to the production environment are all planned outages. Unplanned outages might occur as a result of database corruption, component failure, or human mistake. Natural or man-made catastrophes such as flood, fire, earthquake, and pollution are another sort of occurrence that might result in data loss. The bulk of outages are scheduled, as seen in Figure 11-1. Although planned outages are known and arranged, they nonetheless result in data becoming inaccessible. Statistically, fewer than 1% will be the outcome of an unexpected calamity.

Measuring the Availability of Information

The availability of information is dependent on the availability of a data center's hardware and software components. Failure of these components may cause information availability to be disrupted. A failure occurs when a component's capacity to execute a needed function is lost. The capability of the component may be restored by executing an external remedial action, such as a manual reboot, repair, or replacement of the failed component (s). Repair is the process of restoring a component to a state that allows it to fulfil a specific function within a defined time frame utilising processes and resources. Proactive risk analysis as part of the BC planning process takes into account the component failure rate and average repair time, as assessed by MTBF and MTTR:

Mean Time Between Failure (MTBF): The average time for a system or component to complete regular operations between failures.

Mean Time To Repair (MTTR): This is the average amount of time needed to repair a faulty component. When calculating MTTR, it is assumed that the problem that caused the failure has been appropriately detected and that all necessary parts and staff are accessible. It should be noted that a fault is a physical flaw at the component level that might lead to data loss. The time necessary to notice the defect, deploy the maintenance crew, diagnose the fault, get spare parts, repair, test, and restore normal operations is included in the MTTR.

Catastrophe Recovery as a Service (DRaaS): In the case of a disaster, a disaster recovery as a service (DRaaS) provider relocates an organization's computer processing to their own cloud infrastructure. Businesses pay for this service through subscription or pay-per-use. One benefit of DRaaS is that organisations may continue to function normally from the vendor's location even if their own servers fail. Choosing a local DRaaS supplier will result in lower latency; but, if the vendor's servers are too near to the disaster site, their own servers may be impacted by the same catastrophe.

Physical Disaster Recovery Technologies: Except for cyber assaults, physical disaster recovery tools may help lessen the consequences of some kinds of calamities. Fire suppression technologies to enable data and computer equipment survive a fire and a backup power source to support companies during short-term power outages are two physical aspects that may aid with business continuity.

Virtualization: One of the most difficult aspects of a business continuity plan is backing up an IT system. Virtualization is one of the few methods for creating a functioning clone of a company's whole computer environment. Businesses may also automate certain disaster recovery operations using off-site virtual computers that are not impacted by real catastrophes, allowing everything to be restored more quickly. Frequent data and workload transfers are required for virtualization to be an effective disaster recovery solution. IT teams must have a clear and up-to-date view of how many virtual machines are active in a company at any one moment.

Business Continuity Planning Lifecycle

Business continuity is an essential component of any scalable operations strategy, but many firms may not understand how necessary it is until they experience their first major incident. Only then does business continuity management move to the forefront of planning activities, forcing stakeholders to reflect on what went wrong, why it went wrong, and if they can prevent it from occurring again or be better prepared if it does. The full lifespan of business continuity management starts long before an event occurs. Consider the following six crucial actions to prepare for and recover from catastrophes and critical occurrences, whether you want to start preparing now or update your present systems to ensure they match contemporary requirements.

Reduce Risk: Preparing your company to adapt to unanticipated occurrences requires proactive preparation. This stage is critical for identifying, understanding, and prioritising potential business continuity risks, as well as developing reaction and mitigation measures. Consider which incidents are more probable and risky, and use that information to drive your planning. If your workplace is in an earthquake-prone area, for example, earthquake preparation should be a top concern. Alternatively, if your workforce is fully remote, examine how your business continuity strategy may conflict with the equipment supplied to workers or applications installed on their devices. Mitigation is proactively minimizing the possibility of an occurrence as well as reducing the effect of an event if it occurs. While certain catastrophes, such as natural disasters, cannot be avoided, others, such as IT failures, may be avoided with suitable solutions and backups in place.

Prepare: Whatever the business continuity danger, you must react swiftly and efficiently once it happens. Communication amid disruptive situations is important to preserving your company's survival. Because you never know where essential individuals will be when an incident occurs, your communications must be strong and adaptable, with elements such as: Allow users to transmit and receive vital alerts through mobile devices. Email, SMS messages, ChatOps technologies like Slack or Microsoft Teams, voice recordings, push alerts, and even phone calls are all possibilities. Customizable: Integrate with your organisation chart and use each person's schedule and device preferences. Targeted: During significant events, people get overwhelmed, particularly if their phones or inboxes are overflowing with messages. Targeted alerts aid in reducing alert fatigue.

Respond: Communication is crucial throughout the reaction phase, and response teams should determine the most suitable protective action for each threat to guarantee employee safety. During and after an emergency, the communications process should contain norms and procedures for

alerting first responders, such as public emergency services, trained personnel, and management. The response team must assess employee well-being, offer status updates on difficulties, and keep management up to date on new developments. With a simple click on a mobile device, a top communications platform may allow these actions.

Resolve: A communications strategy is required to ensure that responders can communicate with one another. Interoperability of communication equipment, techniques, and systems is required. Prior to an event, create an integrated voice and data communications system, including equipment, systems, and protocols.

Recover: Recovering after an incident is a distinct process based on the nature of the event, but maintaining in touch with diverse groups of concerned individuals is critical regardless of what occurred. Customers, workers, regulators, vendors, suppliers, shareholders, emergency services, and other stakeholders must all be kept up to date during a business continuity event. How you handle crises may have a significant influence on how the public, media, consumers, and regulators view your firm. When the organisation is ready to resume regular business activities, messages must be sent out detailing the estimated reopening timeframes for affected facilities.

Resume: Because of social media, 24-hour news, and other rapid modes of communication, communicating clearly and simply with workers, consumers, and business partners has become more crucial. Social media, text messaging, phone calls, 800 numbers, and corporate websites might all be used as communication conduits. After a catastrophe, communicate with suppliers and vendors and request flexibility and understanding. They may donate vital gear or software, or they may be ready to set up alternate payment or delivery alternatives until your company gets back on its feet. Maintain constant contact with government agencies and regulatory authorities to secure permissions for resuming building occupancy or rebuilding the facility.

Business impact analysis: When deciding whether to continue with one's firm continuity process, it is critical to consider the influence the business has had not only on the workers but also on the environment in which it was founded. One should endeavour to do a thorough investigation to determine if the business's impact to the environment has been beneficial. If its contribution was well received by the public, one should not be hesitant to continue with the procedure.

Identification of vital systems and components: The crucial and most important systems in a company must be recognised in order for the business continuity process to avoid causing undesirable modifications and impacts on such components.

Business continuity planning and testing: This is an activity that comprises doing a study of all of a company's operations, interdependencies, and vulnerabilities in order to aid in determining priorities for the company and strategic planning for its recovery.

Risk assessment is yet another critical notion in business continuity. In this instance, one must seek out all of the best risk assessment professionals and actuaries to determine the risk associated in business continuity. With the assistance of such individuals, it is simple to establish if the continuity process will result in a loss or profit for one's organisation.

Continuity of operations: It is critical that the usual functioning of the firm is not disrupted during business continuity. As a result, all company activities must continue as normal without interruption.

Disaster recovery: Disaster recovery is yet another critical idea in business continuity. This is a notion aimed at guaranteeing that all information and data in a firm that may have been lost is restored for future use. Most companies and organisations depend substantially on some of this data, thus it is critical that it be retrieved so that it may be utilised in the future.

Succession planning is essentially the notion of deciding who will take over different positions in one's firm. In this scenario, among many other individuals, one could desire to have the finest substitute for one's financial manager.

High availability is a term that basically suggests that one's systems will always be available no matter what occurs. However, redundancy does not guarantee that a system will always be accessible. With the case of redundancy, the other component must be manually enabled when one fails, however in high availability, the system is deemed always available, hence there is no need for manual activation. High availability also implies that many distinct components may be operating together. In a network, for example, there may be several firewalls, wide area networks, switches, and routers all working together so that if one portion fails, the others can simply take up.

Redundancy: Redundancy is a notion that is primarily concerned with keeping things operating in one's organisation despite the absence of one critical component. One goal of redundancy is to keep things operating and up and running. With redundancy, one must ensure that all network components and resources are operational and that we can utilise all available resources. This signifies that one's organisation is still operating regularly and as usual. In this instance, one must guarantee that there is no hardware breakdown. In this instance, redundant servers or power supply may be used. In the event of a power loss, all of one's systems will continue to function normally since there is another power source accessible. With such redundancies, one may be certain that if one component fails, another is present and ready to take its place. It is also necessary to ensure that there are no software faults. This is made feasible by installing software that can notify one of a problem as soon as it occurs. In this situation, other software may be running on a separate portion of the network so that if one fails, other programme on the network can take over. Finally, make certain that there are no system faults. For example, if you want your network to run best, you need have redundant switches, firewalls, and routers. With such redundant mechanisms in place, we are certain that our systems will remain operational at all times.

