Storage Concepts

Introduction

Welcome and Introductions

- Student Introductions
  - Name
  - Position
  - Experience
  - Your expectations

Course Description

This 2 day instructor-led course provides a comprehensive introduction to storage technology concepts, terminology and technologies of today's storage industry. The course examines the need for storage solutions to manage and optimize an IT infrastructure to meet business requirements.

The course also examines the major components of a storage system, common storage architectures and the various means of connecting storage elements. It compares network attached storage (NAS) and storage area network (SAN) implementations and data protection issues. It provides detail on industry-defined tiered storage, virtualization, and storage management strategies.

Prerequisites

- Understanding of basic computer concepts
- Experience working with PCs or servers (Windows or UNIX)

Course Objectives

- By completing the course, you will gain an understanding of:
  - Storage industry concepts and technologies
  - Industry-defined tiered storage, virtualization, and storage management strategies

Course Topics

Modules	Activities
1. Introduction to Data Management and Storage Systems 1a. Overview of Storage Concepts 2. Storage Components and technologies	Learning activities appear throughout the course.
3. Business Continuity and Replication 4. Virtualization of Storage Systems 5. Archiving and File and Content Management	
6. Storage System Administration 7. Business Challenges 8. Storage Networking and Security (Optional)	

Storage Concepts

Introduction to Data Management and Storage Systems

Module Objectives

- Upon completion of this module, you should be able to:
  - Explain the different types of data
  - Explain what Cloud Computing is
  - Explain a storage systems' view of data
  - Distinguish between the physical and logical levels of data processing
  - Understand and explain the basic concepts of data consistency and data integrity

Module Topics

- Introduction to Data Management – Data Types
- Structured and Unstructured Data
- Data versus Information
- Data Processing Levels
- A Storage Systems View of Data
- Data Consistency and Data Integrity Concepts and Principles

Introduction to Data Management

Common Users: Data Examples

- Photos
- Movies
- Documents
- Email
- Personal web pages
- Data backed up to online storage such as Microsoft Sky Drive
- Increasingly popular cloud systems and on-line applications

Business Sector Data

- Accounting, invoices, and financial records
- Databases that contain data about clients
- Email communication
- Digitalization of printed documents
- Archiving

State Institutions Data

- Databases
- Audio and video records
- Hospital records
- Confidential and classified data
- Archiving and digitalization

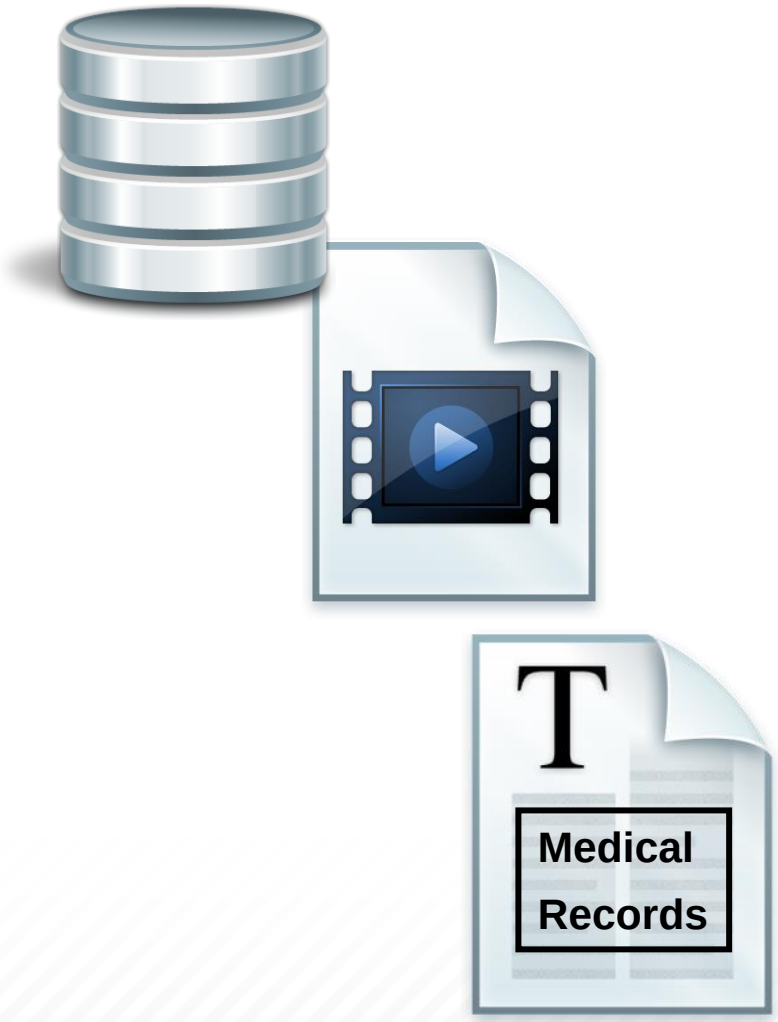

Data Lifecycle

- Data Lifecycle applies to all data that comes into existence
  - Data Retention Period applies to certain types of data and is governed by law

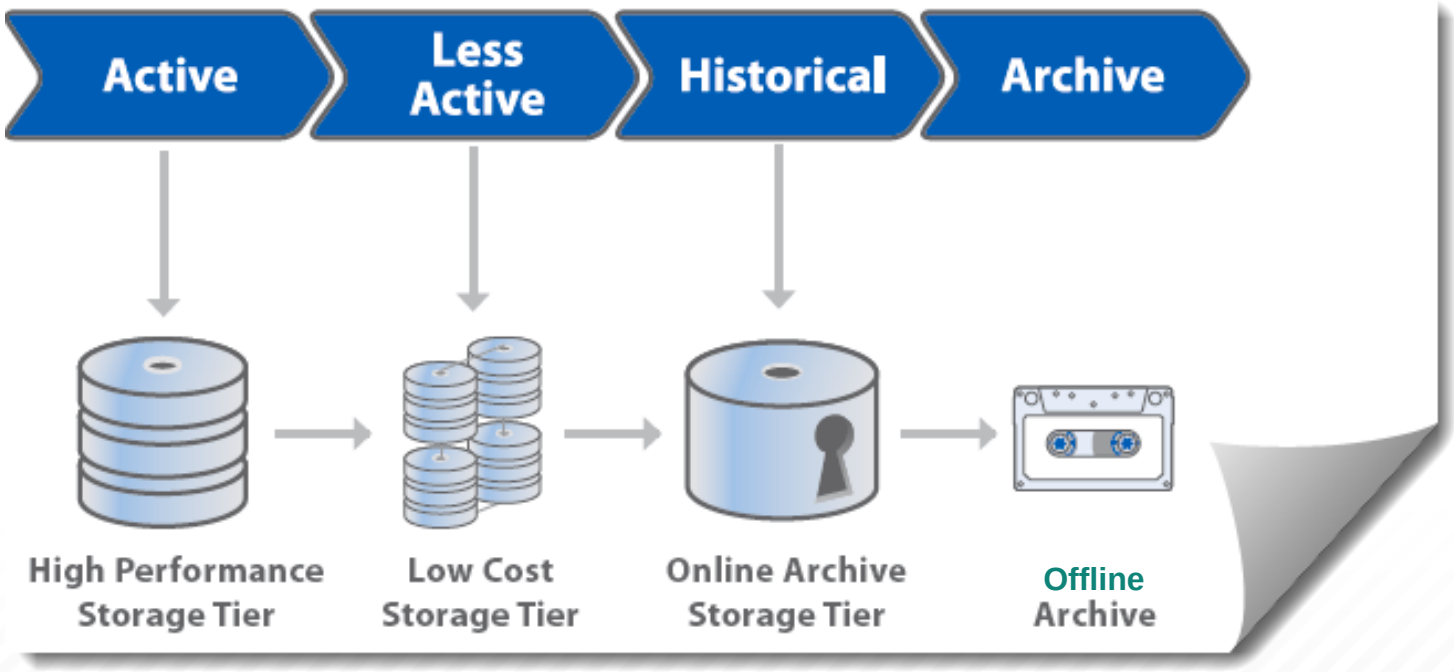

Structured and Unstructured Data

- Structured Data
  - Databases

- Unstructured Data
  - Medical Images – MRI scans
  - Photographs
  - Digital documents – check images
  - Satellite images
  - Biotechnology
  - Digital video
  - Email

The Changing Forms of Data

- New data types and business models compound the problem of exploding storage. Consider the demands brought on by these new business models as a result of the ubiquity of the Internet.

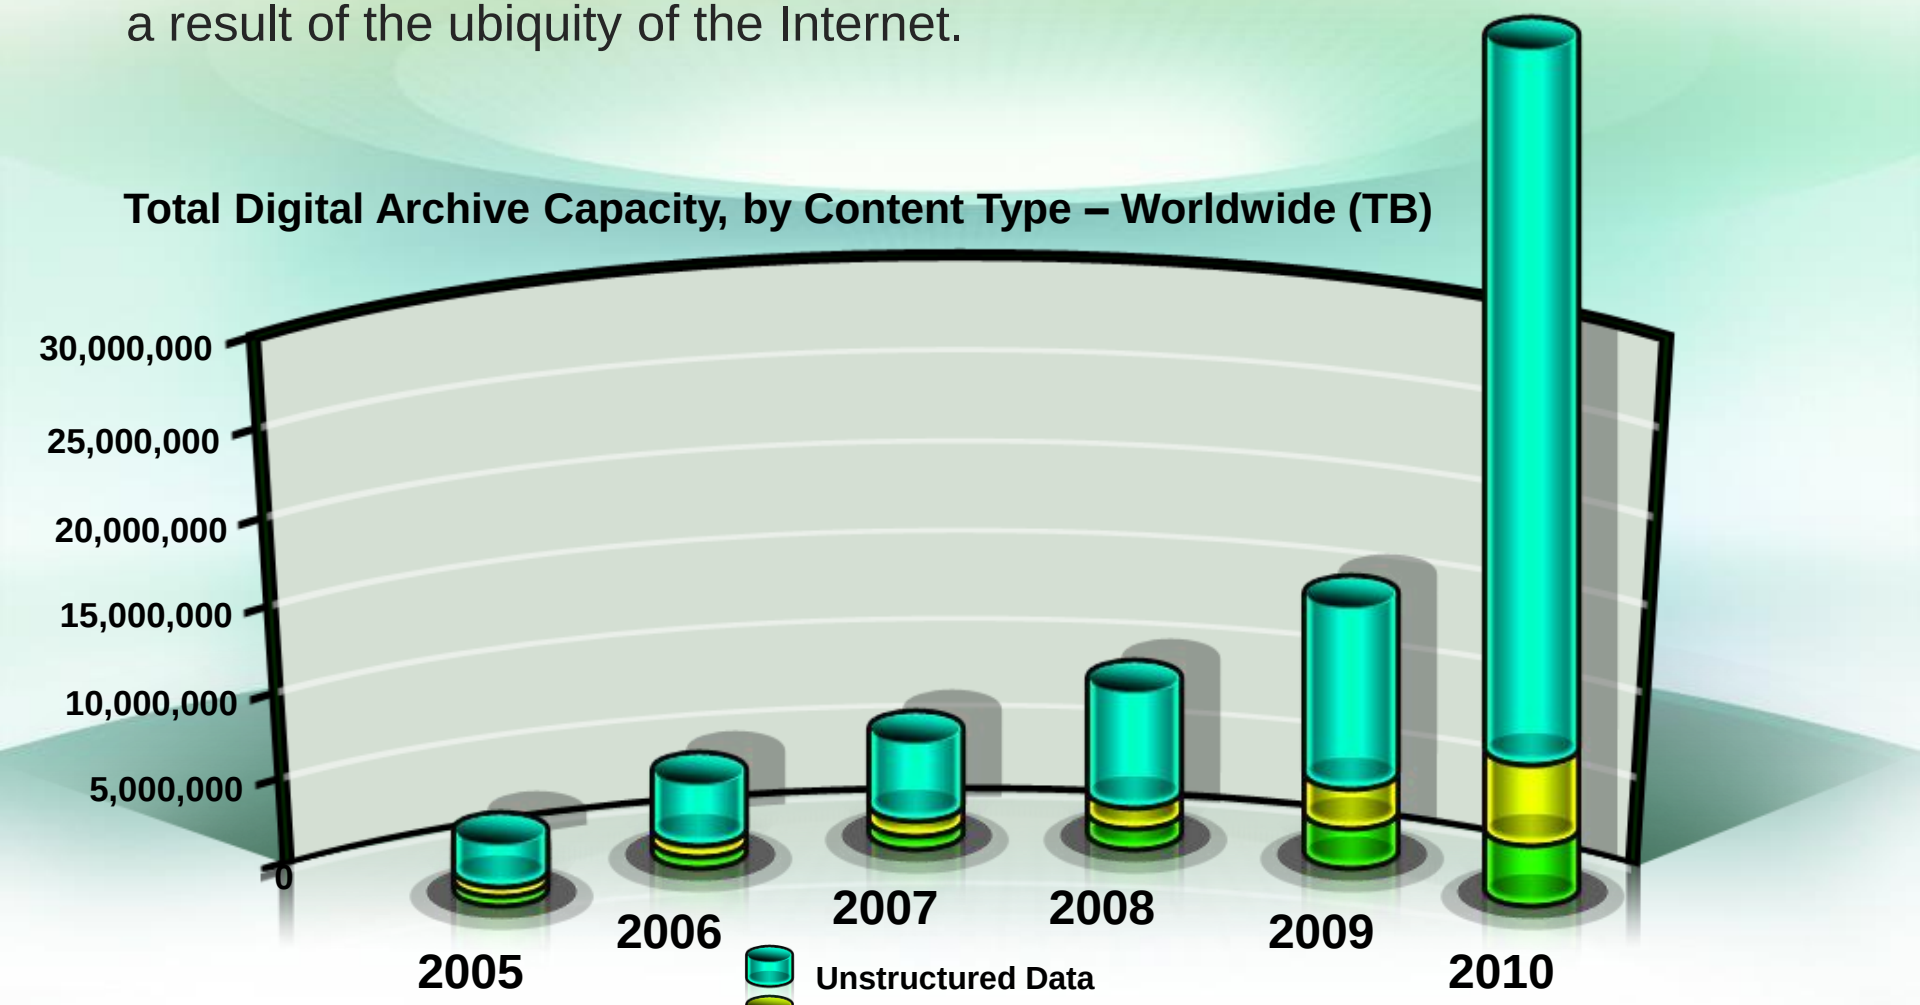

ESG Research Report: Digital Archiving: End-User Survey & Market Forecast 2006-2010

Question: What's The Difference Between Data and Information?

- Data

- Information

Data Versus Information

- Data
  - A physical and written representation of information and knowledge
  - Succession of written characters, which can be represented by numbers, letters, or symbols
- Information
  - Meaningful interpretation of data
  - Does not have to be written in characters

Data in binary code…

Data is stored to preserve and pass information on

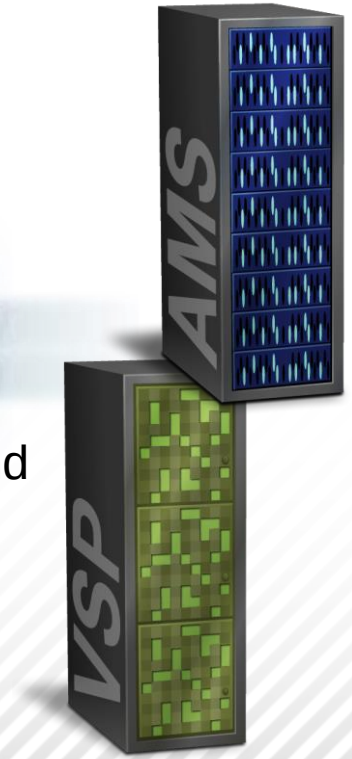

New Trends in Data Management – Cloud Computing

Video – What Is Cloud Computing?