Tolerance for flaws

Hardware: The idea of fault tolerance necessitates the use of redundant hardware components. One might choose to have several power supply and gadgets for usage.

RAID: The Redundant Array of Independent Drives is another option for having numerous discs. RAID may be installed on a single server, which means that even if one disc fails, the system will continue to function normally and no one will notice.

Clustering: To guarantee that resources on a server are available and functional, consider clustering servers together. That way, if a motherboard fails, the system gets disconnected, or a machine fails, the other systems in the cluster can keep everything functioning. Because the cluster computers can interact with one another, they can detect outages and take over resources to ensure that everyone is able to operate the services that they need.

Load balancing: Another key idea is load balancing, which is when all of one's systems are functioning at the same time in order to balance the load and not overwork a single system. If one is lost, the other may be readily switched to without interruption. Without load balancing, more resources on the original system may be required to keep everything working.

Servers: Server clustering is critical for ensuring that servers are constantly operational and offer high availability. One may pick active/active server clustering, in which all end users visit various servers in the system. The servers are always active and talking with one another. One may also select an active/passive server clustering, in which one server is active and the other is passive, waiting for a failure to pick up. Active/passive server clustering is less difficult to build than active/active server clustering.

Concepts for Disaster Recovery

Plans/policies for contingencies: When it comes to data and information backup, there are several ways that may be used. Backup technologies are one of the strategies that may be used. In this instance, standard cassettes and discs may be used. Although disc backup or optical media backup may be used in certain instances, this is the most cost-effective and economical method of backup. Using optical storage medium, enormous volumes of data may be stored in a little amount of space. There are other two approaches for backing up a database. The first is replication, which involves copying all data and information from one system to another. As a result, there is always an identical clone of data on the devices, so that if data on one system is destroyed, the same data may be obtained from another machine. The other technique of backup is online database backup. Email database backups not only give backup capabilities, but also legal and compliance needs for storing such data over time. This backup technique allows you to back up not only the whole email server, but also individual mailboxes and messages. This implies that if a user deletes a message from their mailbox, they may access the database and restore the lost message using recovery software. Snapshots may also be used. This is a software supplied by Windows that allows you to back up files that are open and being utilised by the Windows operating system via the Volume Shadow Copy Service. Even if the files are in use at the moment, an identical replica may be created. Another backup approach used in residential setups and big enterprises is the image copy, which involves making an identical duplicate of everything on the system and then copying it to an image file.

Backup execution/frequency: Backup execution is typically performed in order to provide the possibility to recover a file even if it is lost. A backup replicates all of one's data and generates a copy of it. A complete backup execution might be chosen. This is a kind of backup in which every file on your system is copied. One may also do an incremental backup, which backs up any files that have changed since the previous incremental backup. There is also the differential backup notion. This differs from incremental backups in that it produces a backup of all files that have changed since the previous complete backup. Each backup solution has advantages and disadvantages. Backup execution may also be approached in terms of data retention length. First, one might choose for a short-term backup. This suggests that files that have recently modified should be kept for version control. Long-term backup is another alternative in which files and information are kept for legal needs or corporate standards.

Cold site: A cold site backup facility is an empty structure that serves as a disaster recovery location. This indicates that there might be some cooling and backup systems, but no hardware. This indicates that if an emergency occurs, it is one's obligation to bring all necessary materials.

Because this is a facility with no data stored, one must be familiar with methods for bringing data to the site. Because there are no personnel in this disaster recovery centre, one may opt to bring their own team.

Hot site: A hot site is a recovery facility where everything is duplicated, all systems are operational, and a full replica of a data centre exists. In this situation, while purchasing hardware, one purchases copies of hardware from the hot site. It's the same as purchasing everything in pairs. Applications and software are regularly updated, resulting in an automated reproduction of every component. In the case of a hot site, switching resources from one site to the other in a short amount of time is relatively simple.

Warm site: A warm site differs from a cold site in that it is a place where all of one's equipment is present, but all of the hardware is kept in a separate room. This implies that if there is a crisis, people will arrive at the facility and begin placing stuff in racks since it is a recovery location with ample space. The most essential thing in this instance is to have all of the recovery data and software so that it can be loaded into the systems and the recovery procedure can begin.

Managing the Storage Infrastructure

Management of the storage infrastructure necessitates the use of clever procedures and equipment. This guarantees compliance with regulations, increased data safety and security, availability and performance of all storage infrastructure components, and centralized auditing. Additionally, it ensures resource consolidation and improved usage, reducing the need for irrational technology investment and assisting in the effective exploitation of already available resources. Various tasks are involved in managing the storage infrastructure, including managing availability, capacity, performance, and security. In order to maximize the return on investment, all of these factors must cooperate. The storage management environment has been drastically altered by virtualization technologies, which have also made it easier to manage the storage infrastructure. One of the key elements that serves as the foundation for controlling a storage infrastructure is monitoring. Monitoring offers data on the condition of various storage elements and information needed to carry out crucial management tasks.

Storage Infrastructure Management key Functions:

Storage infrastructure management performs two key functions

- A. **Infrastructure discovery:** Infrastructure components are inventoried through infrastructure discovery, which also offers details on the components' configuration, connectivity, functionalities, performance, capacity, availability, utilisation, and dependencies between the physical and virtual worlds. It offers the visibility required to manage and watch over the infrastructure's parts. A customised tool is used for discovery, and it frequently communicates with infrastructure components utilising their native APIs. It gathers data from the infrastructure elements through interaction. An independent programme that transmits discovered data to a management software, a discovery tool coupled with the software-defined infrastructure controller, or any of these combinations. When a change is made to the storage infrastructure, an orchestrator or an administrator may also start a discovery process.

B. Operations management: Operations management entails ongoing management tasks for the deployment of services and the storage infrastructure. It guarantees that the promised services and service standards are provided. Multiple management procedures are involved in operations management. To guarantee operational agility, operations management should ideally be automated. The majority of management operations can typically be automated by management solutions. The management procedures are presented along with these automated operations. Additionally, orchestration allows for the logical integration and sequencing of the automated management tool actions.

. The key functions of Storage Operations Management are

1. **Configuration Management** - Configuration management is in charge of keeping track of configuration item information (CI). In order to supply services, CIs must be managed, including services, process papers, infrastructure parts (hardware and software), personnel, and SLAs.
2. **Capacity management** - Capacity management guarantees that storage infrastructure resources are adequately available to deliver services and satisfy SLA requirements. Regardless of dynamic resource use and seasonal surges in storage need, it establishes the ideal quantity of storage needed to satisfy a service's needs. Without sacrificing service levels, it also maximizes the utilization of existing capacity and reduces spare and stranded capacity.
3. **Performance management** - Performance management makes sure that all infrastructure components are operating as efficiently as possible so that storage services can reach or surpass the necessary performance level. Specialized management tools gather, evaluate, and publish performance-related data, such as reaction time and component throughput. If a component performs as intended, it can be determined through the performance analysis. Additionally, these technologies proactively notify administrators of potential performance problems and may suggest a course of action to resolve them.
4. **Availability Management** – Establishing appropriate guidelines based on the specified availability levels of services is the responsibility of availability management. In order to achieve or surpass both present and future service availability requirements at a reasonable cost, the guideline provides the procedures and technical elements needed. A storage infrastructure's availability-related problems and problem areas are all identified by availability management.
5. **Incident Management** - An incident is an unanticipated occurrence, like an HBA failure or an application issue, which may interrupt services or decrease service quality. All incidents in a storage system must be found and documented via incident management. Investigative teams for incident management look into incidents that have been elevated by the incident management software or service desk. They offer methods to restore the services within the SLA's predetermined timeframe. Error-correction tasks are shifted to problem management if the support services are unable to identify and address an incident's underlying causes. The incident senior management in this instance offers a temporary fix (workaround) for the situation.

6. **Management of Problems:** When several episodes share one or more similar symptoms, a problem is identified. For the purpose of identifying issues with a storage infrastructure, problem management examines each incident and its history. It pinpoints the underlying root cause of a problem and offers the best remedy, including potential preventative measures. Despite being different management processes, incident and problem management use interconnected incident and problem management systems and require automated interaction with one another. These technologies could aid in tracking individual incidents, designating them as problems, and transferring them to problem management for additional inquiry.
 7. **Security Management:** Information security policies must be created by security management in order to guide the organization's approach to information security management. It establishes the security architecture, procedures, controls, tools, user obligations, and standards required to efficiently uphold the information security policy. It also makes sure that the necessary security procedures and controls are correctly put in place. Information in a storage infrastructure is kept secure, intact, and readily available thanks to security management. It stops security-related incidents or actions from taking place that could harm infrastructure elements, management procedures, information, and services. Additionally, it complies with internal and external regulatory or compliance requirements for information protection at reasonable or acceptable costs.
-

CHAPTER 11

BACKUP AND RECOVERY

Brijraj Singh Solanki, Assistant Professor,
School of Computer & System Sciences, Jaipur National University, Jaipur, India
Email Id- brijraj.solanki@jnujaipur.ac.in

A backup is a replica of the production data that is made and kept for the sole purpose of restoring lost or damaged data. Organizations must back up a growing amount of data due to rising commercial and regulatory expectations for data storage, retention, and availability. The need for reliable data backup and speedy restore grows throughout the company, which may be dispersed across various sites, making this work more difficult. Furthermore, businesses must perform backup with the least amount of resources and at the lowest possible cost. Organizations must make sure that the appropriate data is available at the appropriate time and location. For the backup and recovery solution to be implemented successfully, it is crucial to assess backup technologies, recovery options, and retention needs for data and applications. As needed by the business, the system must provide simple recovery and retrieval from backups and archives. Backup is the process of transferring physical or virtual files or databases to a backup place for preservation in the event of equipment failure or disaster. Backing up data is critical to a good disaster recovery strategy. Enterprises back up data that they believe is susceptible in the case of faulty software, data corruption, hardware failure, malicious hacking, human mistake, or other unanticipated occurrences. Backups take and synchronise a point-in-time snapshot, which is subsequently used to restore data to its original condition. Backup and recovery testing investigates an organization's data security and replication policies and technology. The purpose is to guarantee that data can be retrieved quickly and reliably whenever the need arises. File restoration is the process of recovering backed-up data files. Although the phrases data backup and data protection are sometimes used interchangeably, data protection incorporates the larger aims of business continuity, data security, information lifecycle management, and malware and computer virus avoidance.

The significance of data backup: Data backups are critical infrastructure components in every company because they protect against data loss. Backups allow you to recover deleted data or files that have been mistakenly overwritten. Furthermore, backups are often an organization's best choice for recovering from a ransom ware attack or a severe data loss catastrophe, such as a data center fire.

Critical databases or associated line-of-business applications are backed up. Predefined backup rules establish how often data is backed up and how many duplicate copies – known as replicas – are necessary, as well as service-level agreements (SLAs) that indicate how soon data must be recovered. Best practises recommend scheduling a complete data backup at least once a week, preferably on weekends or after business hours. In addition to weekly full backups, businesses generally plan a series of differential or incremental data backup tasks that backup just the data that has changed since the previous complete backup.

Backup storage media evolution

Key data is often backed up by enterprises to specialized backup disc appliances. Backup software automates the process of transferring data to disc appliances, which may be embedded into the appliances or operate on a separate server. Backup software manages data deduplication and other techniques that decrease the amount of physical space necessary to store data. Backup software also imposes restrictions governing how often certain data is backed up, how many copies are generated, and where backups are kept. Prior to the early 2000s, when disc became the primary backup medium, most enterprises utilised magnetic tape drive libraries to store data centre backups. Tape is still used today, although mostly for historical data that does not need to be recovered rapidly. Some businesses now use detachable external drives instead of tapes, but the essential notion of backing up data to removable media remains the same. Organizations were able to accomplish continuous data security thanks to disk-based backups. Previously, companies would generally generate a single nightly backup. Initially, all nightly backups were complete system backups. The backup files grew in size with time, yet the backup windows stayed the same size or even decreased. As a result, several firms were obliged to maintain nightly incremental backups.

Platforms for continuous data protection totally prevent these issues. The systems execute an initial complete backup to disc, followed by incremental backups every few minutes when new or updated data is produced or edited. These backups may safeguard both structured data (that is, data stored on a database server) and unstructured or file data. Back in the day, disc backup software was meant to operate on a separate server. This programme managed backups and wrote backup data to a storage array. These systems quickly gained popularity because they functioned as online backups, which meant that data could be backed up or restored on demand without the need to mount a tape. Despite the fact that some backup packages continue to utilise separate backup servers, backup companies are rapidly shifting to integrated data protection appliances. An integrated data appliance, at its most basic, is a file server supplied with HDDs and backup software. These plug-and-play data storage systems often offer automated functions such as disc capacity monitoring, extensible storage, and preloaded tape libraries.

Some backup manufacturers have also started to provide backup solutions based on hyper-converged systems. These systems are made up of groups of standardised servers that have been clustered together to conduct backup-related operations. One of the primary advantages of hyper-converged systems is their ease of scalability. A hyper-converged system's nodes each have their own integrated storage, computation, and network resources. Administrators may increase the backup capacity of the business by simply adding additional nodes to the cluster. Most disk-based backup solutions, whether hyper-converged or not, allow copies to be shifted from spinning storage to magnetic tape for long-term retention. Because of growing tape density and the introduction of the Linear Tape File System, magnetic tape systems are still in use.

Early disc backup systems were referred to as virtual tape libraries (VTLs) because they incorporated disc drives that functioned similarly to tape drives. Backup software designed to write data to tape may therefore consider disc as a real tape library. VTLs fell out of favour once backup software companies geared their products for disc rather than tape. Because of cost and durability problems, solid-state drives (SSDs) are seldom employed for data backup. SSDs are included by

certain storage suppliers as a caching or tiering technique for managing writes with disk-based arrays. This is particularly prevalent in hyper-converged systems. Data is first cached in flash storage before being written to disc. Flash drives may have some usage for backup when companies introduce SSDs with more capacity than disc drives.

For main storage, local backup vs. offline backup

Modern main storage systems have developed to provide more robust native data backup features. Advanced RAID protection schemes, limitless snapshots, and tools for replicating snapshots to secondary or even tertiary off-site backup are among the capabilities available. Despite these advancements, main storage-based backup is often more costly and lacks the indexing features of classic backup programs. Local backups save data copies on external hard disc drives (HDDs) or magnetic tape systems, which are often kept in or near an on-premises data center. The information is sent through a secure high-bandwidth network connection or business intranet. The ability to backup data behind a network firewall is one benefit of local backup. Local backup is also significantly faster and gives you more control over who has access to your data. Offline or cold backup is similar to local backup, however it is more often linked with database backup. Because the backup procedure happens when the database is removed from its network, an offline backup incurs downtime.

Cloud backup and storage

Off-site backup sends copies of data to a distant location, which may be a company's secondary data centre or a leased colocation facility. Off-site data backup is increasingly being equated to subscription-based cloud storage as a service, which offers low-cost, scalable capacity and removes the need for a client to acquire and maintain backup gear. Despite its increasing popularity, backup as a service (BaaS) requires customers to encrypt data and take additional precautions to ensure data integrity.

Cloud backup is classified as follows:

- A. **Free cloud storage:** Users send data to a cloud services provider, who charges a monthly membership fee depending on the amount of storage used. There are extra expenses for data entrance and outflow. The three major public cloud providers are AWS, Google Cloud, and Microsoft Azure. Smaller managed service providers host backups on own clouds as well as manage client backups on major public clouds.
- B. **Personal cloud storage:** Data is often backed up to many servers inside a company's firewall, typically between an on-premises data centre and a backup disaster recovery location. As a result, private cloud storage is also known as internal cloud storage.
- C. **Cloud storage that is hybrid:** A business employs both on-site and off-site storage. Typically, enterprises utilise public cloud storage sparingly for data archiving and long-term preservation. For speedier access to their most sensitive data, they employ private storage for local access and backup. Most backup solutions allow customers to backup local apps to a dedicated private cloud, thereby considering cloud-based data backup as an extension of their physical data centre. Catastrophe recovery as a service occurs when a mechanism allows apps to fail over in the event of a disaster and then fail back. Cloud-to-

cloud (C2C) data backup is an option that is gaining traction. C2C backup safeguards data on SaaS systems like Salesforce and Microsoft Office 365. This data is often only available in the cloud, yet SaaS suppliers sometimes demand exorbitant costs to retrieve data lost due to client fault. C2C backup copies SaaS data to another cloud, where it may be recovered if any data is lost.

PC and mobile device backup storage

Local backup from a computer's internal hard disc to an associated external hard drive or removable media, such as a thumb drive, are options for PC users. Consumers may also back up their data from smartphones and tablets to personal cloud storage, which is offered from providers such as Box, Carbonite, Dropbox, Google Drive, Microsoft OneDrive, and others. These services are often utilised to give a specific capacity for free, with the option for clients to buy extra storage as required. These consumer-based cloud products, unlike commercial cloud storage as a service, often do not provide the degree of data protection that enterprises demand.