YouTube video

- **http://www.youtube.com/watch?v=ae_DKNwK_ms&feature=youtu.be**

Three Key Cloud Characteristics

> "Cloud is a way of using technology, not a technology in itself – it's a self-service, on-demand pay-per-use model. Consolidation, virtualization and automation strategies will be the catalysts behind cloud adoption."
>
> *– The 451 Group*

The 3 key characteristics of a cloud are:

- Self-service
- Pay-per-use
- Dynamic scale up and down

Levels of Data Processing – Logical And Physical

Hard Drive Partitioning (Physical Level)

- For example, a laptop running Microsoft Windows software:
  - Installed Hard Drives — CD / DVD / memory card drives
  - Each drive assigned a letter = valid address within the Operating System
    - For example, hard drives have *C*: & *D*: addresses, DVD drive has *E*: address

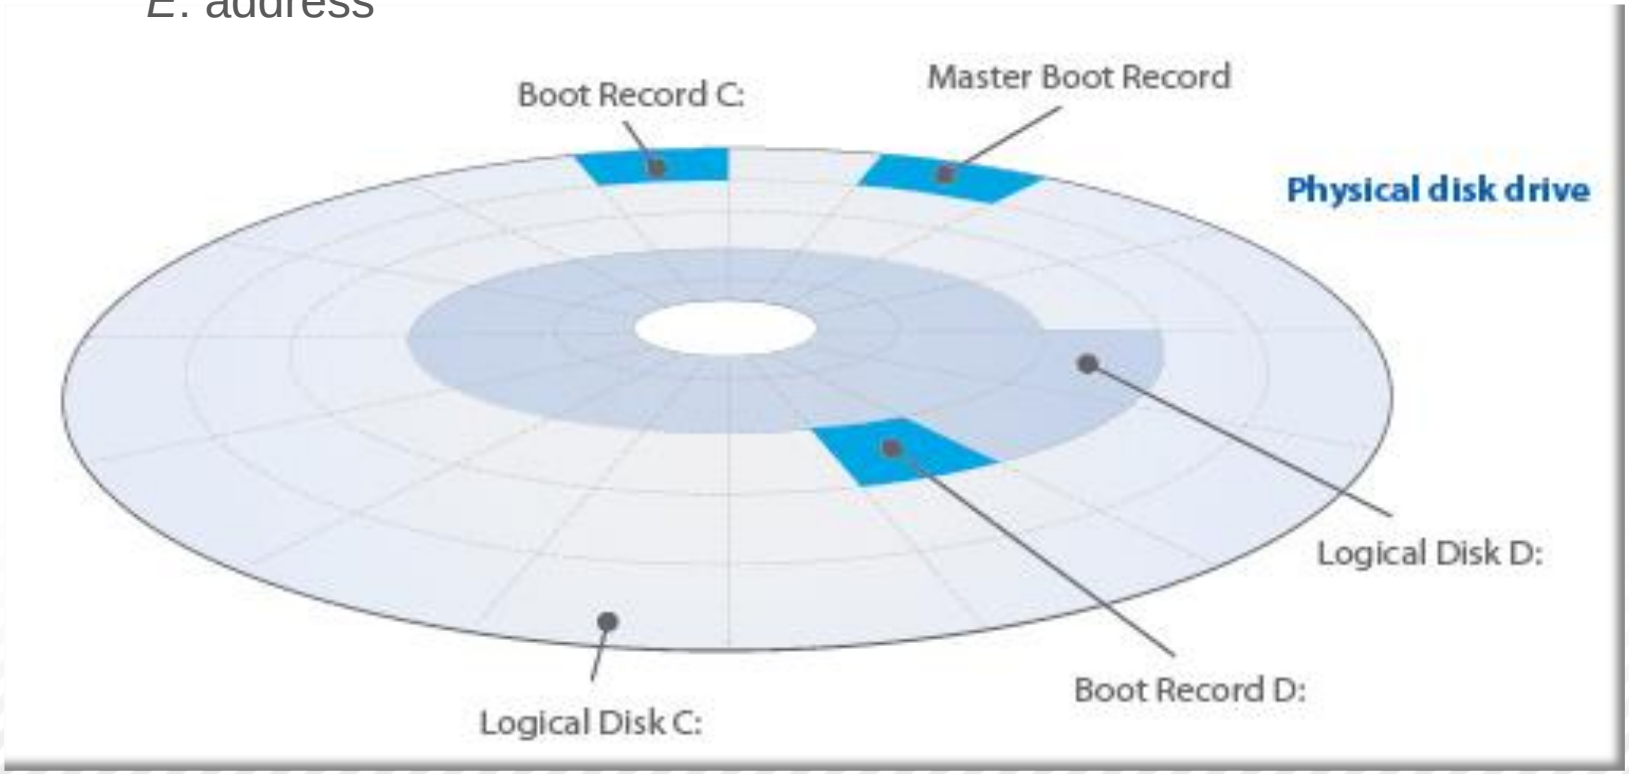

Exercise: Can You Think of Some Reasons for Partitioning a Hard Drive?

- A user may want to encrypt a partition that will contain critical data

- Other reasons:

You can also send your answers using the CHAT

Logical Level – Volumes / File Systems

- Volume – logical interface used by an operating system to access data stored on a particular media while using a single instance of a file system
- File System – the way in which files are named, organized, and stored on the hard drive

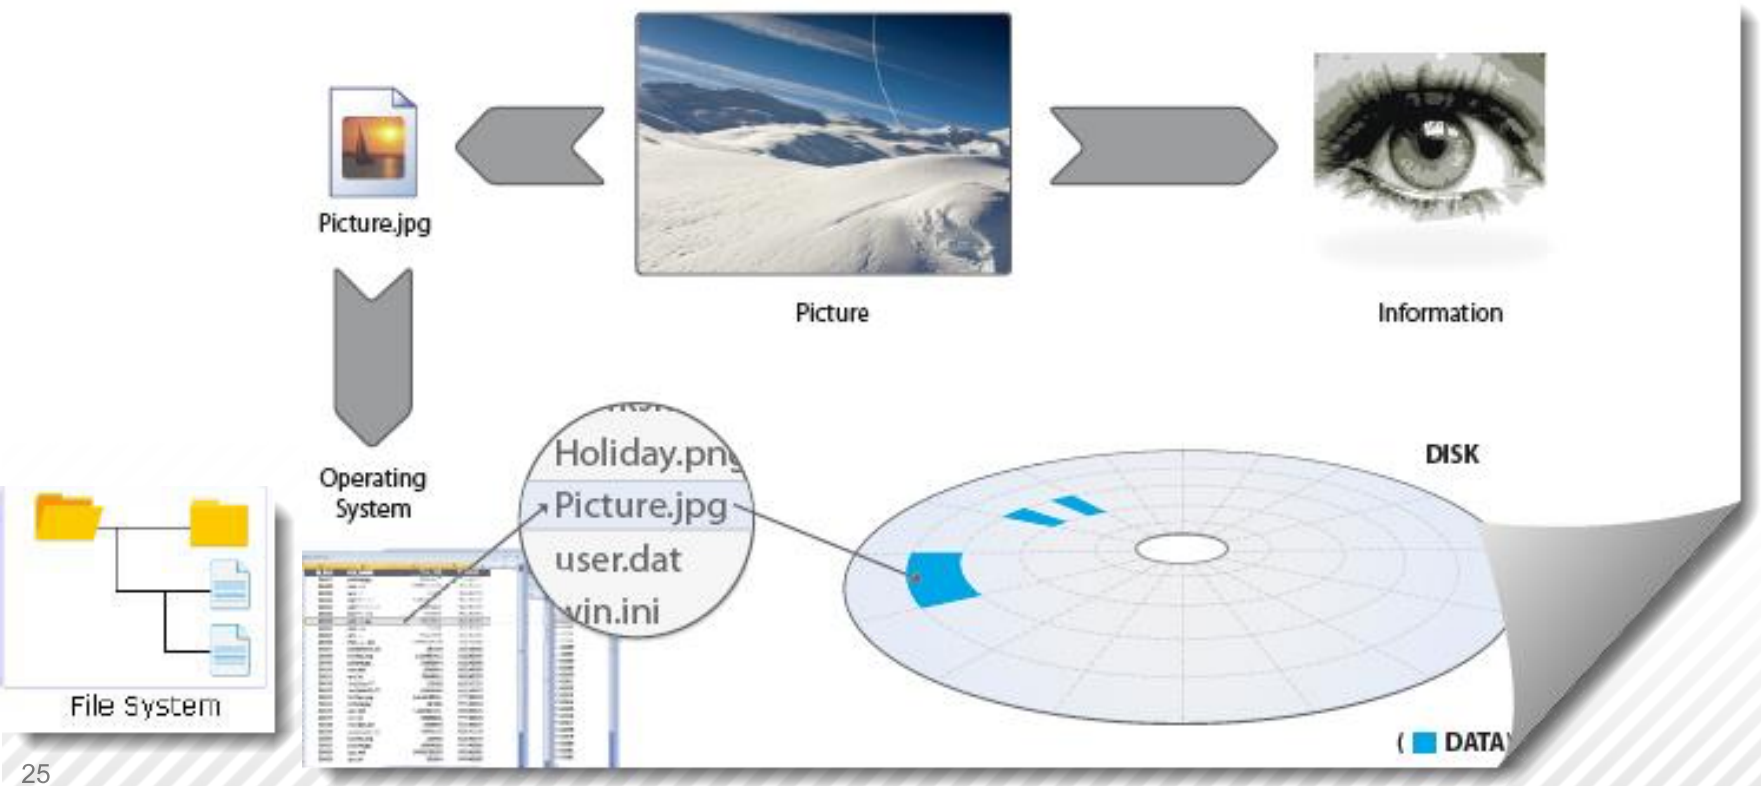

Summary – File System

- Stores information about where data is physically located
- Stores metadata, containing additional information
- Maintains data integrity and allows users to set access restrictions and permissions
- Installed on a homogenous storage area called a volume
- Applications access a file system using an application programming interface (API)
- Sets up the way files are named and organized

How the NTFS Works

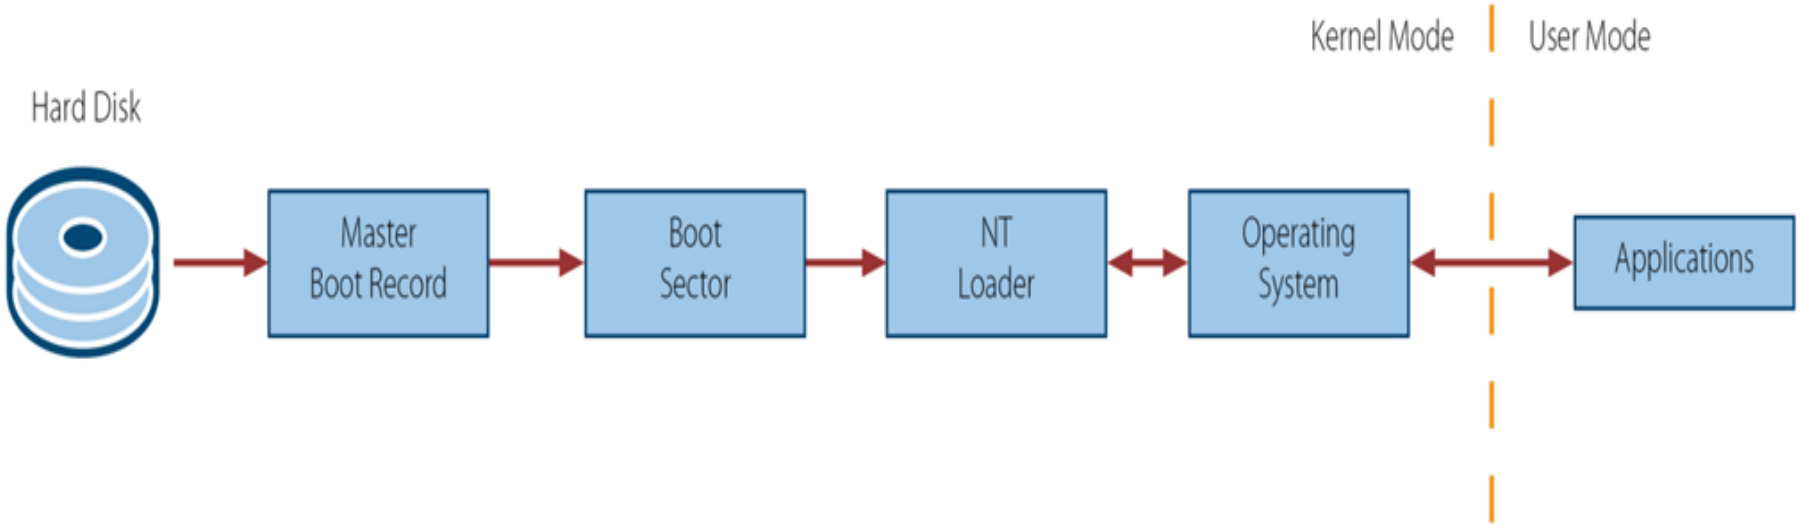

Logical Units on a Storage System

- A storage system is partitions of logical units striped across a large number of hard drives

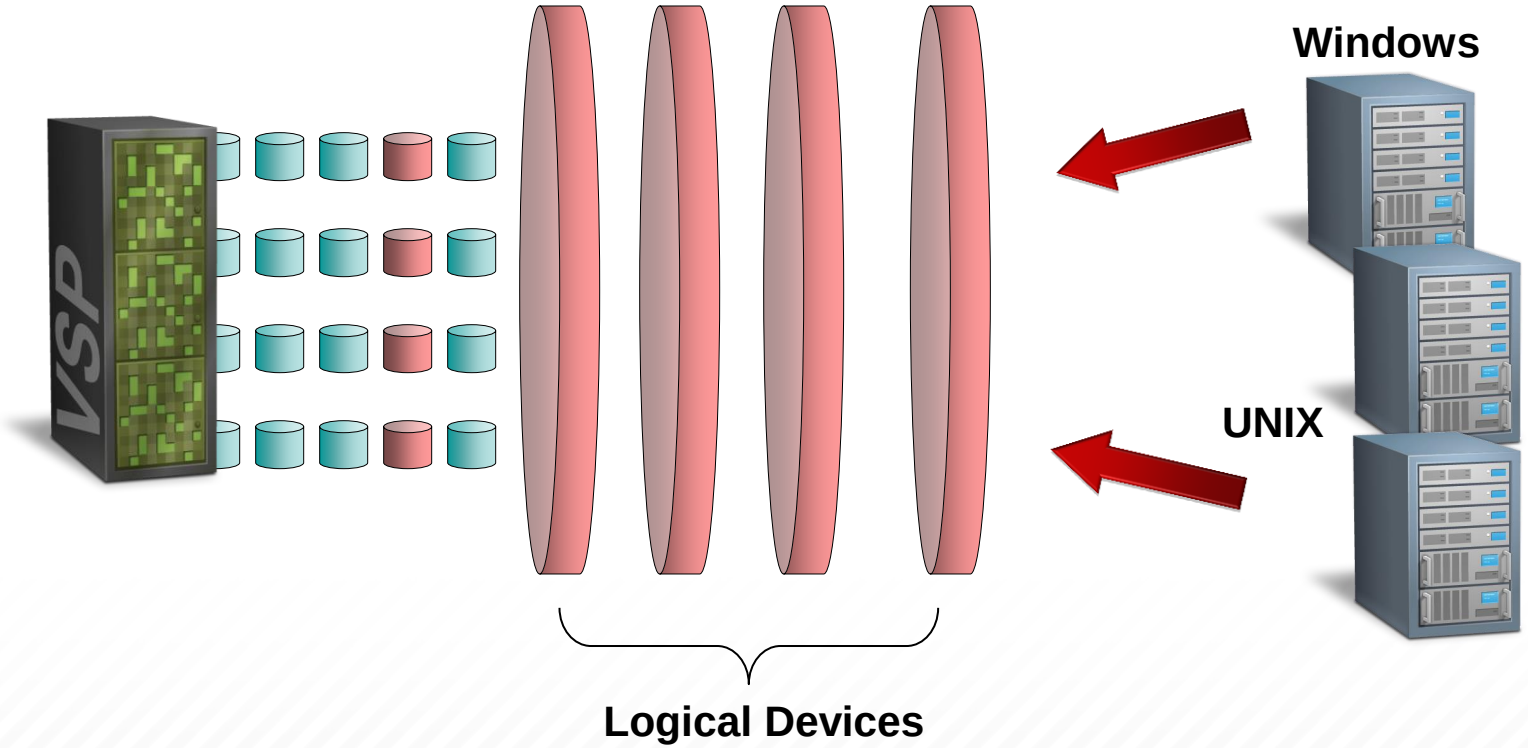

Logical Unit Number (LUN) Concept

- A LUN is a logical device mapped to a storage port

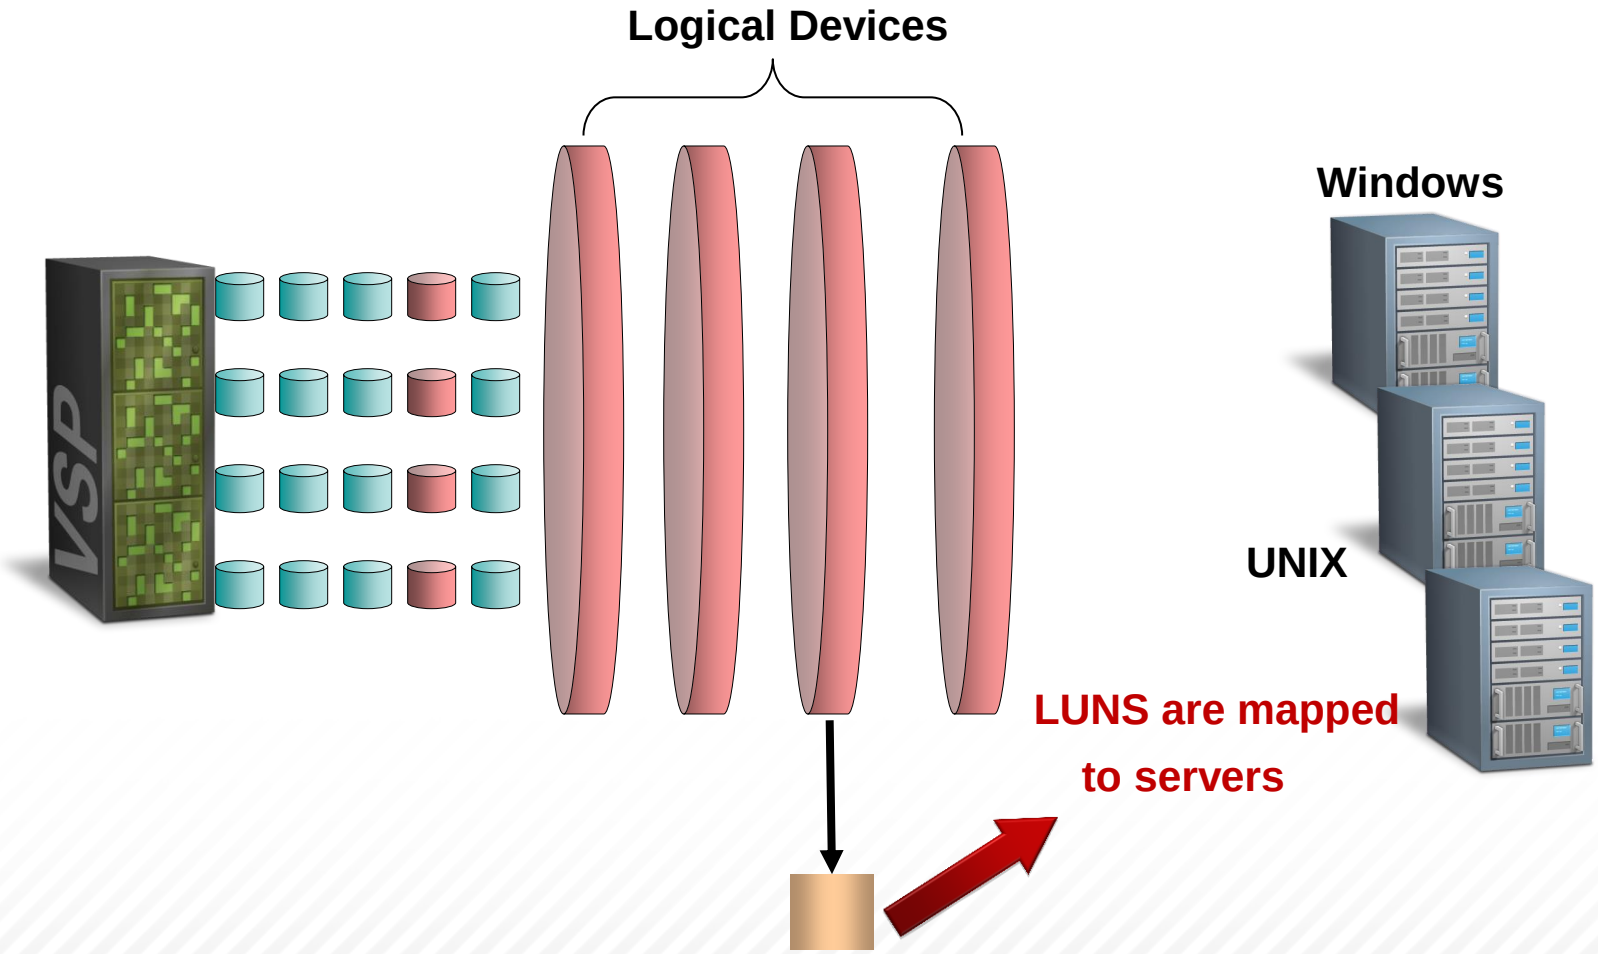

Some Definitions

Microcode and Firmware

- Microcode: built-in software that works on the lowest layer of instructions, directly controlling hardware equipment
  - Most basic software – no graphical user interface (GUI)
  - Types:
    - Microcode stored on individual hard drives
    - Microcode on a storage system also containing an integrated interface, either a command line interface (CLI) or a GUI – firmware
- Firmware: contains microcode and some kind of user interface (menus, Icons)

Data Consistency

- Data Consistency: means you have valid, usable and readable data
  - Point-in-time consistency: data is consistently as it was at any single instant in time
    - For example, synchronous (continuous) data replication
  - Transaction consistency: preventing “lost” transactions
  - Application consistency

Data Integrity

- Data Integrity: describes accuracy, reliability and correctness in terms of security and authorized access to a file
  - Policies: containing rules that govern access, preventing possible data alterations
  - Permissions and restrictions of access tools

Module Summary

- Upon completion of this module, you should have learned to:
  - Explain the different types of data
  - Explain what Cloud Computing is
  - Explain a storage systems' view of data
  - Distinguish between the physical and logical levels of data processing
  - Understand and explain the basic concepts of data consistency and data integrity

Storage Concepts

Differentiate between common basic storage architectures

Module Objectives

- Upon completion of this module, you should be able to:
  - Differentiate between common basic storage architectures

Storage Architecture Connectivity

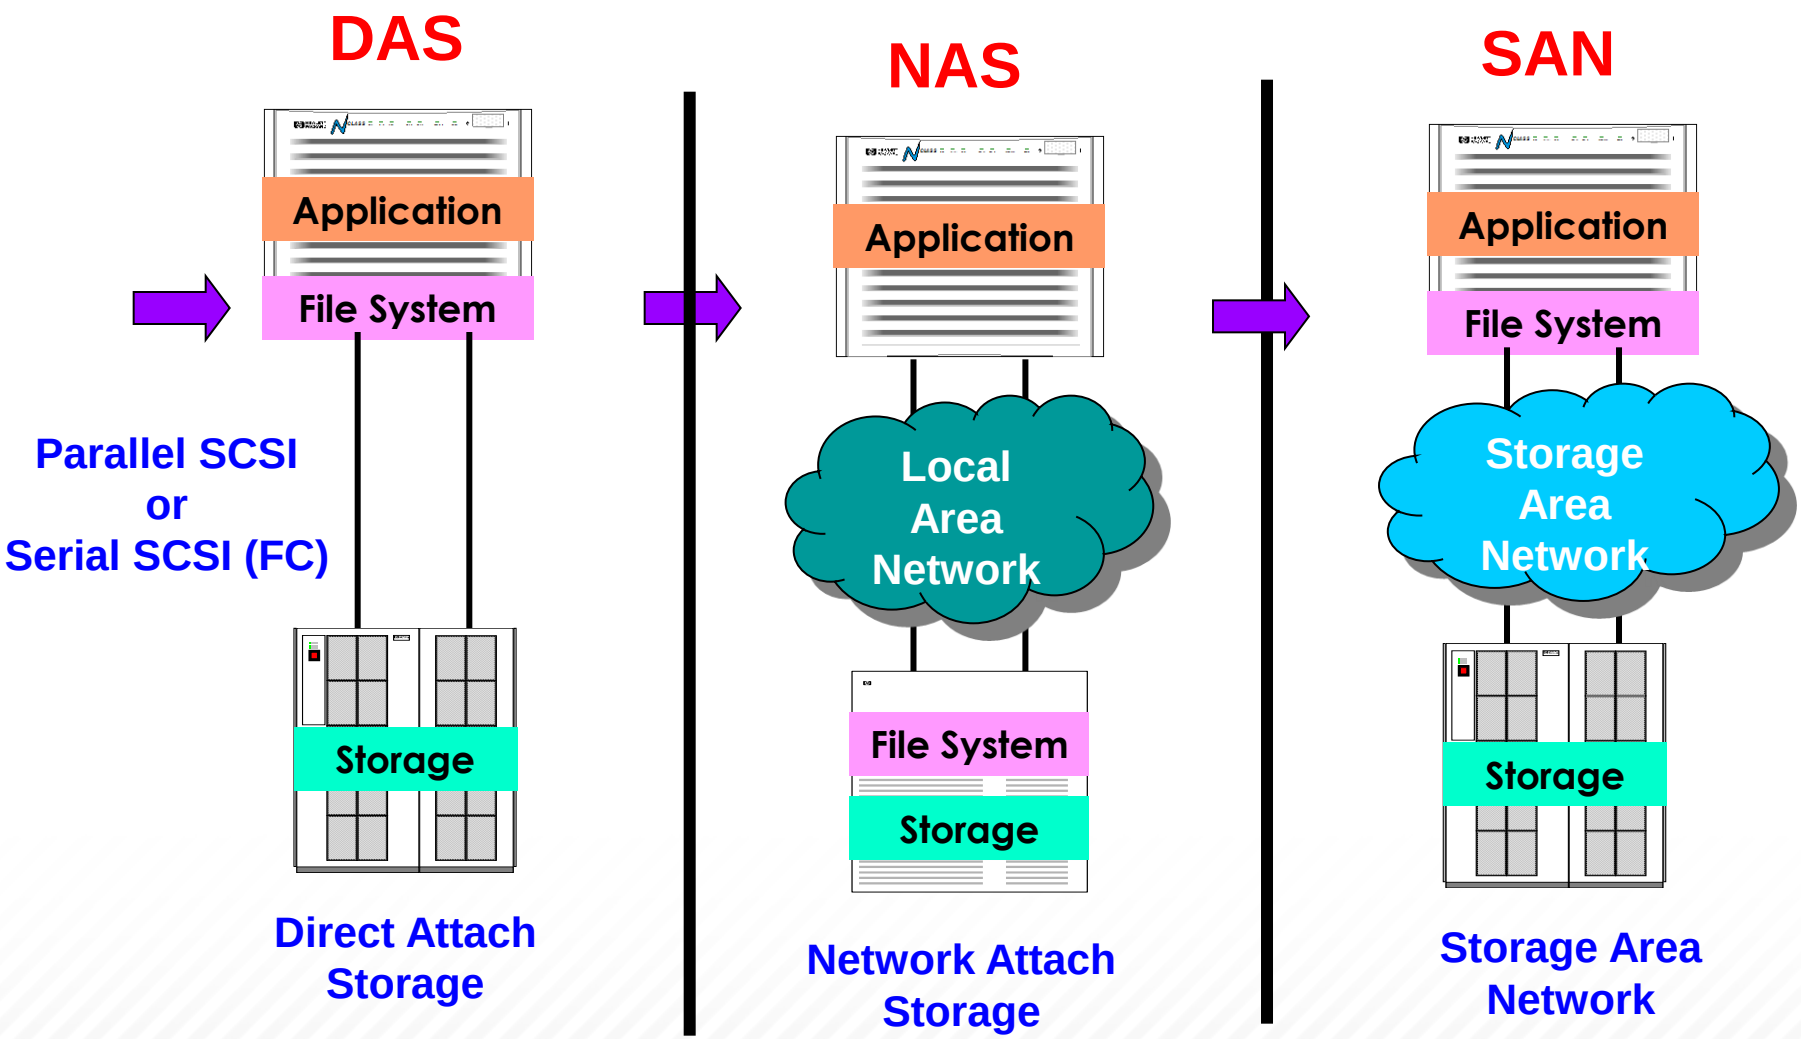

Direct Attached Storage

- Storage is directly attached to the server
- No other device on the network can access the stored data
- Example of DAS – A PC with an attached external disk drive, another; A Mainframe direct connect to SAN storage array
- Best used for accessing personal data or high speed non-shared access

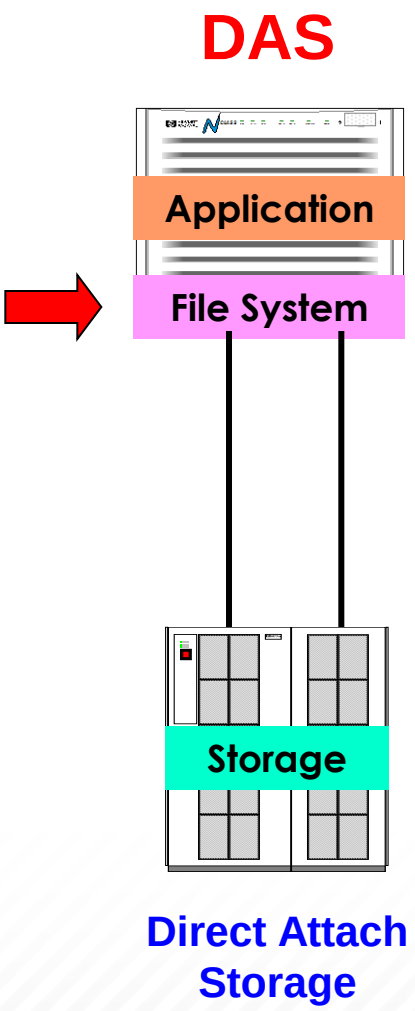

Network Attached Storage

- File oriented data access
- Optimized for file serving
- Easy Installation and Monitoring
- No server intervention (or layer) required for data access
- Low Total Cost of Ownership
- Use existing Network/cabling
- Multiple protocol support (file sharing) using NFS, SMB, CIFS, HTTP, etc.
- NAS Heads
- NAS Blades (New HDS-G-Series 400-800)
- NAS Filers (HDS-HNAS 3000/4000 Series)