The Microsoft Resilient File System (ReFS) is included into the Microsoft Windows Server OS to automatically identify and restore damaged data. While not exactly a data backup solution, Microsoft ReFS is designed to protect file system data against damage. VMware vSphere includes backup software for data protection, high availability, and replication. The VMware vStorage API for Data Protection (VADP) allows VMware or third-party backup applications to take complete and incremental backups of VMs in a secure manner. Backups are implemented by VADP using hypervisor-based snapshots. In addition to data backup, VMware vSphere live migration allows VMs to be transferred across platforms to reduce the impact of a disaster recovery event. VMware Virtual Volumes may also help with VM backup.

Backup kinds have been specified.

- A. A full backup creates a copy of the complete data collection. Although considered the most dependable backup option, conducting a complete backup is time-consuming and necessitates the use of several discs or tapes. Most businesses only do complete backups on a regular basis.
- B. Incremental backup, which backs up just the data that has changed since the previous complete backup, is an alternative to full backups. The disadvantage is that if an incremental-based data backup copy is utilised for recovery, a complete restoration takes longer.
- C. A differential backup replicates data that has changed since the last complete backup. This allows for a faster complete restoration by needing just the most recent full backup and the most recent differential backup. For example, if you generate a complete backup on Monday, the Tuesday backup will be equivalent to an incremental backup at that moment. The differential that has changed from Monday's complete backup would then be backed up in Wednesday's backup. The disadvantage is that the steady expansion of differential backups has a negative impact on your backup window. A differential backup creates a file by fusing an earlier full copy with one or more incremental copies made later. The assembled file is not a straight replica of any single current or previously generated file, but rather a composite of the original file and any later changes to that file.

- D. A variant of differential backup is synthetic full backup. The backup server creates an extra full copy based on the original full backup and data gathered from incremental copies in a synthetic full backup.
- E. Incremental-forever backups reduce the backup window while enabling quicker data recovery access. An incremental-forever backup captures the whole data set and then supplements it with incremental backups indefinitely. Delta differencing is another term for backing up just modified blocks. Full backups of data sets are often saved on the backup server, automating the restoration process.
- F. Reverse-incremental backups are modifications performed between two mirror instances. After the first full backup, each incremental backup adds any changes to the current full backup. This effectively creates a new synthetic full backup copy for each incremental update, while also allowing reverting to earlier full backups.
- G. Hot backup, also known as dynamic backup, is used for data that is accessible to users while the update is being performed. This strategy avoids user downtime and lost productivity. The danger with hot backup is that if the data is changed while the backup is running, the subsequent backup copy may not be accurate.

Techniques and technology that may be used to supplement data backup

Continuous data protection (CDP) refers to layers of related technologies that are aimed to improve data security. When a modification is made, a CDP-based storage system backs up all business data. Multiple copies of data may be made using CDP techniques. A built-in engine in many CDP systems replicates data from a main backup server to a secondary backup server and/or tape-based storage. A common design for CDP systems is disk-to-disk-to-tape backup. Unlike array-based vendor snapshots, which are taken each time new data is written to storage, near-continuous CDP takes backup snapshots at predetermined intervals. Data minimization reduces the size of your storage footprint. Data compression and data deduplication are the two main ways. These strategies may be used alone, although suppliers often mix them. Data reduction has an impact on backup windows and restoration timeframes. Disk cloning is the process of replicating the contents of a computer's hard disc, storing it as an image file, and transferring it to storage media. Disk cloning may be used for provisioning, system provisioning, system recovery, and restarting, as well as restoring a system to its original configuration. Erasure coding, also known as forward error correction, originated as a scalable replacement for classic RAID systems. Object storage is the most common use for erasure coding. To guarantee redundancy and robustness, RAID striped data writes over many drives, employing a parity drive. The technique fragments data and encodes it with other pieces of superfluous data. These encoded pieces are distributed over various storage medium, networks, and geographical locations. The linked pieces are utilised to rebuild corrupted data using an oversampling approach. Flat backup is a data protection strategy that involves moving a straight copy of a snapshot to low-cost storage without the use of typical backup software. The original snapshot preserves its natural format and location; if the original becomes unavailable or unsuitable, the flat backup clone is mounted. Mirroring stores data files on many computer servers to guarantee that consumers may access them. Data is written to both the local and distant discs at the same time in synchronous mirroring. Writes from local storage are not recognised until a confirmation from remote storage is received, guaranteeing that the two locations have the same data copy. Asynchronous local writes, on the other hand, are completed before confirmation from the remote server is delivered. Replication allows customers to choose

the number of replicas, or copies, of data necessary to maintain or restore business activities. Data replication transfers data from one place to another, giving an up-to-date copy to speed up disaster recovery. Recovery-in-place, often known as quick recovery, allows users to run a production application briefly from a backup VM instance, preserving data availability while the original VM is being restored. Mounting a real or virtual machine instance directly on a backup or media server may accelerate system-level recovery to minutes. Because backup servers are not designed for production workloads, recovering from a mounted image degrades performance. Storage snapshots record a series of reference markers on disc for a particular database, file, or storage volume. Users use the markers, also known as pointers, to recover data from a certain moment in time. Individual storage snapshots are instances, not entire backups, since they are derived from an underlying source volume. As a result, snapshots do not safeguard data against hardware failure. Snapshots are classified into three types: modified block, clones, and CDP. Snapshots initially surfaced as a storage array management tool. Virtualization introduced hypervisor-based snapshots. Backup software or even a virtual machine may do snapshots.

Copy data management, as well as file sync and sharing

Copy data management is only tangentially connected to backup (CDM). This is software that gives insight into the numerous data copies that an organisation may generate. It permits separate groups of users to operate from a single copy of the data. Although not strictly a backup solution, CDM helps businesses to manage data copies more effectively by detecting redundant or underused copies, hence decreasing backup storage space and backup windows. Employees' mobile devices are protected by file sync and sharing technologies. These programmes essentially replicate changed user files from one mobile device to another. Although this safeguards the data files, it does not allow users to return to a previous point in time if the device fails.

Selection of Best Backup Option

When determining which sort of backup to utilise, you must examine many important factors. Enterprises often blend several data backup options, as mandated by data priority. The SLAs that apply to an application should govern a backup plan in terms of data access and availability, recovery time targets, and recovery point objectives. The adaptability of a backup programme, which should ensure that all data is backed up and enable replication and recovery while developing quick backup procedures, also influences backup selection.

Developing a backup strategy

Most firms develop a backup strategy to control the techniques and kinds of data protection they use and to guarantee that vital company data is consistently and frequently backed up. As the department responsible for securing all of the organization's sensitive data, the backup policy also establishes a checklist that IT can monitor and follow. A backup policy should contain a backup schedule. The rules are written so that others can back up and restore data if the primary backup administrator is not accessible. Data retention rules are often included in backup plans, particularly for firms in regulated sectors. Predefined data retention rules might result in the automatic destruction or transfer of data to other media once it has been stored for a certain length of time. Individual users, departments, and file types may all have their own data retention policies. A

backup strategy should include an initial full data backup as well as a series of differential or incremental data backups in between full backups. At least two complete backup copies should be kept, one of which should be kept off-site. Backup policies should prioritise recovery over backup, since backed-up data is useless if it cannot be retrieved when required. And recovery is critical to disaster recovery. Backup plans used to be primarily concerned with getting data to and from tape. Most data is now backed up to disc, and public clouds are often utilised as backup destinations. Because the process of transferring data to and from disc, cloud, and tape differs for each destination, the policy should reflect this. Backup operations may also change based on the application; for example, a database may need a different approach than a file server.

Recovery Considerations

When developing a backup system, RPO and RTO are critical factors to consider. RPO indicates the time period between two backups and sets the accepted limit of data loss for an organisation. In other words, the RPO controls the frequency of backups. For example, if application A has an RPO of one day, the data must be backed up at least once each day. A backup's retention duration is also determined by the RPO selected for operational recovery. Users of application "A," for example, may request that the application data be restored from its operational backup copy, which was created a month ago. This defines the backup's retention time. Based on operational recovery requirements, the RPO for application A might therefore vary from one day to one month. However, due to internal rules or external circumstances such as regulatory regulations, the business may opt to keep the backup for a longer length of time. If backups have short retention periods, it may be impossible to recover all of the data required for the intended recovery point since some data may be older than the retention term. Long retention periods may be configured for all backups, allowing any RPO to be met within the given retention periods. However, this necessitates a huge storage area, which increases the cost. As a result, it is critical to establish the retention time based on an analysis of all previous restore requests as well as the assigned budget. RTO refers to the amount of time required for the recovery procedure. To satisfy the required RTO, the company may opt to deploy a variety of backup solutions to reduce recovery time. RTO determines the kind of backup media that should be utilised in a backup system. Recovery from multiplexed data streams on tape, for example, takes longer than recovery from unmultiplexed tapes. Because of recovery limits, organisations conduct more complete backups than they really need. Cumulative and incremental backups are predicated on a complete backup. Several tapes are required to properly restore the system while recovering from tape media. Recovery may be accomplished with a lower RTO and fewer tapes when using a complete backup.

Back Up methods

When it comes to backup, there are several solutions available. The process of deciding which backup technique to use might be daunting for some. The main aim is to keep your data secure, but how you do so varies greatly depending on the backup strategy you choose. This article should provide you with a complete grasp of the various server backup techniques, allowing you to make an informed choice about how to configure your backup tasks based on your specific requirements as they relate to the kinds of backup available.

Complete Backup: A complete backup is the most basic kind of backup, including all of the folders and files that you choose to backup. Because it only needs access to a single backup file set, a complete backup is also the simplest sort of backup to recover. Full backups are often saved in a compressed, proprietary format that can only be restored using the programme that produced the backup. Full backups are often conducted as the first backup, followed by differential or incremental backups. Full backups are often bigger in size and need more storage space since they include all of the files and folders that were chosen for the backup process rather than just the updated files. A complete backup of a virtual machine is a backup of the whole virtual machine. Because these backups might be rather substantial, it's critical to be aware of how much storage space is available on your storage destination device.

Backup in Increments: Because they only backup files that have been created or updated since the previous full or incremental backup, incremental backups save a significant amount of storage space. Because incremental backups are quicker, they need a considerably shorter backup window. Incremental backups are often used in tandem with complete backups. A typical backup plan would be to do a complete backup once a week, followed by a daily incremental backup the following day. For example, if you ran a Full backup on Monday, you may then perform incremental backup operations from Tuesday through Friday (if you are not working weekends). With this option, your Tuesday backup would only include new or updated files that were created since the Full backup on Monday. On Wednesday, an incremental backup would be performed, but this time it would only backup any new or updated files since the last incremental backup on Tuesday. This plan would be followed for the remainder of the work week, providing you 1 complete backup and 4 incremental backups. The cycle would resume the next week. Because it only keeps data since the previous backup - regardless of kind - incremental backup has a shorter backup window and requires less storage space. When it comes time to restore, incremental backups take longer because you must first restore your entire backup and then each consecutive incremental backup in order to get the precise file iteration you want. Due to the nature of incremental backups, they also tend to demand more computational power since they must compare each source file to the most recent full backup and then to each successive incremental backup to discover whether any changes were made to any of your files. If you choose this kind of data backup, you should consider rotating media devices so that you always have a separate, disconnected (and ideally offsite) backup device to recover from in the event of a virus or other calamity.

Backup Differential: Differential backups are a hybrid of full and incremental backups. A differential backup is simply a cumulative backup of all changes made since the last full backup. This implies that Differential backups are bigger than Incremental backups since they are a rolled-up version of all Incremental backups performed since the previous Full backup. Given the nature of Differential backup, you might configure your Differential backups to overwrite your previous Differential backup in order to conserve storage space. To save storage space, several applications make this the default setting for Differential backups. Differential backups, like incremental backups, need extra network bandwidth to compare current files to those that have previously been backed up in order to discover and save only modified files. Because each Differential backup is independent of the others, they are much quicker to recover than incremental backups. That is, just the Full backup and the chosen Differential backup are needed to recover a specific backup set.

Backup Images: An image-based backup, also known as disaster recovery, disc image, or system image backup, enables you to generate a complete disc backup of your whole system (or one or more partitions), including your operating system, apps, and data, rather than simply your files and folders. This form of backup is often stored as a single file known as an image. When you need to restore your complete system to a new machine after a catastrophe, a full system image backup is quite useful. When you build an image backup of your server, you may easily restore your complete server to a new server, even if the hardware is different. A Windows image backup provides a safe way to store a more secure picture of your whole system, which may be saved to several storage media, providing you with a backup of your backup. Image backups are created using disc imaging software and provide the quickest recovery option possible for restoring your complete system. Image-level backups may also be restored to a real or virtual system. A system Image backup may also be mounted to restore a specific file rather than the full system.

Job Duplication: Copy jobs are distinct from backup tasks in that they produce the identical set of files, stored in their native and uncompressed file format. These copy files are merely "copied" or moved to your selected storage destination device on a manual or scheduled basis, but they should never be considered backups or be used in lieu of regular backups. Copy jobs allow you to specify the directories and files you wish to replicate in their native file format. Copy jobs produce files that are exactly the same size as the original files, so be sure you have adequate storage space on the destination device before beginning the Windows file copy process. Copy jobs are useful when you want to produce a quick file copy to have a secondary copy set of your data that you can access at any time without the need for backup software.

Backup Procedure

A backup system is built on a client/server architecture that includes a backup server and many backup clients. The backup server handles backup operations and maintains the backup catalogue, which comprises backup process and backup metadata information. The backup server relies on backup clients to collect data for backup. The backup clients might be on the server or on another server, presumably to back up the data visible to that server. To carry out its functions, the backup server receives backup metadata from backup clients.

Figure depicts the backup procedure. The storage node is in charge of writing data to the backup device (a storage node is a host that handles backup devices in a backup environment). The storage node is often connected with the backup server and housed on the same physical platform. A backup device is directly connected to the host platform of the storage node. Because it links to the storage device, some backup architecture refers to the storage node as the media server. Because they may be used to combine backup servers, storage nodes play a significant role in backup planning. The backup procedure is depending on the backup server's rules, such as the time of day or the conclusion of an event. The backup server then starts the process by sending a request to a backup client (client-initiated backups are also possible). This request directs the backup client to deliver its metadata to the backup server and the data to the appropriate storage node. When the backup client receives this request, it transmits the information to the backup server. This information is saved on the backup server's metadata catalogue. The backup client also delivers the data to the storage node, which copies it on the storage device.




Figure 11.1: Backup Procedure

When all of the data has been backed up, the storage node disconnects from the backup device. Backup completion status is sent to the metadata catalogue by the backup server. Backup software also has comprehensive reporting features that are based on the backup catalogue and log files. This information might include the quantity of data backed up, the number of successful backups, the number of partial backups, and the sorts of problems that may have occurred. Depending on the backup program utilised, reports may be altered.

Types of Backups

Backups are frequently divided into three groups:

Full backups - Imagine this procedure as pushing all of the data held on a production system onto a backup system for safekeeping, similar to filling up an extra tyre at the gas station. Full backups safeguard all of the information on a single server, database, virtual machine (VM), or other network-connected data source. Depending on how much data needs to be preserved, these backups may take many hours or even days to complete. An information management solution needs to do fewer full backups the more up to date it is, and when it does, it does it more quickly.

Incremental backups - Visualize incremental backups as re-visiting the gas station periodically and adding a tiny bit more air, just in case, so that you are always prepared to change your tyre. Only new data since the last full incremental is captured by an incremental backup. However, a backup solution must first perform a full backup before beginning its first incremental backup. Then, based on the most recent incremental action, it can perform them automatically.

Differential backups - These add more air, similar to incremental backups, but the delta comes from the most recent complete backup rather than the most recent incremental. Consider the differences between this backup and the last time you ever inflated the tyre. Once more, this is only possible if a full backup has been done first. Organizations often set regulations about the amount of data and the frequency of incremental and differential backups.

Data Recovery Types

- A. **Granular recovery of files, folders, and objects (also known as file-level or object-level recovery):** This technique swiftly restores one or a small number of particular data sets from a large number of volumes.

- B. **Instant mass restoration** - Using this method, IT workers may quickly and efficiently restore hundreds of virtual machines (VMs) to any point in time, saving time and resources.
- C. **Volume recovery:** it is a technique used by teams to quickly restore an unlimited number of VMs at once, such as all of the VMs in an application group.
- D. **Virtual Machine Disk (VMDK) recovery:** With this procedure, all of a VM's data and applications are swiftly restored.
- E. **Bare machine recovery** - Restoring the entirety of the operating system (including all programmes, apps, and data) in a single step
- F. **Instant mounts of volumes** – Teams can save time by using a remote backup as a target to reestablish an entire volume to a Windows virtual machine.
- G. **Instant restores of VMs** – This procedure restores a significant number of VMs to any previous recovery moment with backup copies fully hydrated and immediately accessible.