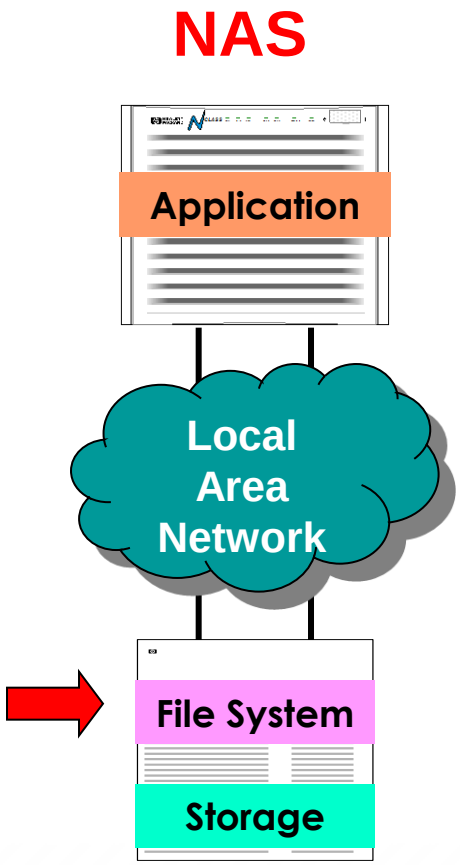

Network Attached Storage

- Targets midsize customers and remote, branch offices of large organizations
- Ideal for customers with a collaborative environment that requires sharing of files such as project management teams, law offices, and design firms
  - File serving
  - Software development
  - CAD/CAM
  - Rich media
  - Publishing and broadcast
  - Archiving
  - Near-line Data Storage to meet regulatory requirements

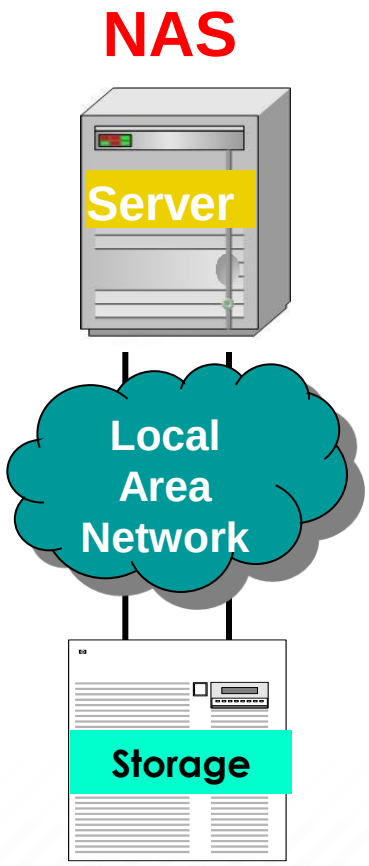

Introduction to SAN

- **Storage Area Network**

  - Is a separate network that includes computer servers, disks, and other storage devices
  - Allows networking concepts to be applied to a server/storage model
  - Has its own connections rather than using a fixed backbone network
  - Has connections that utilize Fibre Channel equipment
  - Allows very fast access among servers and storage resources
  - Enables many servers to share many storage devices
  - Designed for very high speed shared data storage up to 16 million nodes

Storage Area Network

- Designed to attach computer storage devices such as disk array controllers and tape libraries to servers
- Primary purpose is the transfer of data between computer systems and storage elements (switches/ directors [large switches])
- Consists of a communication infrastructure and a management layer so that data transfer is secure
  and robust
- SAN and NAS can co-exist on the same network infrastructure
- Complicated (relatively) and expensive to implement

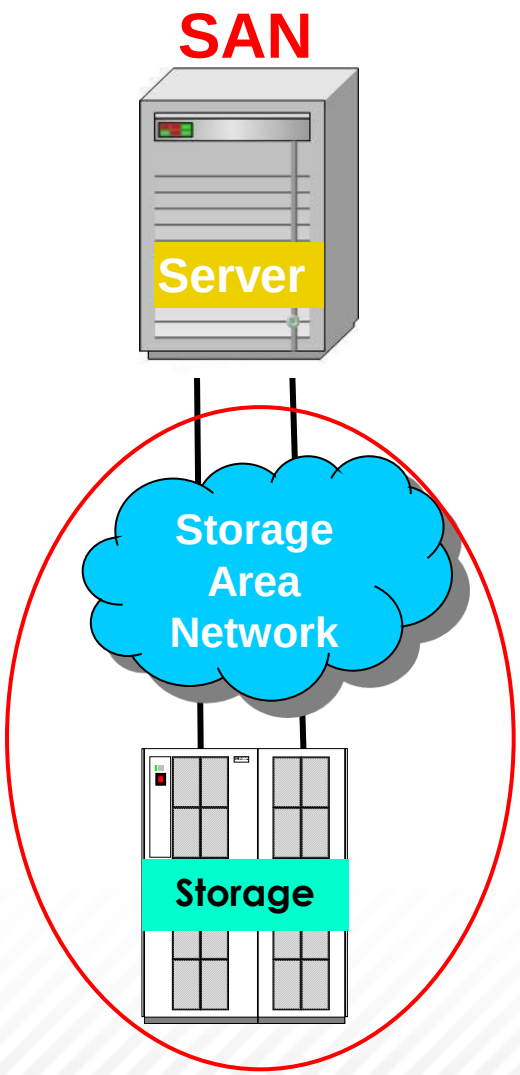

Storage Area Network

- **Best suited for complex data centers that require high availability, scalability, reliability, and performance**
  - Storage hosting providers
  - International organizations that have multiple data centers
  - Businesses that implement Service Level Agreements (SLAs)
  - Disaster Recovery and, or Business Continuity requirements

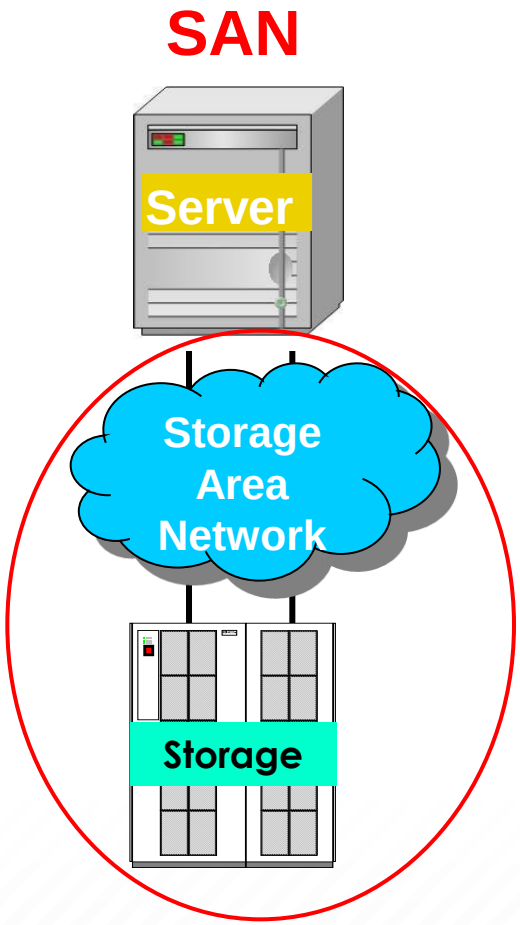

Summary

- **Direct Attached Storage (DAS) – Storage is directly attached to the application or file server**
  - Only one computer can access the storage
  - A PC with an externally-attached hard disk drive
- **Network Attached Storage (NAS) – Multiple computers can access and share the storage devices**
  - Accessed over an IP network
  - Ideal for collaborative file sharing
- **Storage Area Networks (SAN) – Designed to attach computer storage devices to servers**
  - Complex communication infrastructure organizes the connections, storage elements, and computer systems so that data transfer is secure and robust
  - Best suited for:
    - Storage hosting providers
    - International organizations that have multiple data centers
    - Anyone with need for high speed shared data access

Storage Concepts

Storage Components and Technologies

Module Objectives

- Upon completion of this module, you should be able to:
  - Identify and describe the major components of a storage system
  - Explain the different RAID levels and configurations
  - Describe the midrange storage system architecture and its components
  - Explain the factors directly influencing the performance of a storage system
  - Understand and explain the differences between a midrange and an enterprise storage system

Module Topics

- How a hard drive works
- What means of hardware redundancy we have
- How to describe a midrange storage system architecture and components

Module Flow

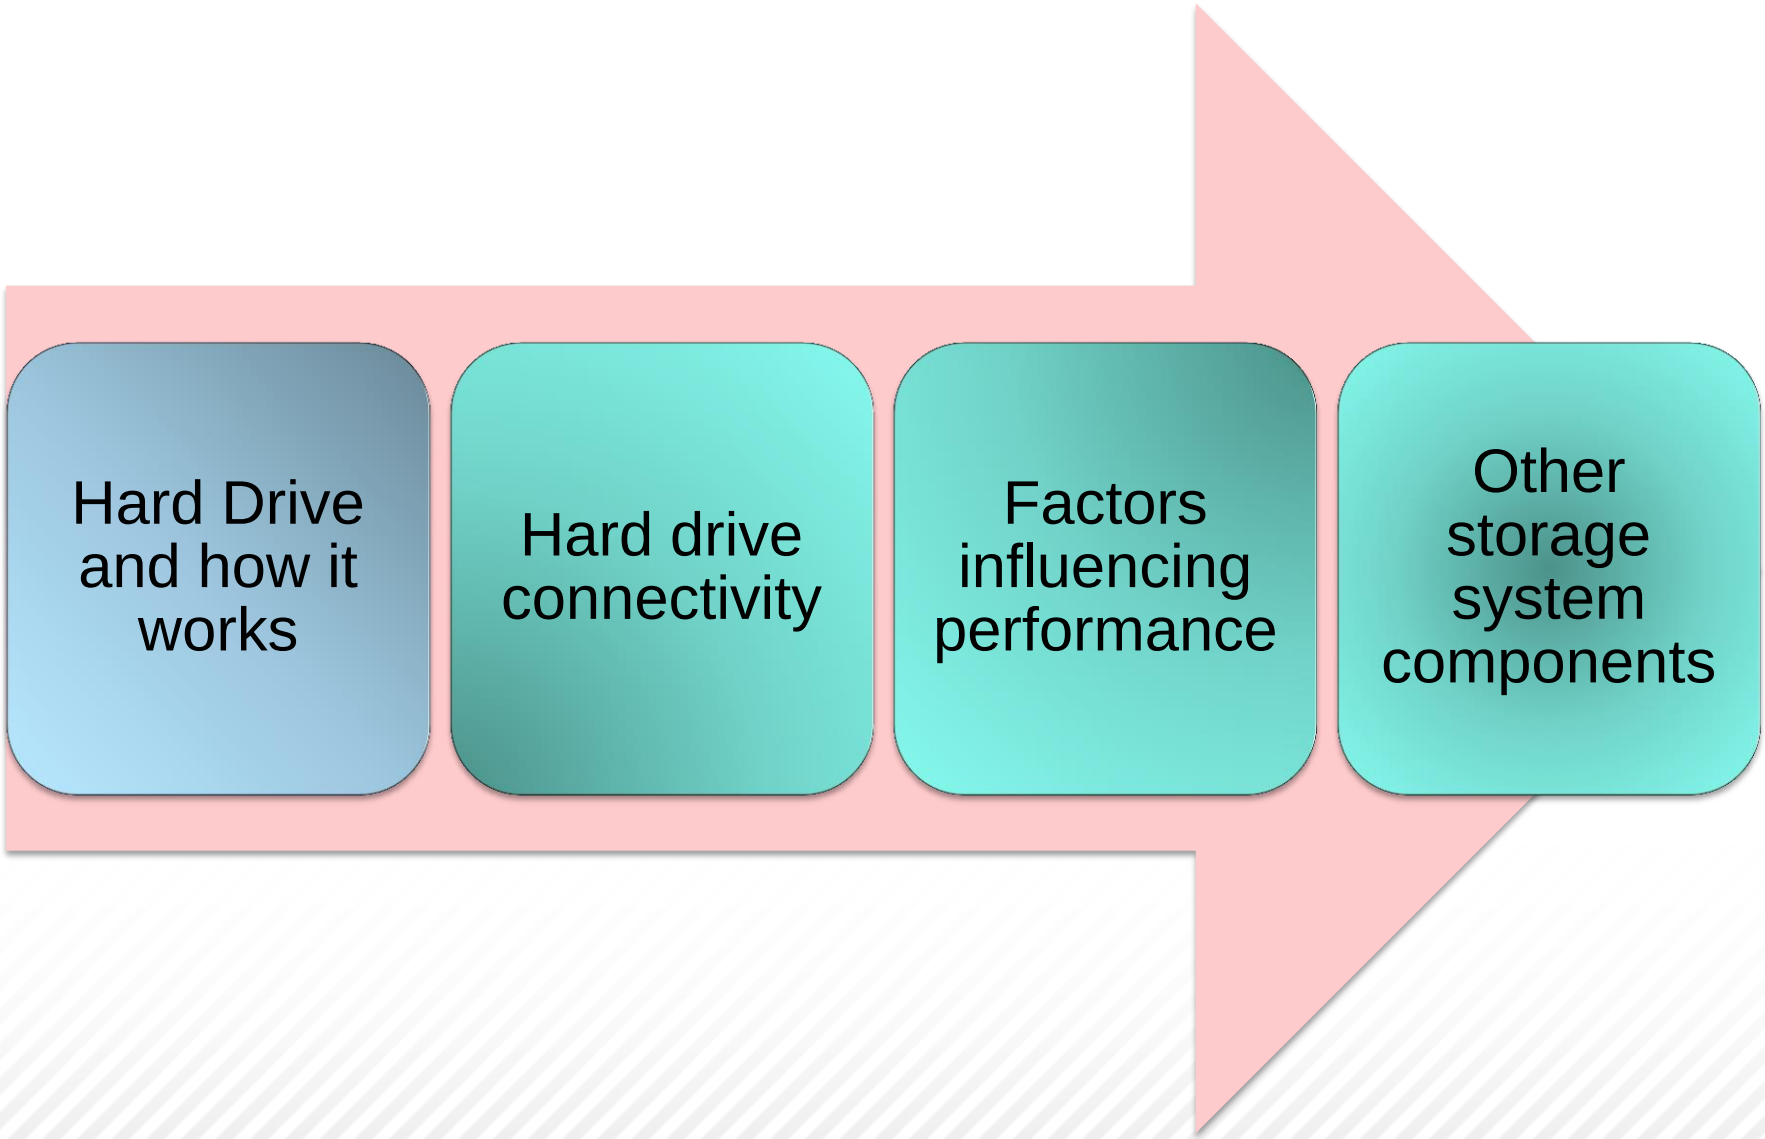
Hard Drive and how it works
Hard drive connectivity
Factors influencing performance
Other storage system components

Overview of Disk Array Components

Disk Drive Components and Connectivity

- A hard disk drive (HDD) is a **nonvolatile storage device** that stores data on a magnetic disk.
- Key components of a disk drive:

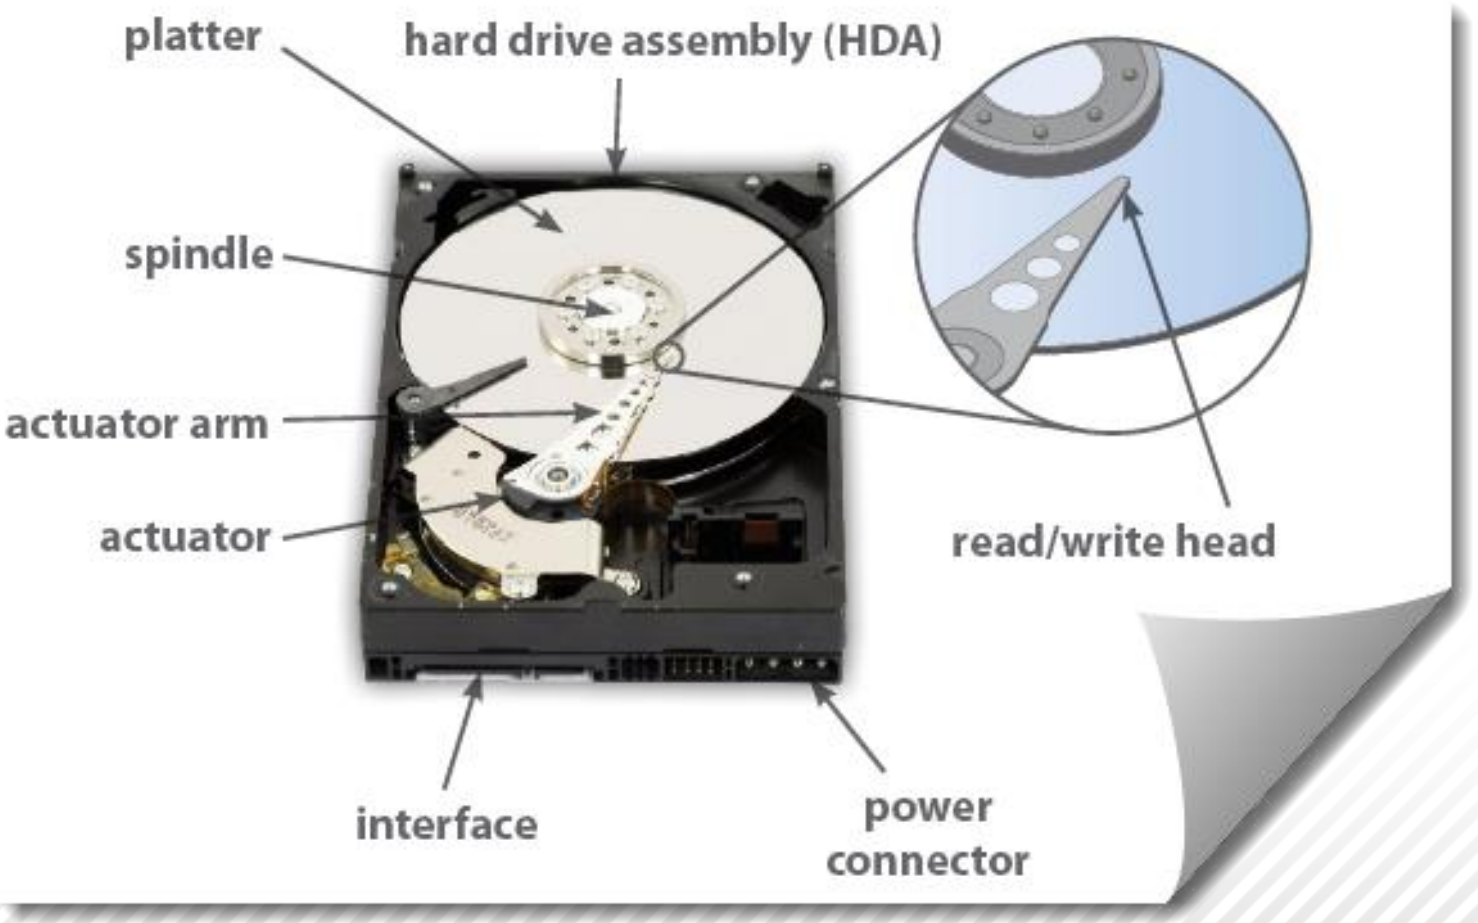

Disk Drive Components and Connectivity

- **Spindle** – A spindle holds one or more platters. It is connected to a motor that spins the platters at constant revolutions per minute (RPM)
- **Platter** – A platter is the disk that stores the magnetic patterns. It is made from a nonmagnetic material, usually glass, aluminum, or ceramic, and has a thin coating of magnetic material on both sides.

- A platter can spin at a speed of 7,200 to 18,000 RPM. The cost of an HDD increases for a higher speed.

Disk Drive Components and Connectivity

- **Head –** The read-write head of an HDD reads data from and writes data to the platters. It detects (when reading) and modifies (when writing) the magnetization of the material immediately underneath it. Information is written to the platter as it rotates at high speed past the selected head.

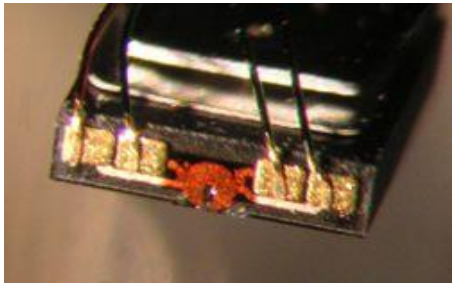

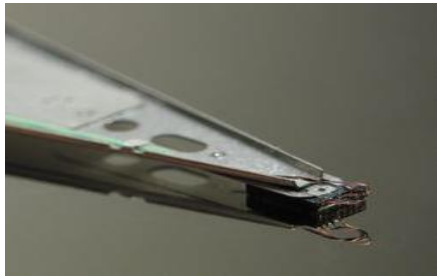

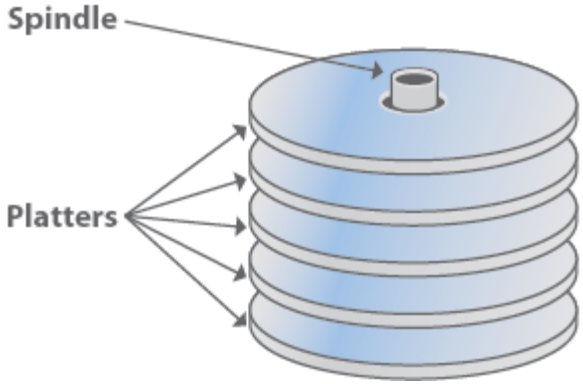

- There is one head for each magnetic platter surface on the spindle; these are mounted on a common **actuator** arm.
- **Actuator** – An actuator arm moves the heads in an arc across the spinning platters, allowing each head to access the entire data area, similar to the action of the pick-up arm of a record player.

Disk Drive Components and Connectivity

- The performance of an HDD is measured using the following parameters:

  - **Capacity –** The number of bytes an HDD can store. The current maximum capacity of an HDD is 4TB.

  - **Data transfer rate –** The amount of digital data that can be moved to or from the disk within a given time. It is dependent on the performance of the HDD assembly and the bandwidth of the data path.
    - The average data transfer rate ranges between 50-300 MB per second.

  - **Seek time –** The time the HDD takes to locate a particular piece of data. The average seek time ranges from 3 to 9 milliseconds.

Disk Drive Components and Connectivity

- Transfer Rates – Performance

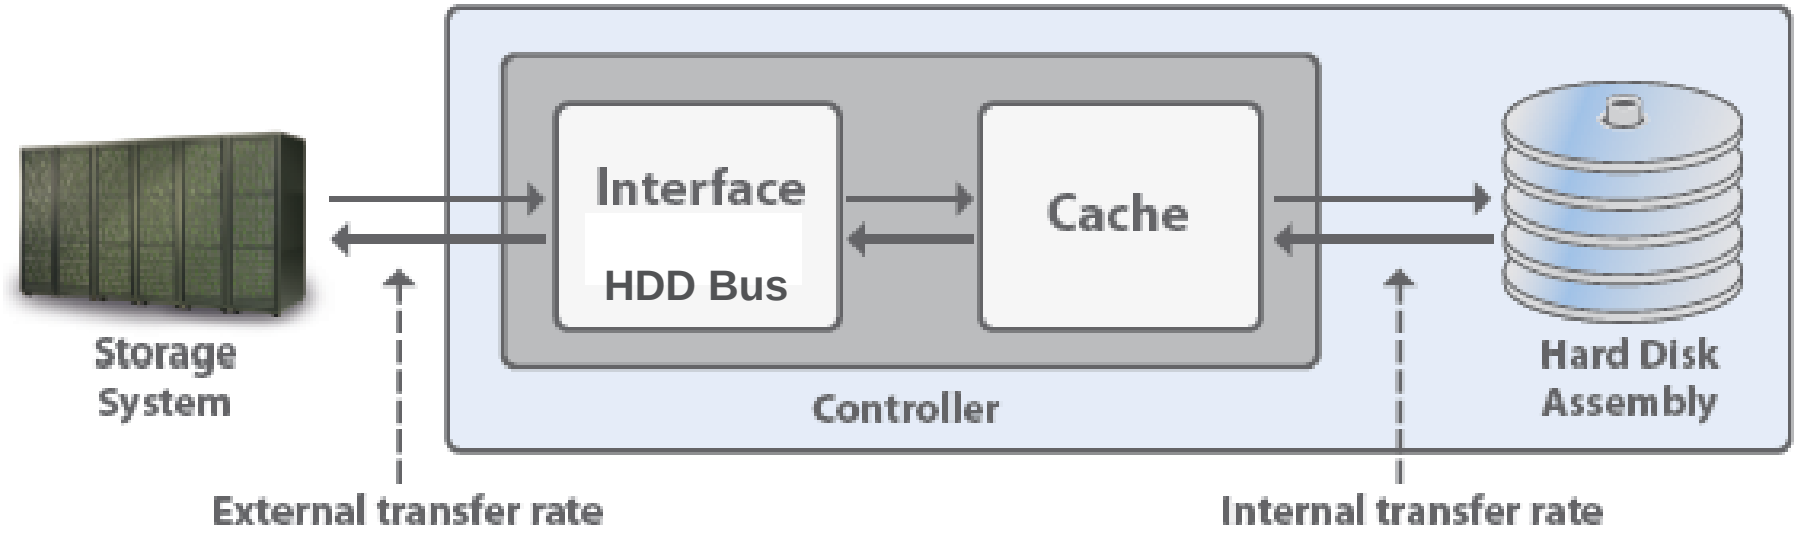

HDD Connectivity Interfaces

Disk Drive Components and Connectivity

- **Bus Cables** connect the storage central processing unit (CPU) to the HDD interface.
- **Interface** is a device that enables the connection of electrical circuits together.
- Interfaces use the following standards:
  - Parallel advanced technology attachment (PATA)
  - Serial advanced technology attachment (SATA)
  - Small computer systems interface (SCSI)
  - Serial attached SCSI (SAS)

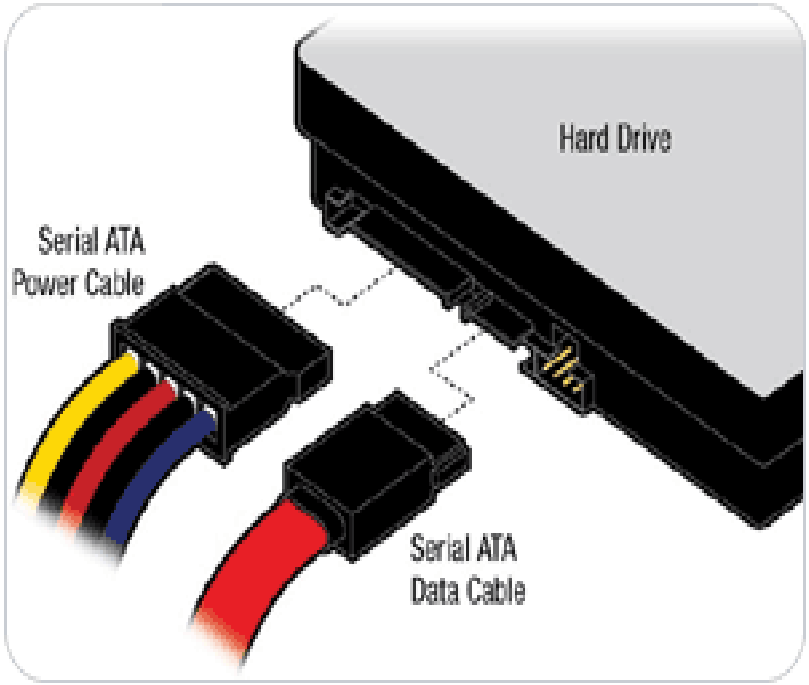

Disk Drive Components and Connectivity

- **PATA**

PATA is a standard used to connect HDDs to computers, based on parallel signaling technology. PATA cables are bulky and can be a maximum of 18 inches long, so they can be used only in internal drives.

Disk Drive Components and Connectivity

- **SATA**

SATA evolved from PATA. It uses serial signaling technology. SATA is a standard used to control and transfer data from a server or storage appliance to a client application. Compared to PATA, SATA has the following advantages:

- Greater bandwidth
- Faster data transfer rates – up to 600GB/sec
- Easy to set up and route in smaller computers
- Low power consumption
- Hot-swap support

SATA does not perform as well as SAS

Disk Drive Components and Connectivity

- **SCSI**

  A parallel interface standard used to transfer data between devices on both internal and external computer buses.

- SCSI advantages over PATA and SATA:
  - Faster data speeds
  - Multiple devices can connect to a single port
  - Device independence; can be used with most SCSI compatible hardware

- SCSI has the following disadvantages:
  - SCSI interfaces do not always conform to industry standards.
  - SCSI is more expensive than PATA and SATA.

Disk Drive Components and Connectivity

- **Serial Attached SCSI (SAS)**

  Serial Attached SCSI has evolved from the previous SCSI standards as it uses serial signaling technology. SAS is a standard used to control and transfer data with SCSI commands from a server or storage appliance to a client application.

- SAS advantages over SCSI:
  - Greater bandwidth
  - Faster data transfer rates
  - Easy set up and routing in smaller computers
  - Low power

SAS Cables

- SAS cables are similar to SATA cables.

Redundant Array of Independent Drives

RAID

- Redundant Array of Inexpensive/Independent Drives (RAID) – A method of storing data on multiple disks by combining various physical disks into a single logical unit. A logical disk is a combination of physical disks
- RAID provides the following advantages:
- Data consistency and integrity (security, protection from corruption)
- Fault tolerance
- Capacity
- Reliability
- Better Speed

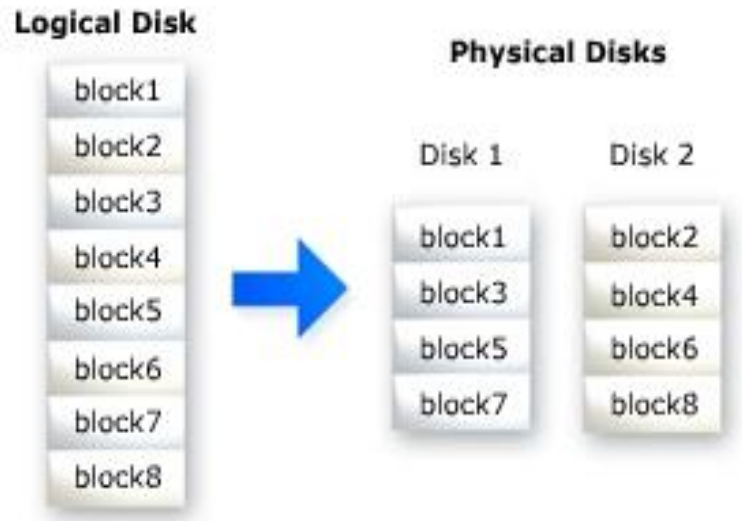

- Different types of RAID can be implemented, according to application requirements.

RAID

- The different types of RAID implementation are known as RAID *levels*.

- Common RAID levels:
  - RAID-0
  - RAID-1
  - RAID-1+ 0
  - RAID-5
  - RAID-6

RAID-0 (Data Striping)

- **RAID-0**
  - RAID-0 implements striping; data is spread evenly across two or more disks.
- RAID-0 Benefits:
  - Easy to implement
  - Increased performance in terms of data access (more disks equals more heads, which enables parallel access to more data records)
- RAID-0 Disadvantage:
  - This RAID level has no redundancy and no fault tolerance. If any disk fails, data on the remaining disks cannot be retrieved, which is the major disadvantage.

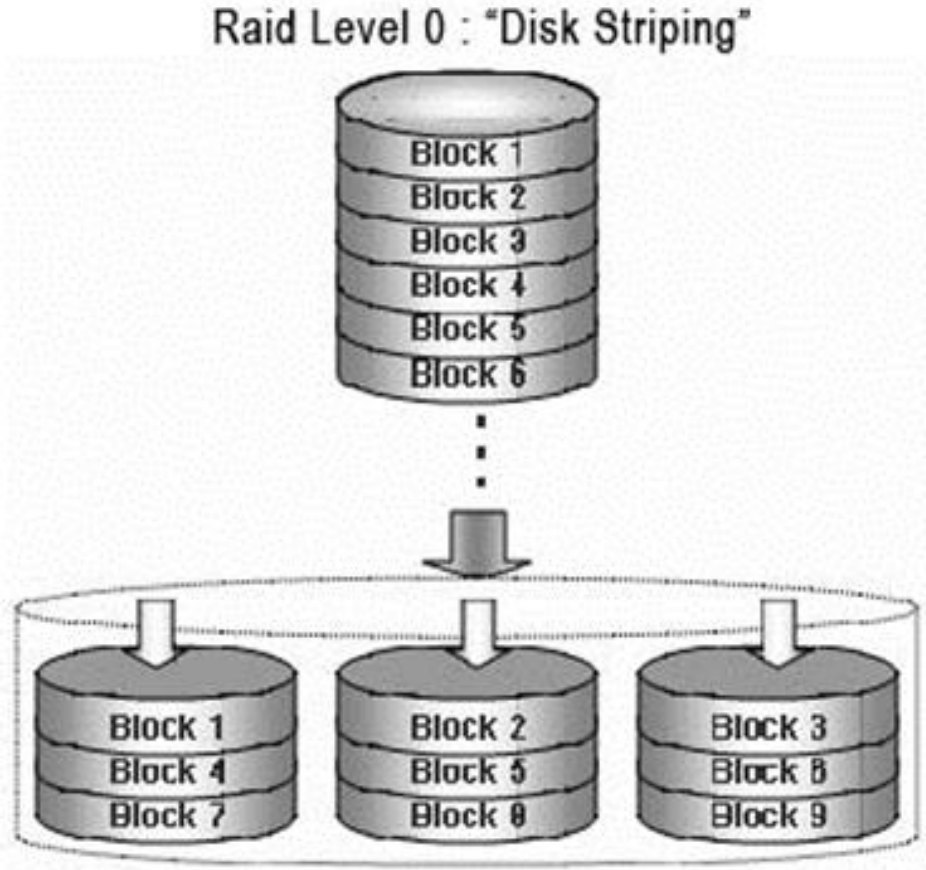

RAID-0 uses only **data striping**.

RAID-1 (Data Mirroring)

- **RAID-1**
  - Implements mirroring to create exact copies of the data on two or more disks.
- RAID-1 Advantages:
  - Reduces the overhead of managing multiple disks and tracking the data.
  - Read time is fast because the system can read from either disk.
  - If a disk fails, RAID-1 ensures there is an exact copy of the data on the second disk.
- RAID-1 Disadvantages:
  - The storage capacity is only half of the actual capacity as data is written twice.
  - RAID-1 is expensive; doubles the storage.

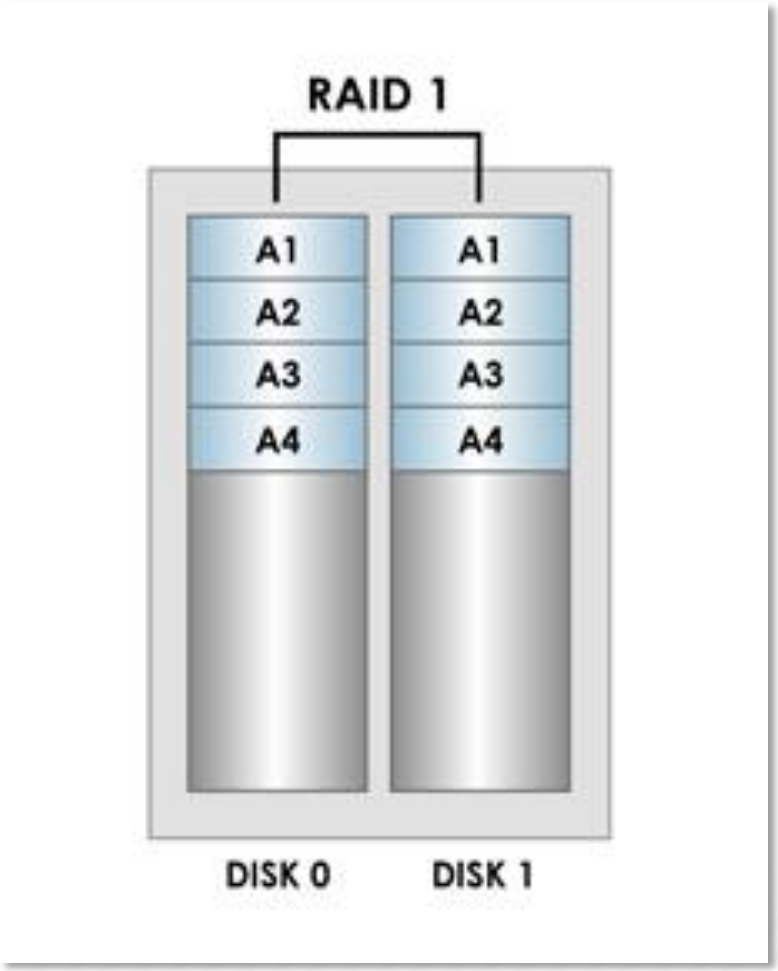

RAID-1+

- **RAID-1+**
  - RAID-1+0 is an example of multiple or **nested** RAID levels.
  - A **nested** RAID level combines the features of multiple RAID levels. The sequence in which they are implemented determines the naming of the nested RAID level.
  - For example, if RAID-0 is implemented before RAID-1, the RAID level is called RAID-0+1. RAID-1+0 combines the features of RAID-0 and RAID-1 by mirroring a striped array.

- RAID-1+ has the following advantages:
  - Easy to implement
  - Fast read/write speed
  - Data protection

- RAID-1+ has the disadvantage of high cost to implement.

RAID-1+

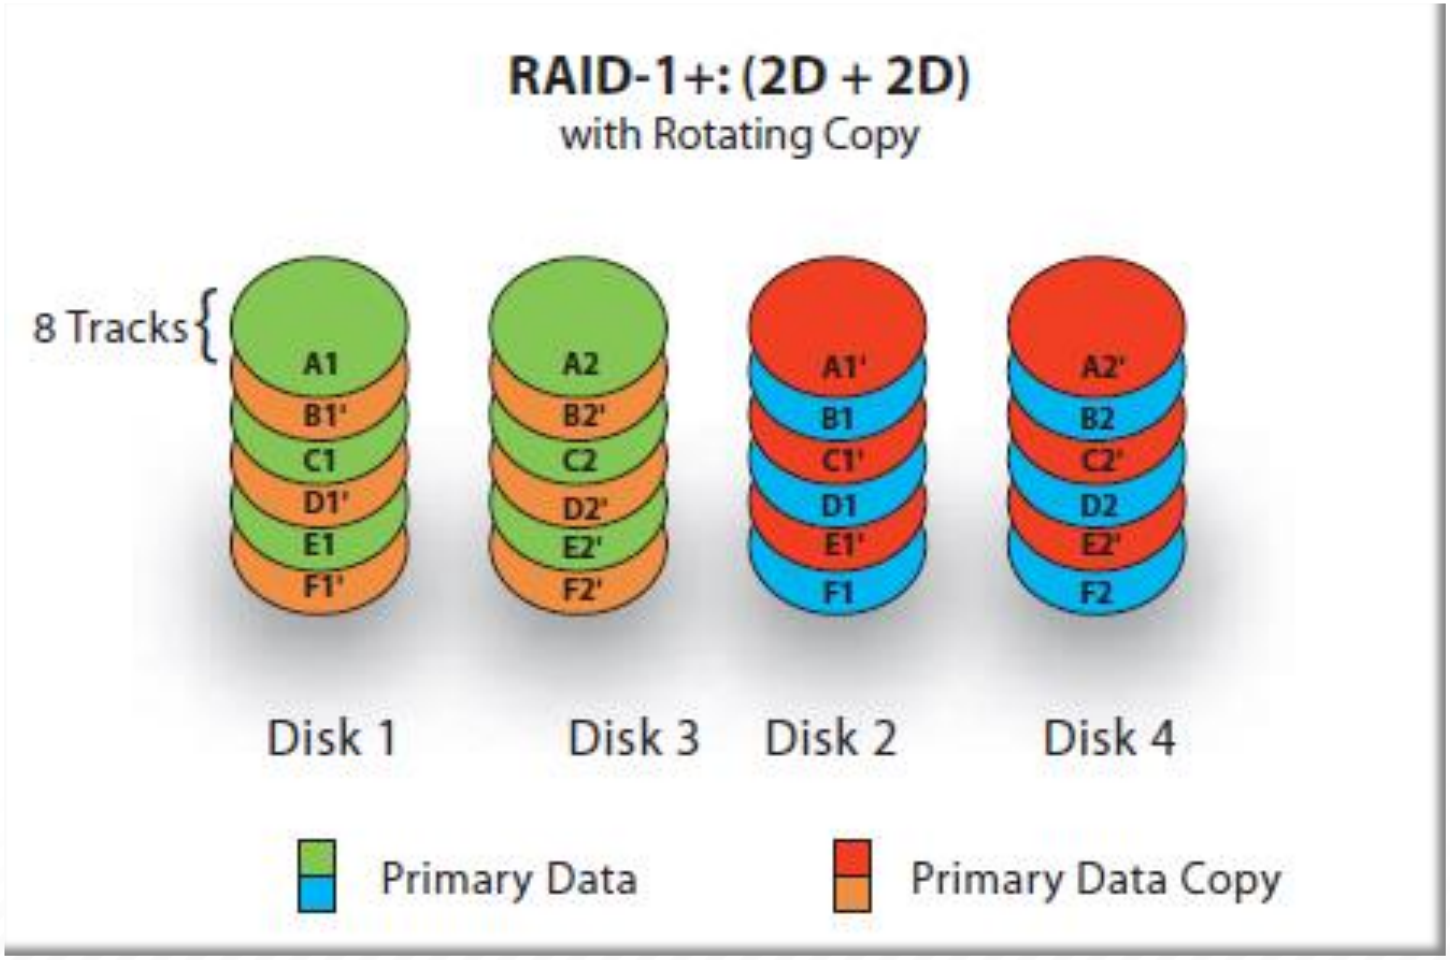
RAID-1+: (2D + 2D)
with Rotating Copy
8 Tracks
A1
B1'
C1
D1'
E1
F1'
A2
B2'
C2
D2'
E2'
F2'
A1'
B1
C1'
D1
E1'
F1
A2'
B2
C2'
D2
E2'
F2
Disk 1
Disk 3
Disk 2
Disk 4
Primary Data
Primary Data Copy

RAID-5

- **RAID-5**
  - RAID-5 consists of a minimum of three disks (two data and one parity)
  - RAID-5 distributes parity information across all disks to minimize potential bottlenecks if one disk fails, in which case parity data from the other disks is used to recreate the missing information.

- RAID-5 has the following advantages:
  - The most common and secure RAID level
  - Fast read speed
  - Ensures data recovery if a disk failure occurs

- RAID-5 has the following disadvantages:
  - Extra overhead required to calculate and track parity data
  - Slower writes because it has to calculate parity before writing data

RAID-5

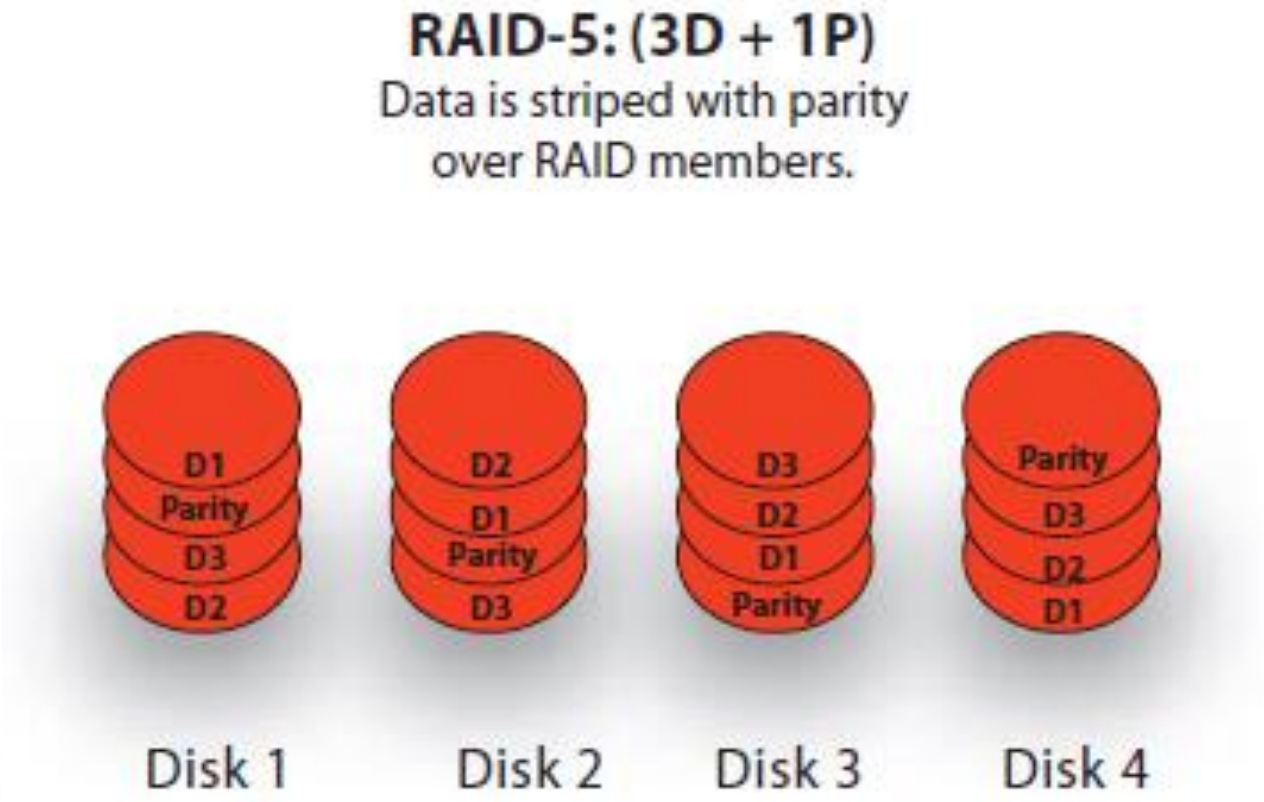
RAID-5: (3D + 1P)
Data is striped with parity
over RAID members.
D1
Parity
D3
D2
D2
D1
Parity
D3
D3
D2
D1
Parity
Parity
D3
D2
D1
Disk 1
Disk 2
Disk 3
Disk 4

RAID-6

- **RAID-6**
  - RAID-6 is similar to RAID-5 with an additional parity disk
  - In RAID-5, if a second disk fails before the first failed disk has been rebuilt, data can be irretrievably lost. The additional parity drive in RAID-6 provides a solution to this problem.