The procedure of backing up your data in case of loss and establishing secure systems that enable you to retrieve your data as a result is known as backup and recovery of data. To provide access to computer data in the event of data loss or corruption, data backup entails the copying and preservation of computer data. Only data that has been backed up using a trustworthy backup device can be restored.

Any sane disaster recovery plan must include data backup since it is one type of disaster recovery. All of the data and settings for your company operating systems may not always be recoverable from backups. For example, different types of disaster recovery may be required for computer clusters, database servers, or active directory servers since backup and recovery efforts may not fully restore them.

These days, you don't need to archive your data on a local system's hard disc or external storage because you may back it up utilizing cloud storage.

Key terms

Better evaluate backup and disaster recovery options and impact your strategic decisions by understanding a few key phrases.

1. The length of time needed to resume regular business activities following an outage is known as the recovery time objective (RTO). You should think about how much time you're willing to lose and the effect that time will have on your bottom line when determining your RTO. The RTO may differ significantly between different business types. For instance, if a public library loses access to its catalogue system, it will probably be able to operate manually for a few days until the systems are repaired. • But even 10 minutes of downtime—and the ensuing income loss—would be intolerable if a significant online shop lost control of its inventory system.
2. The quantity of data you are willing to lose in the event of a disaster is known as the recovery point objective (RPO). In order to prevent data loss in the event of an outage, you might need to regularly copy data to a remote data center. Alternately, you can determine that losing data for five minutes or an hour is acceptable.

3. Automatically offloading duties to backup systems in a way that is seamless to users is the disaster recovery procedure known as failover. It's possible to fail over from your primary data centre to a backup location that has redundant systems ready to take over right away.
4. Failback is the process of going back to the old systems during a disaster recovery. You should be able to fail back flawlessly once the crisis is over and your major data center is back up and running.
5. The process of restoring involves moving backup data to your main system or data center. Restoration is typically seen as a backup step rather than a catastrophe recovery one.

A backup is a duplicate of one or more files made as a backup in case the original material is lost or becomes useless. In a backup environment, three fundamental topologies are used: direct connected backup, LAN-based backup, and SAN-based backup.

Direct-attached backup: A mixed topology is created by merging LAN and SAN topologies. A backup device is directly connected to the client. Through the LAN, just the metadata is delivered to the backup server. This setup clears backup traffic from the LAN. Sharing backup devices across many servers is an effective strategy. In this case, the client also serves as a storage node, writing data to the backup device.

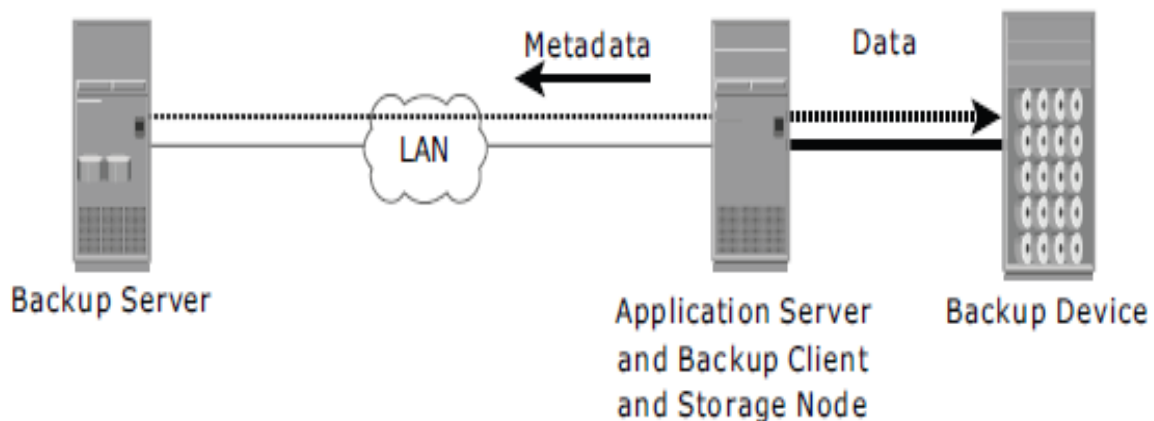


Figure 11.2: Direct-attached backup

Backup across a LAN: All servers are linked to the LAN in LAN-based backup, and all storage devices are directly connected to the storage node. The data to be backed up is sent across the LAN from the backup client (source) to the backup device (destination), which may impact network performance. This effect may be reduced by using a variety of steps, such as designing separate backup networks and deploying dedicated storage nodes for particular application servers.




Figure 11.3: Backup across a LAN

Backup to a SAN: SAN-based backup is sometimes referred to as LAN-free backup. When a backup device must be shared across clients, the SAN-based backup architecture is the best choice. The backup device and clients are connected to the SAN in this situation. In this example, clients read data from the SAN's mail servers and send it to a SAN-attached backup device. Backup data traffic is limited to the SAN, but backup metadata is sent over the LAN.




Figure 11.4: Backup to a SAN

Topological mix: the hybrid topology incorporates both LAN and SAN topologies. This architecture may be used for a variety of reasons, including cost, server placement, administrative overhead reduction, and performance factors.

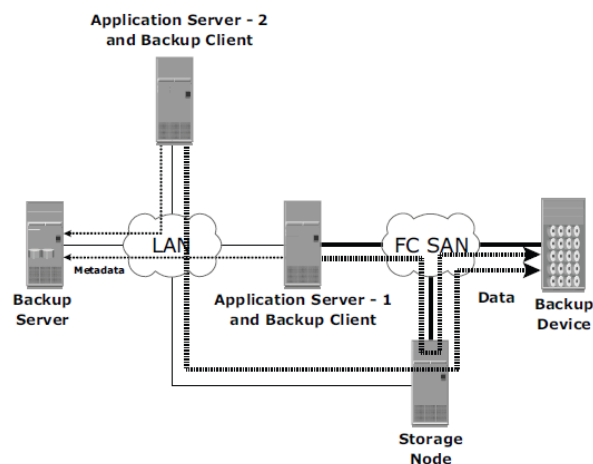


Figure 11.5: Topological mix

NAS with Cloud Backup

Cloud storage may be used to backup NAS file storage or to completely replace it. When deciding between them, there are three major variables to consider. These are the cost, security, and storage constraints:

Cost: The price of a NAS for a small company ranges from \$500 to \$1000 or more, depending on the functionality and storage capacity required. You'll have to pay for extra hard drives if you require more storage. Cloud file storage options, on the other hand, are based on a monthly cost. The more storage space you need, the more you will have to spend. 1TB of storage costs around \$95 per month. The cloud option eliminates the need for an initial investment in a NAS as well as the expenses involved with acquiring additional hard drives, but you will be charged a monthly storage fee.

Security: Because you host the files on your own hardware while using NAS, you have complete control over data security. NAS provides choices for data redundancy and protection, as well as data encryption and user access restrictions. When you utilise a cloud service, the cloud provider is responsible for the security of your data. When you use the cloud, you relieve yourself of data security, but when you give control of your data security to another party, there is always the possibility that your data may be stolen or harmed. For further redundancy and safety, some NAS suppliers offer the option of backing up your data to the cloud.

Storage constraints: NAS devices have storage constraints that are dictated by their architecture. Hard drives may be changed if extra storage is required. There is no need to be concerned about running out of storage space while using cloud storage. When extra storage space is required, you would increase your plan with your cloud provider.

Backup Options

There are several methods for backing up your documents. Choosing the finest solution may help you create the best data backup strategy for your requirements. Six of the most prevalent approaches or technologies are listed below:

Disposable Media: Backup data on removable media such as CDs, DVDs, newer Blu-Ray discs, or USB flash drives is an easy solution. This is useful in smaller contexts, but for greater data volumes, you'll need to backup to many drives, which might complicate recovery. Also, be sure to save your backups in a different area; otherwise, they may be lost in a catastrophe. Tape backups are also included in this category.

Redundancy: Back up a second hard drive that is a clone of the drive in a sensitive system at a precise moment in time, or you can put up a whole redundant system. For example, another email server on standby to back up your primary email server. Redundancy is a valuable tool, but it is difficult to maintain. It requires regular replication across cloned systems and is only beneficial in the event of a particular system failure unless the redundant systems are located at a faraway place.

External Hard Disk Drive: Use archive software to store changes to local files to a high-volume external hard disc in your network. Archive software enables you to recover data from external devices in a matter of minutes. However, when your data quantities expand, one external drive will no longer be enough, and your RPO will skyrocket. The use of an external drive needs its deployment on the local network, which is dangerous.

Hardware Devices: Many suppliers provide full backup appliances, which are commonly implemented as 19" rack-mounted devices. Backup appliances have a huge storage space and backup software pre-installed. Install backup agents on the systems to be backed up, specify your backup schedule and policy, and the data begins to flow to the backup device. As with previous solutions, attempt to isolate the backup device from the local network and, if feasible, locate it in a distant location. Backup software is more difficult to setup and manage than hardware appliances, but it provides more flexibility. They let you to choose which systems and data to backup, distribute backups to the storage device of your choosing, and manage the backup process automatically.

Cloud Backup Solutions

Many manufacturers and cloud providers provide Backup as a Service (BaaS) solutions, which allow you to transfer local data to a public or private cloud and restore data from the cloud in the event of a catastrophe. BaaS systems are simple to use and offer the significant benefit of storing data in a distant place. However, if you use a public cloud, you must maintain compliance with applicable legislation and standards, and keep in mind that data storage costs in the cloud will be substantially greater over time than the cost of establishing identical storage on-premises.

Storage Security and Management

Security Framework for Storage

Accountability, confidentiality, integrity, and availability are the four main services of security that make up the fundamental security framework. This framework includes all security safeguards necessary to lessen risks to these four key security features:

- A. Accountability service:** Accounting for all activities that occur within the infrastructure of a data centre is referred to as an accountability service. The accountability service keeps a record of occurrences that can be later audited or tracked for security reasons.
- B. Confidentiality service:** Provides the necessary information secrecy and guarantees that only authorized users can access data using the confidentiality service. Typically, this service protects both data in transit (data sent over cables) and data at rest (data stored on a backup medium or in the cloud) and authenticates people who need access to information. To maintain its confidentiality, data can be encrypted both in transit and at rest. Confidentiality services incorporate traffic flow security protocols as part of the safety protocol in addition to preventing unauthorised users from accessing information. These security precautions typically involve hiding data transmission frequency, data volume, source and destination addresses. Integrity service: Assures that the data hasn't been changed. The service's goal is to identify and guard against unauthorised information alteration or deletion. Integrity services collaborate with accountability services to identify and authenticate users, much like confidentiality services do. Integrity services specify safeguards for both at-rest and in-transit data.
- C. Availability service:** This guarantees quick and dependable access to data for authorised users. Users can access the necessary computer systems, data, and applications by using these services. Additionally, communication systems used to send data between computers that may be located in various places employ availability services. This guarantees information accessibility in the event of a breakdown at one

specific site. The implementation of these services is required for both physical and electronic data.

Risk Triad:

Risk triad defines risk in terms of threats, assets, and vulnerabilities. Risk develops when a threat agent (an attacker) attempts to get access to resources by taking advantage of a known vulnerability. Organizations manage risks by concentrating on vulnerabilities because they are unable to completely eliminate the danger agents that could come from different sources and take different shapes to harm their assets. Organizations can implement defences to lessen the impact of a threat agent attack, hence decreasing susceptibility.

- A. Assets:** One of any organization's most valuable resources is information. Hardware, software, and the network infrastructure needed to retrieve this information are additional assets. Organizations must create a set of guidelines to preserve these assets and guarantee that the resources are accessible to legitimate users and reliable networks. These specifications apply to organisational policies, network infrastructure, and storage resources. When making plans for asset security, several things need to be taken into account. Security measures have two goals. The first goal is to make sure that authorised users may readily access the network. Additionally, it should be dependable and stable under a variety of environmental circumstances and consumption levels. Making it extremely challenging for potential attackers to enter and compromise the system is the second goal. These techniques ought to offer sufficient defence against unwanted access to resources, viruses, worms, Trojan horses, and other harmful software applications. To reduce the amount of potential vulnerability holes, security procedures should also encrypt crucial data and shut unnecessary services. Regular installation of operating system and other software updates must be made possible by the security measure. In addition, it must offer sufficient redundancy by replicating and mirroring the production data to avoid catastrophic data loss in the case of an unanticipated malfunction. Make sure that all users are aware of the rules governing network use in order for the safety system to operate effectively. Two factors can be used to assess a storage security methodology's performance. One, just a small portion of the value of the protected data should be spent on the system's implementation. Two, it should be more expensive for a possible attacker to breach the system in terms of resources and time than the value of the data that is being secured.
- B. Threats:** Threats are possible assaults that could be made against an IT infrastructure. There are two types of attacks: active and passive. Attacks that are passive in nature aim to breach the system's security. Information confidentiality is threatened by them. Data manipulation, DoS, and repudiation assaults are examples of current attacks. They endanger the availability and integrity of data. An unauthorised user tries to change data for malevolent purposes in a modification attack. Data in transit or at rest can both be the target of a modification attack. The integrity of the data is threatened by these attacks. Attacks that cause denial of service (DoS) prevent authorised users from using resources. In most cases, these assaults do not entail gaining access to or altering data on the computer system. Instead, they endanger the availability of data. One type of DoS attack is the deliberate flooding of a network or website to deny lawful access to authorised users. Rejection is an assault on the veracity of the information. It makes an effort to spread false

information by pretending to be someone else or denying that an action or transaction took place.

- C. Vulnerability:** Potential attacks are most likely to occur along routes that lead to information. Different access points that offer varied access levels to the data and resources may be present along each of these paths. Adequate security measures must be put in place at each access point along an access path. Defence in depth refers to the implementation of security measures at each access point along each access path. Defence in depth advises securing all entry points in a setting. This lessens vulnerability to an adversary who could access storage resources by getting beyond insufficient security measures put in place at the weak single point of entry. The safety of information assets may be compromised by such an attack. When determining how vulnerable an environment is to security threats, three elements should be taken into account: attack surface, attack vector, and work factor. The many entry points that an attacker can utilise to start an assault are referred to as the attack surface. Every part of a storage network has the potential to be vulnerable. An attacker can exploit any of the external interfaces that the component supports, including the hardware interfaces, the protocols it supports, and the management and administrative interfaces, to carry out a variety of attacks.

The attack surface for the attacker is formed by these interfaces. If enabled, even inactive network services may join the attack surface. A step or set of procedures required to carry out an assault is known as an attack vector. For instance, a hacker could use a flaw in the administration interface to launch a snoop attack, changing the storage device's configuration to provide access to the traffic from additional hosts. Data in transit can be watched using this redirected traffic. Work factor is a term used to describe the time and effort needed to use an attack vector. For instance, when attempting to retrieve sensitive data, attackers take into account the time and effort needed to carry out a database assault. This could involve creating SQL queries, figuring out the database schema, and identifying privileged accounts.

CHAPTER 12

LOCAL REPLICATION

Brijraj Singh Solanki, Assistant Professor,
School of Computer & System Sciences, Jaipur National University, Jaipur, India
Email Id- brijraj.solanki@jnujaipur.ac.in

The process of producing an identical duplicate of data is known as replication. This procedure also includes the creation and management of duplicate database versions. The replication procedure not only makes precise copies of the database data, but it also synchronises a collection of clones such that changes made to one copy are reflected in all of the others. The beauty of replication is that it enables several users to work on their local copy of the data while keeping the database updated and giving the users the illusion that they are working on a single, centralised database.

Local Replication

- A. Data replication inside the same array or data centre.
- B. Local replication uses the power of Infortrend storage technology to provide snapshot and quantity copy/mirror capabilities to aid users in effectively securing information.
- C. Examples of host-based local replication technologies include file system replication and LVM-based replication. Full-volume mirroring, pointer-based full-volume replication, and pointer-based virtual replication are three ways for storage array-based replication.
- D. The two primary methods used for local replication are host-based and storage-based replications.

Replications depending on hosts

- 1. In host-based replication, the local replication process is handled by logical volume managers (LVMs) or file systems.
- 2. Examples of host-based local replication include LVM-based replication and file system (FS) snapshots.

Replication on the basis of LVM

- a) The logical volume manager is in charge of generating and managing the host-level logical volume in LVM-based replication.
- b) Physical volumes (physical disc), volume groups, and logical volumes are the three components of an LVM.
- c) A volume group is formed by combining one or more physical volumes. Within a volume group, logical volumes are generated. There may be numerous logical volumes in a volume group.
- d) In LVM-based replication, as illustrated in Figure, each logical partition in a logical volume is mapped to two physical partitions on two distinct physical discs.

- e) The LVM device driver writes to the two physical partitions when an application writes to a logical partition. This is also referred to as LVM mirroring. Mirrors may be divided and the data stored inside them accessible individually. LVM mirrors may be dynamically added or withdrawn.