- RAID-6 is designed for large environments and offers the following benefits:
  - Provides protection against double-disk failure
  - Has a fast read speed

- RAID-6 has the following disadvantages:
  - Similar to RAID-5 plus the cost of the extra parity disk
  - Slight performance overhead

RAID-6

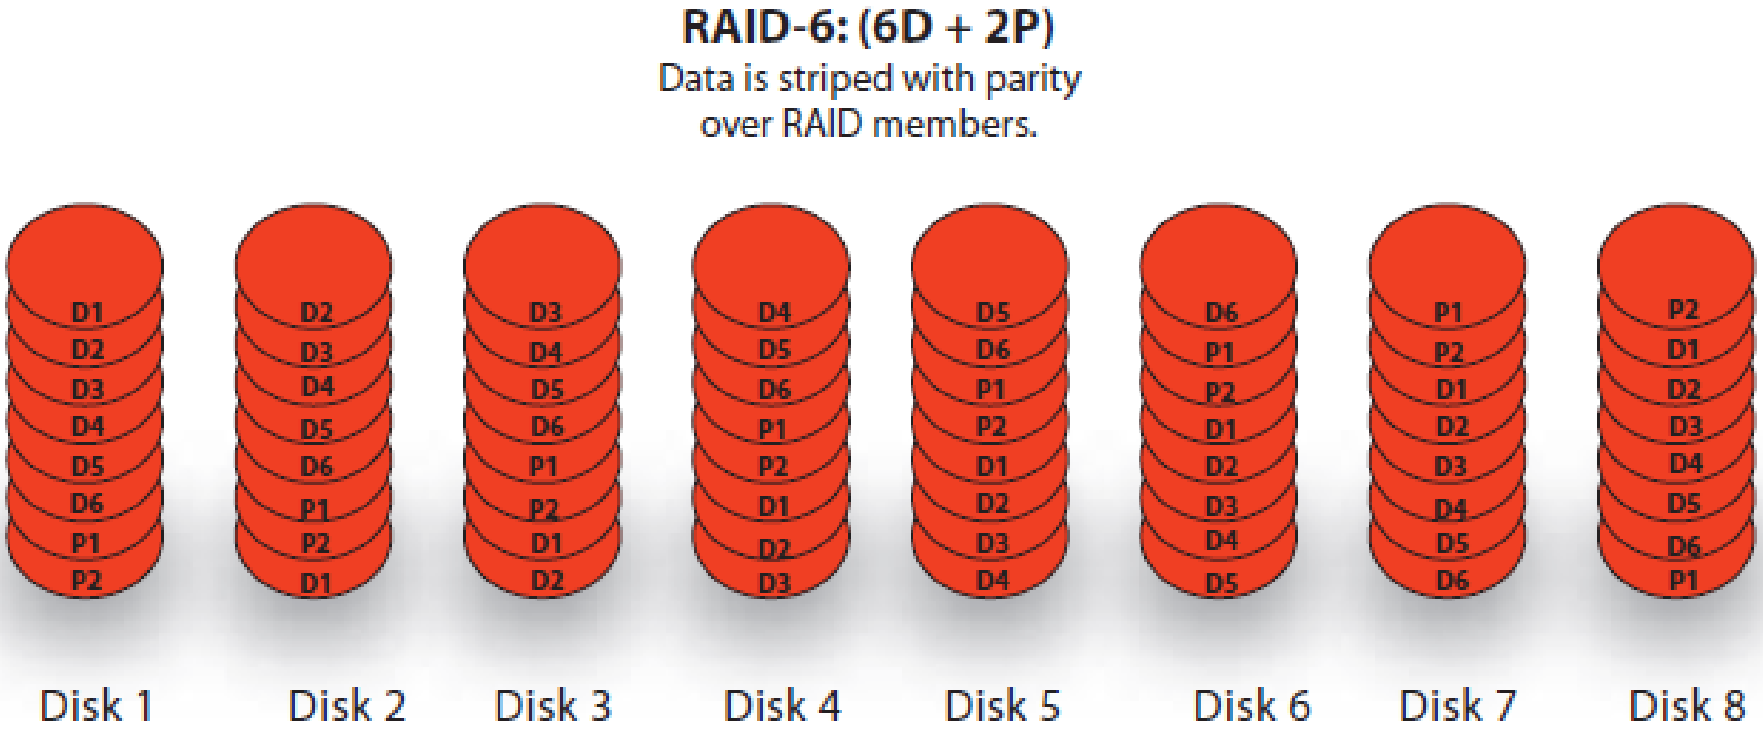
RAID-6: (6D + 2P)
Data is striped with parity over RAID members.
D1 D2 D3 D4 D5 D6 P1 P2
D2 D3 D4 D5 D6 P1 P2 D1
D3 D4 D5 D6 P1 P2 D1 D2
D4 D5 D6 P1 P2 D1 D2 D3
D5 D6 P1 P2 D1 D2 D3 D4
D6 P1 P2 D1 D2 D3 D4 D5
P1 P2 D1 D2 D3 D4 D5 D6
P2 D1 D2 D3 D4 D5 D6 P1
Disk 1
Disk 2
Disk 3
Disk 4
Disk 5
Disk 6
Disk 7
Disk 8

RAID Type Configuration and Usage

	RAID-10	RAID-5	RAID-6
Description	Data Striping and Mirroring	Data Striping with distributed parity	Data Striping with two distributed parities
Number of Disks	4/8	4/8	8
Benefit	Highest performance with data redundancy; higher write IOPS per Parity Group than with similar RAID-5.	The best balance of cost, reliability, and performance.	Balance of cost, with extreme emphasis on reliability
Disadvantages	Higher cost per number of physical disks	Performance penalty for high percentage of Random Writes	Performance penalty for all writes

Spare Disks — Sparing

- **Correction Copy** – Occurs when a drive in a RAID group fails and a compatible spare drive exists. Data is then reconstructed on the spare drive.

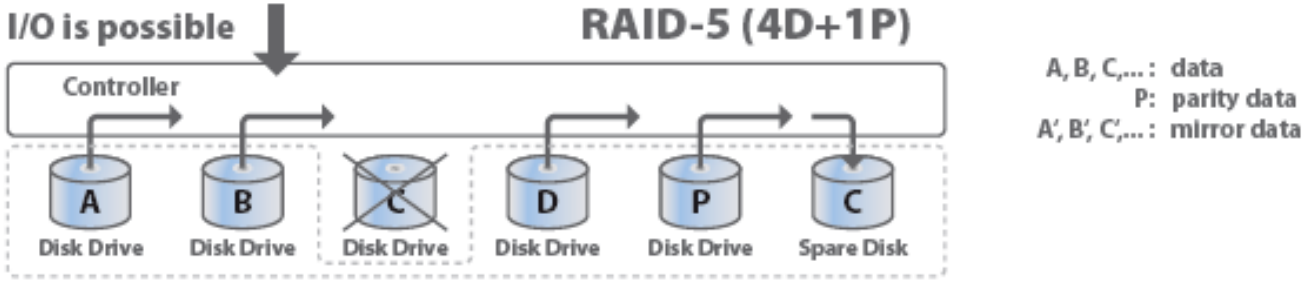

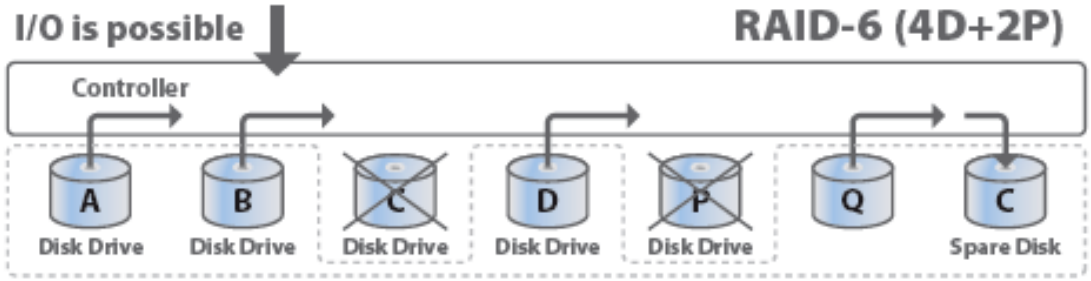

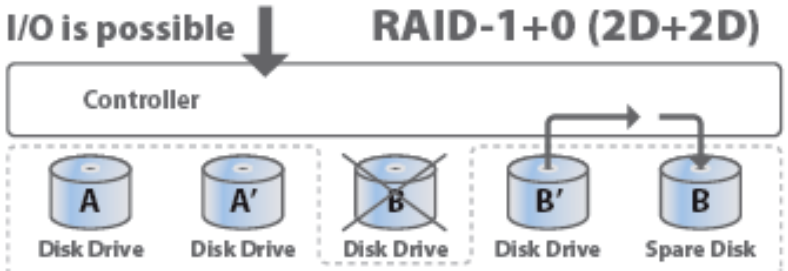

Spare Disks — Sparing

- **Dynamic Sparing** – occurs if the online verification process (built-in diagnostic) determines that the number of errors has exceeded the specified threshold of a disk in a RAID group. Data is then moved to the spare disk, which is a much faster process than data reconstruction

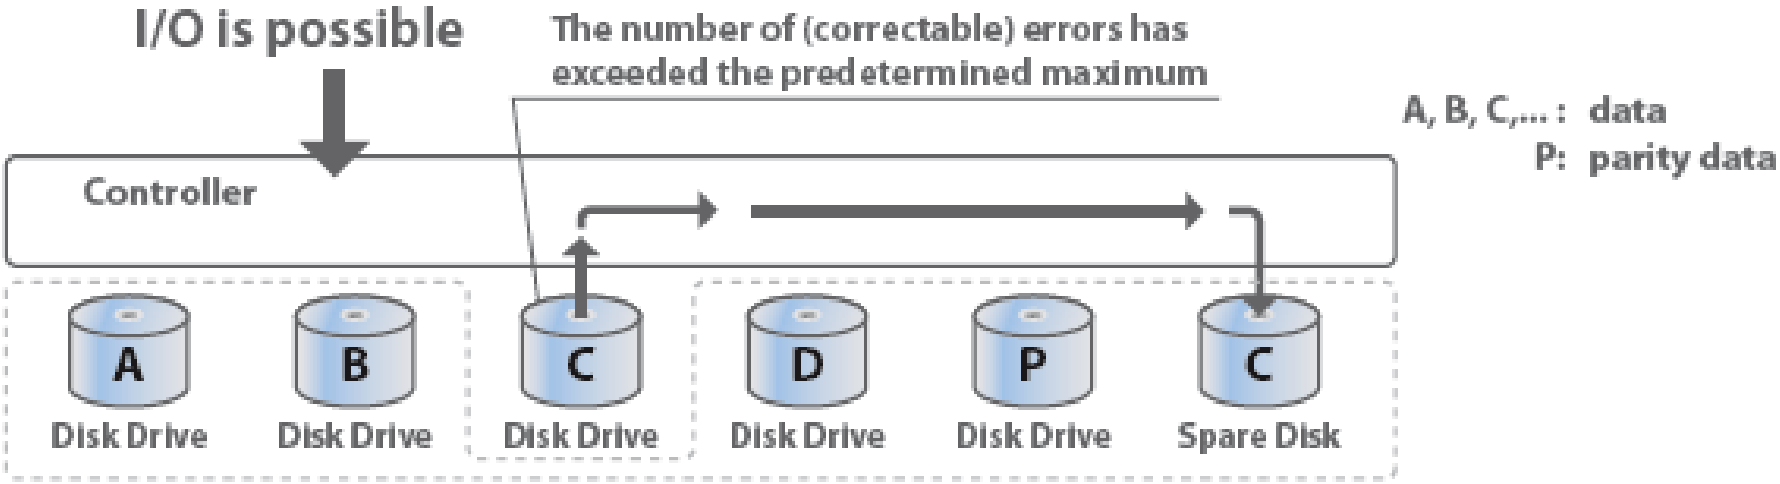

Correction Copy Parameters

- Copy Back

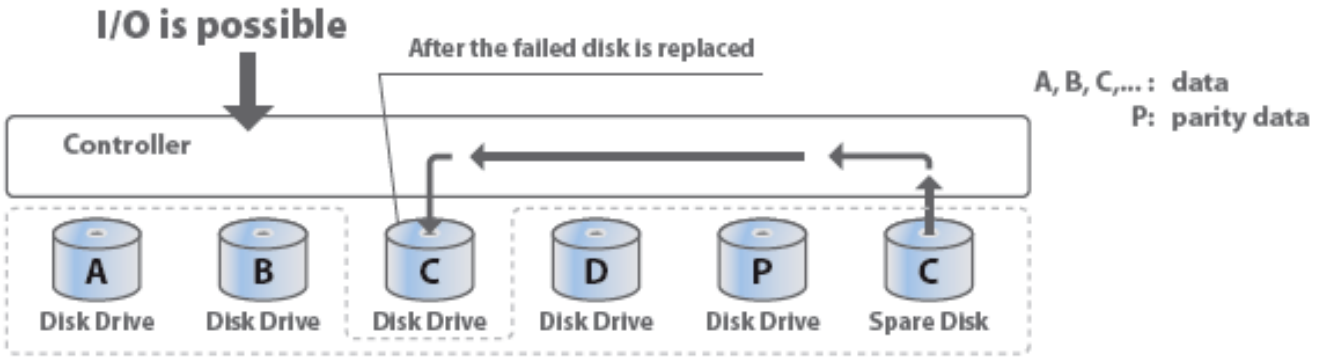

- No Copy Back

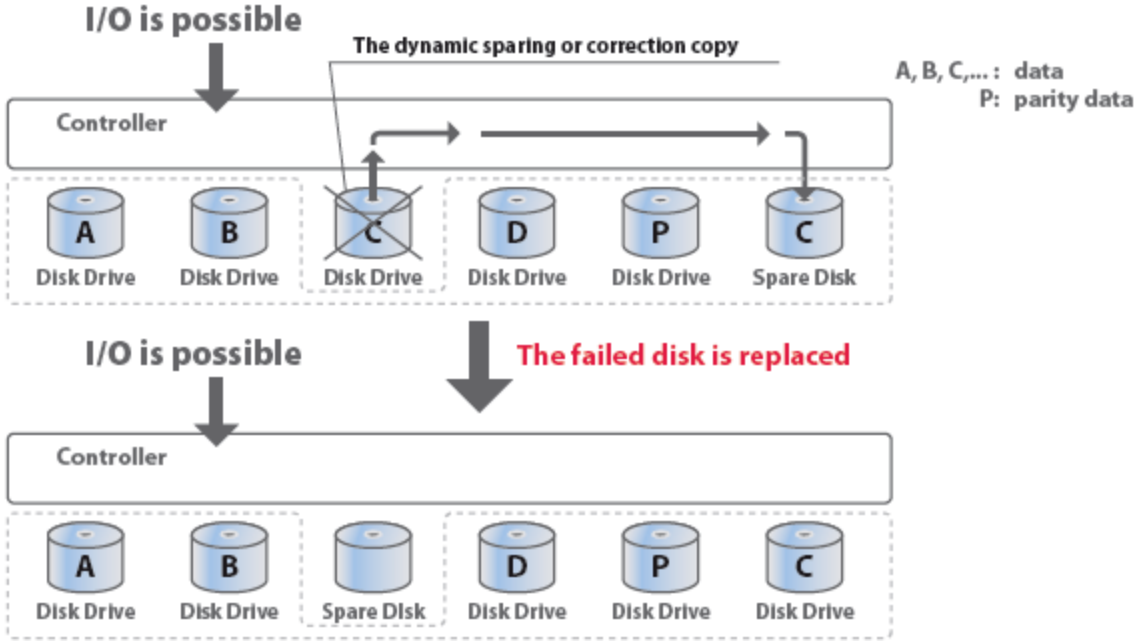

Exercise: RAID Configuration Options

- Considering the following data storage needs, which RAID options would provide the best performance?