The Benefits of LVM-Based Replication

- a) LVM-based replication does not need a vendor-specific storage system. LVM is often included with the operating system, therefore no extra licencing is necessary to implement LVM mirroring.
- b) Array-Based Storage Replication
- c) The array operating environment executes the local replication procedure in storage array-based local replication.
- d) The replication procedure does not require host resources such as CPU and memory. As a result, the host is not burdened by replication processes. An alternative host may access the replica for any business activity.
- e) In this replication, the needed number of replica devices on the same array are chosen, and data is copied between source-replica pairings.
- f) The illustration depicts storage array-based local replication, in which the source and destination are in the same array but are accessible by distinct hosts.
- g) Full-volume mirroring, pointer-based full-volume replication, and pointer-based virtual replication are all types of storage array-based local replication.

Full-Scale Mirroring

- a) Full-volume mirroring involves attaching the target to the source and establishing it as a mirror of the source (Figure [a]). Data from the source is transferred to the destination. Any changes to the source are mirrored on the target. After all of the data has been duplicated and both the source and the target have the same data, the target may be regarded a mirror of the source.
- b) While the target is connected to the source and synchronization is taking place, it is inaccessible to any other host. The production host, on the other hand, has access to the source.
- c) Once the synchronization is complete, the target may be disconnected from the source and made accessible for BC activities. When the target is separated from the source, full-volume mirroring

Technologies for Local Replication

The two primary methods used for local replication are host-based and storage-based replications. Host-based local replication technology includes file system replication and LVM-based replication. Full-volume mirroring, pointer-based full-volume replication, and pointer-based virtual replication are three ways for storage array-based replication.

Local Replication on the Host

Local replication is handled by logical volume managers (LVMs) or file systems in host-based replication. Host-based local replication includes LVM-based replication and file system (FS) snapshots. The logical volume manager is in charge of establishing and managing the host-level logical volume in LVM-based replication. Physical volumes (physical disc), volume groups, and

logical volumes are the three components of an LVM. A volume group is formed by combining one or more physical volumes. Within a volume group, logical volumes are generated. There may be numerous logical volumes in a volume group. Each logical partition in a logical volume is mapped to two physical partitions on two distinct physical volumes in LVM-based replication, as illustrated in Figure 1. The LVM device driver writes an application write to a logical partition to the two physical partitions. This is also referred to as LVM mirroring. Mirrors may be divided and the data stored inside them accessible individually. LVM mirrors may be dynamically added or withdrawn.




Figure 12.1: Local Replication on the Host

Data consistency: Data consistency is the process of keeping information consistent as it flows across a network and across computer applications. Data consistency is classified into three types: point in productive maintenance, transactional consistency, and application continuity. The best strategy to guarantee that data is not lost or damaged while it travels over a computer network is to ensure that all three parts of data consistency are handled. There are no assurances that any piece of information on the system is consistent throughout the whole computer network in the absence of data consistency.

Introduction to Data Replication Techniques

To reduce the risk of business interruption, each firm must safeguard its vital data across physical, virtual, and cloud environments. In the event of a catastrophe, quick data recovery and restart capabilities are necessary to maintain company continuity. As a result, replication is one method of ensuring company continuity in the event of a catastrophe.

Data Replication Techniques

Replication, as defined by EMC, is the process of generating an identical duplicate (replica) of existing data. These copies are used to restore and resume operations in the event of data loss due to catastrophes or malfunctions. For example, if a production server fails, replica data may be utilised to continue production processes with little downtime. Other servers may also be allocated replicas to execute other business processes such as backup, reporting, and testing.

Replication: Data may be copied to one or more locations based on the needs of the organisation. Data may be duplicated inside a data centre, across data centres, from a data centre to a cloud, or between clouds, for example. Organizations have policy-based replication automation in a software-defined data centre architecture. The number of copies to be made and the place where

the data should be kept are determined by policy-based automated replications. Typically, cloud service providers have many data centres across the globe and may provide clients the option of selecting the place to which the data is to be duplicated. A production server often accesses data from one or more LUNs on storage devices. These LUNs are referred to as source LUNs, production LUNs, or just the source. The target LUN, or simply the target or replica, is the LUN to which the production data is duplicated.

Data Replication Characteristics

Recoverability enables the restoration of data from replicas to the source in the event of data loss. Restart ability enables the use of replicas to resume business processes. The replica must be consistent with the source in order to be used for recovery and restart operations. The data on the replica will be an exact copy of the production at a specified moment. The data on the replica is always in sync with the production data. The goal of any continuous replication is to lower the RPO to zero or close to zero. Consistency is a crucial need for replica device usage. In the case of file systems, consistency may be accomplished by either unmounting the file system or keeping the file system online by flushing server buffers before generating a replica. Prior to creating the replica, the server's memory buffers must be flushed to discs to maintain data consistency. The data on the replica will not include the information that was buffered in the server if the memory buffers are not flushed to disc. Similarly, consistency in databases may be accomplished by either taking the database offline to create a consistent replica or by keeping the database online. If the database is online, it is accessible for I/O activities, and transactions to the database continually update the data. When a database is replicated while it is live, all changes made to the database must be applied to the replica in order for it to remain consistent.

Uses for Replication

Replication may be utilised as an alternate backup source. Data is read from the production LUNs and written to the backup device during typical backup processes. This puts extra strain on the production infrastructure since production LUNs are engaged in both production operations and data servicing for backup operations. To circumvent this problem, a replication from the production LUN may be established and utilised as a source for backup operations. This reduces the backup I/O burden on production LUNs. It is useful for quick recovery and restart. Replicas may be taken at short, regular intervals for crucial applications. This enables for simple and quick data recovery. If the production LUN fails completely, the replication solution allows you to resume the production activity on the replica to decrease RTO. It may be used to create reports. Running reports on the replicas decreases the I/O load on the production device significantly. It may be used to put novel business models to the test. Replicas are often used to test new apps or update existing ones. For example, a company may utilise the replica to test a production application update; if the test is successful, the upgrade may be executed in the production environment. Data migration is another use for a replica. Data migrations are undertaken for a variety of reasons, such as upgrading an application from a lesser capacity LUN to a bigger capacity LUN.

Replication Types

Local and distant replication are the two types of replication.

Local replication is the process of duplicating data inside the same storage system or data centre. Local replicas aid in the restoration of data in the case of data loss or allow the programme to be restarted promptly to maintain business continuity. Local replication is possible at the server, storage, and network levels.

Data replication to distant sites is referred to as remote replication. Remote replication assists enterprises in mitigating the hazards of regional outages caused by natural or man-made calamities. During a catastrophe, services may be relocated (failed over) to a distant site to ensure that company operations continue uninterrupted. Remote replication also enables businesses to copy their data to the cloud for disaster recovery purposes. Data may be copied synchronously or asynchronously through remote replication, which can also be done at the server, storage, and network levels.

Considerations for Restore and Restart

Data may be restored to production devices using local replicas. Alternatively, programmes may be relaunched using the replicas' consistent point-in-time copies of the data. In the case of logical corruption of production devices that is, the devices are accessible but the data on them is invalid a replica may be utilised to restore data to them. Accidental deletion of information (tables or entries in a database) is an example of logical corruption, as is erroneous data input and wrong updating of existing information. Restoring from a replica is gradual and has a very short RTO. In rare cases, applications may be restarted on the production devices before the data copy is completed. Access to the production and replica devices should be disabled prior to the restoration procedure. Physical problems, such as production server or physical drive failure, may also cause production devices to become unavailable. Applications may then be resumed using data from the most recent replica. If the production server fails, the most recent information from the replica devices may be restored to the production devices after the problem has been fixed. Applications may continue to execute on replica devices if the production device(s) fail. A fresh PIT copy of the replica devices may be generated, or the replica devices' most recent information can be restored to a new set of production devices. Access to the replica devices should be disabled before resuming apps that use them. To defend against subsequent failures, a "Gold Duplicate" (another copy of the replica device) of the replica device should be generated to maintain a copy of data in the case of replica device failure or corruption. Full-volume replicas may be restored to the original source devices or to a new set of source devices (both full-volume mirrors and pointer-based in Full Copy mode). Restorations to the original source devices may be incremental, but restores to a new set of devices need a full-volume copy. Access to data on the replica is reliant on the health and accessibility of the original source volumes in pointer-based virtual and pointer-based full-volume replication in CoFA mode. These copies cannot be utilised for a restoration or restart if the original source disc is unreachable for any reason.

Questionnaire for Practices

1. What are the key activities include in managing storage infrastructure?
 2. What are the components of storage system environment?
 3. Explain the Raid.
 4. What are the components of an intelligent storage system?
 5. What is DAS storage and how it works?
 6. What are the different types of storage interfaces?
 7. What are the Configuration of SAN?
 8. What are the communication protocols that are used between a NAS and clients?
 9. Explain the components of iSCSI protocol.
 10. What are the advantages of storage virtualization?
 11. What is content-addressed storage (CAS) and how does it work?
 12. Why we need Business Continuity Plan?
-