1. Online transactional database (banking, stock market) where performance and reliability is key _____
2. Search engine, or catalog system for a library _____
3. Overnight billing and inventory system _____

If vILT class, write your answers on blank lines

Building a Midrange Storage System – Components

Expansion Unit / Disk Enclosure

A typical expansion unit or disk enclosure (based on SAS architecture), consists of following components:

- Individual hard drives (SSD, SAS, SATA)
- Expander (buses and wires for connecting drives together)
- Power supplies
- Cooling systems (fans)
- Chassis

Expansion Unit Connectivity

Each expansion unit is connected to both controllers

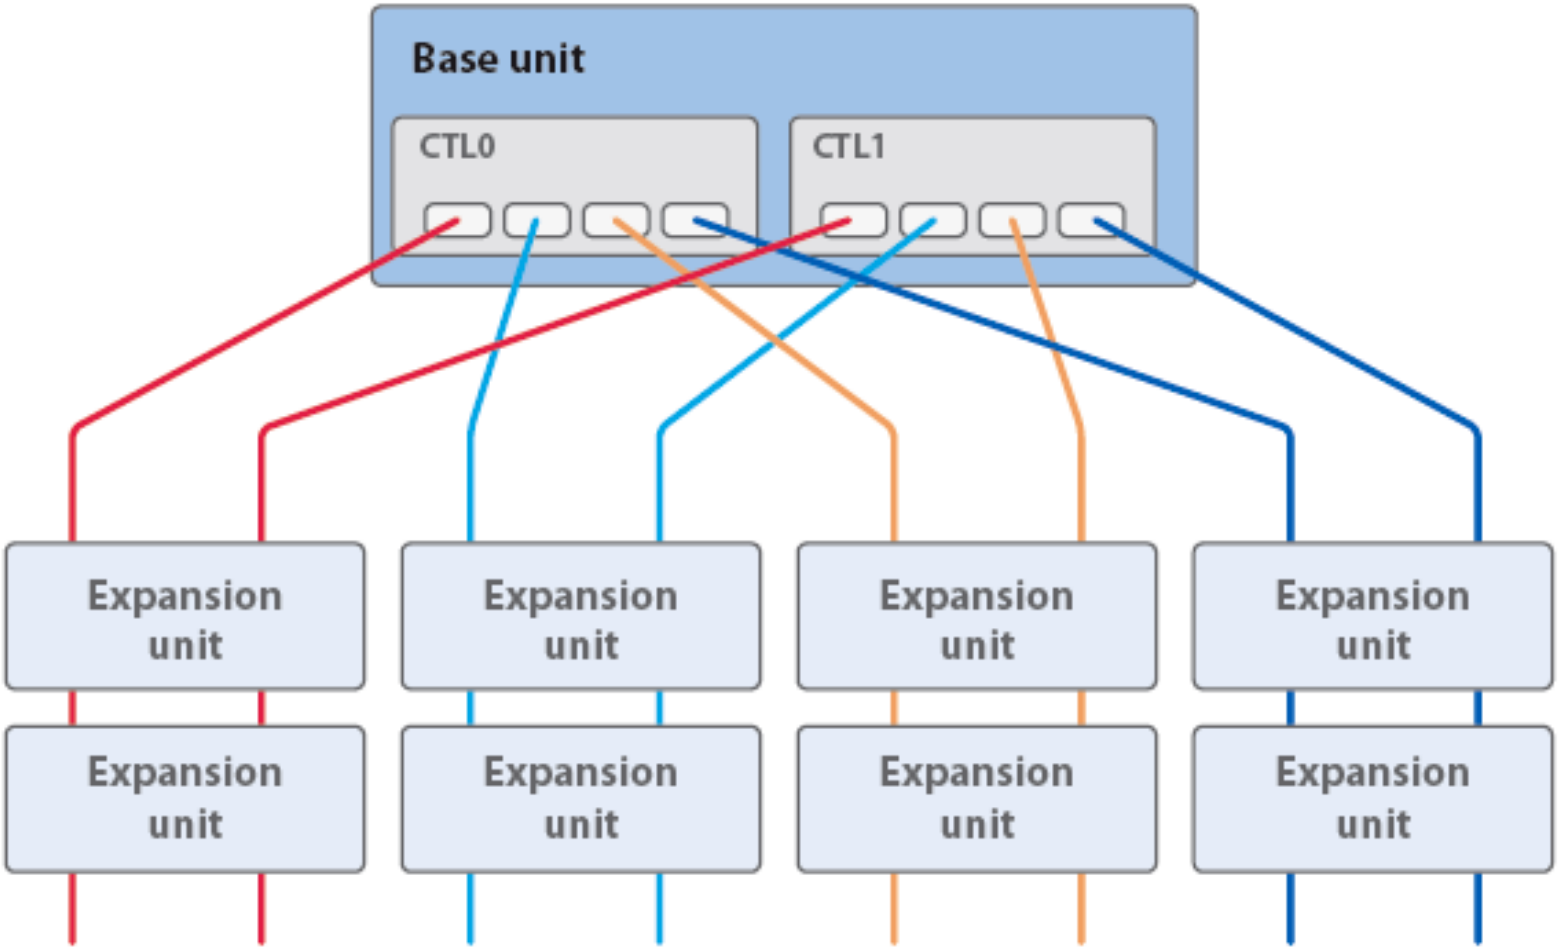

Back-End Architecture – Midrange Storage System

A midrange storage system contains main controller boards that are equipped with components that can be put into three categories:

- Front end (connection to hosts or other storage systems)
- Cache (works as a buffer, has major influence on performance)
- Back end (connections to hard drive enclosures, RAID operations)

Back-End SAS Architecture Example

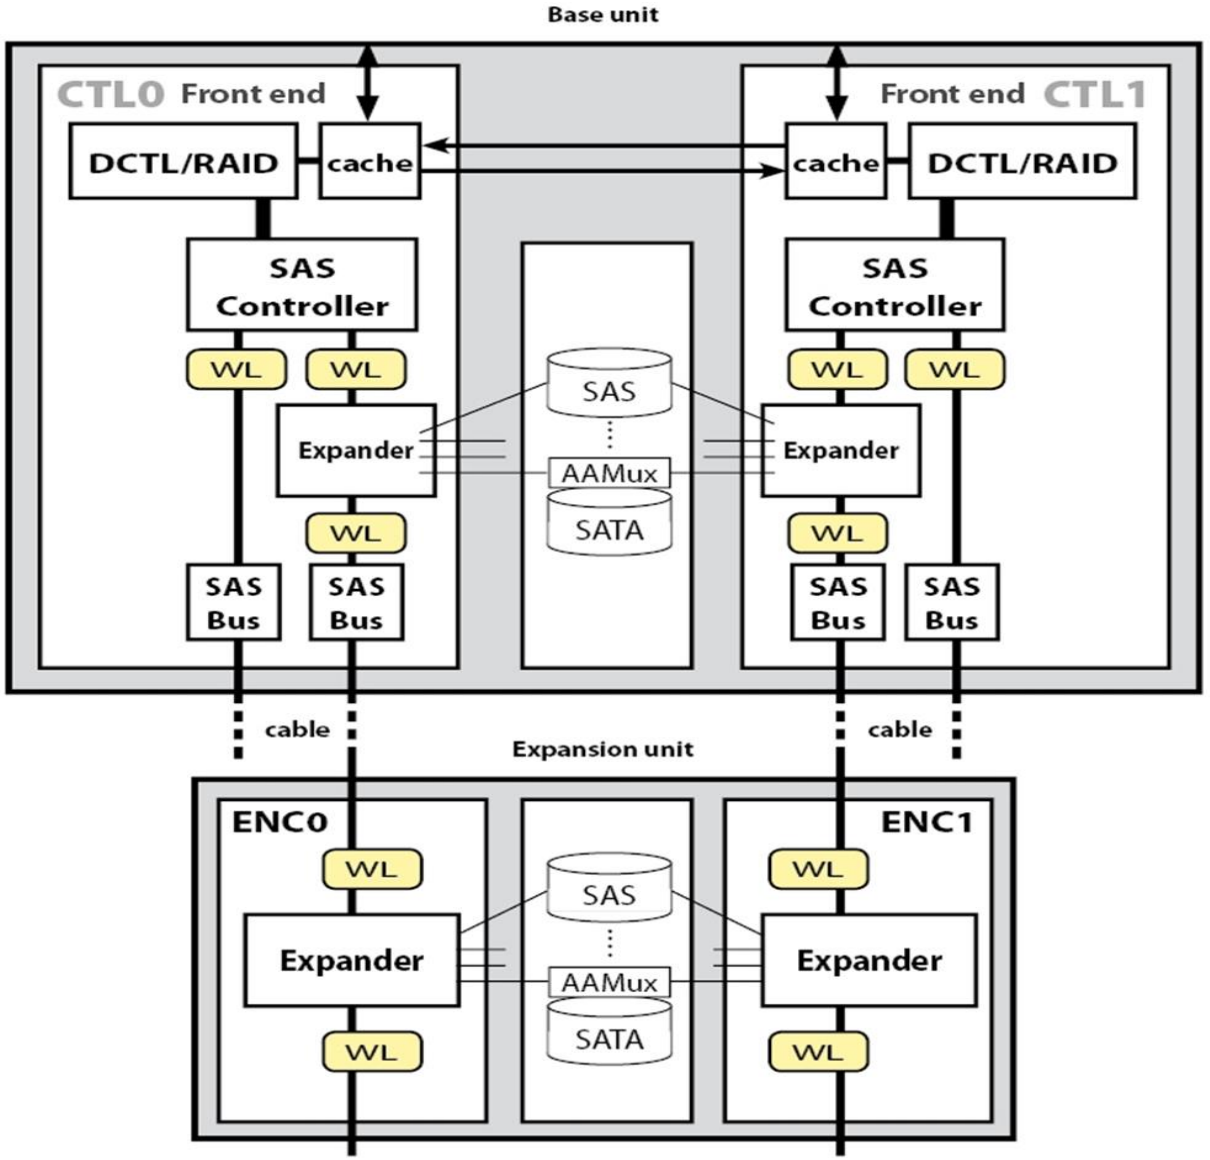

Cache

Cache

Cache is a temporary storage area for frequently used data. A system can access the cached copy instead of the data in the original location. This reduces the time taken to access data.

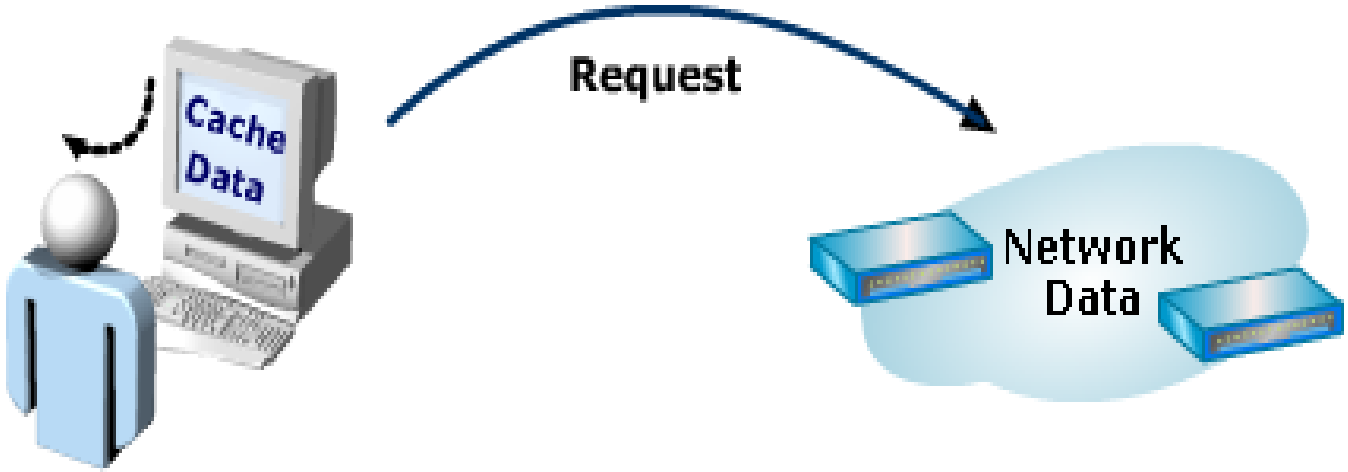

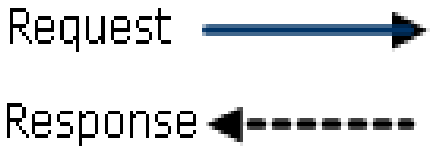

Cache Operations

The storage system's microcode contains algorithms that should anticipate what data is advantageous to keep (for faster read access) and what data should be erased. There are two basic algorithms that affect the way cache is freed up:

Least Recently Used (LRU) — Data that is stored in cache and has not been accessed for a given period of time (i.e., data that is not accessed frequently enough) is erased.

Most Recently Used (MRU) — Data accessed most recently is erased. This is based on the assumption that recently used data may not be requested for a while.

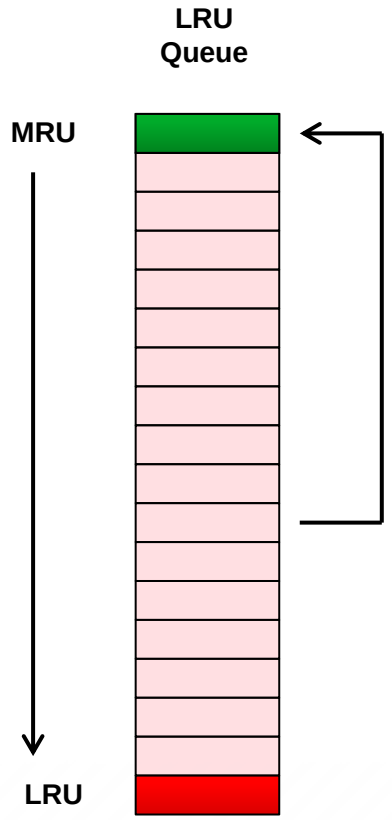

Cache Data Protection

- Cache Mirroring

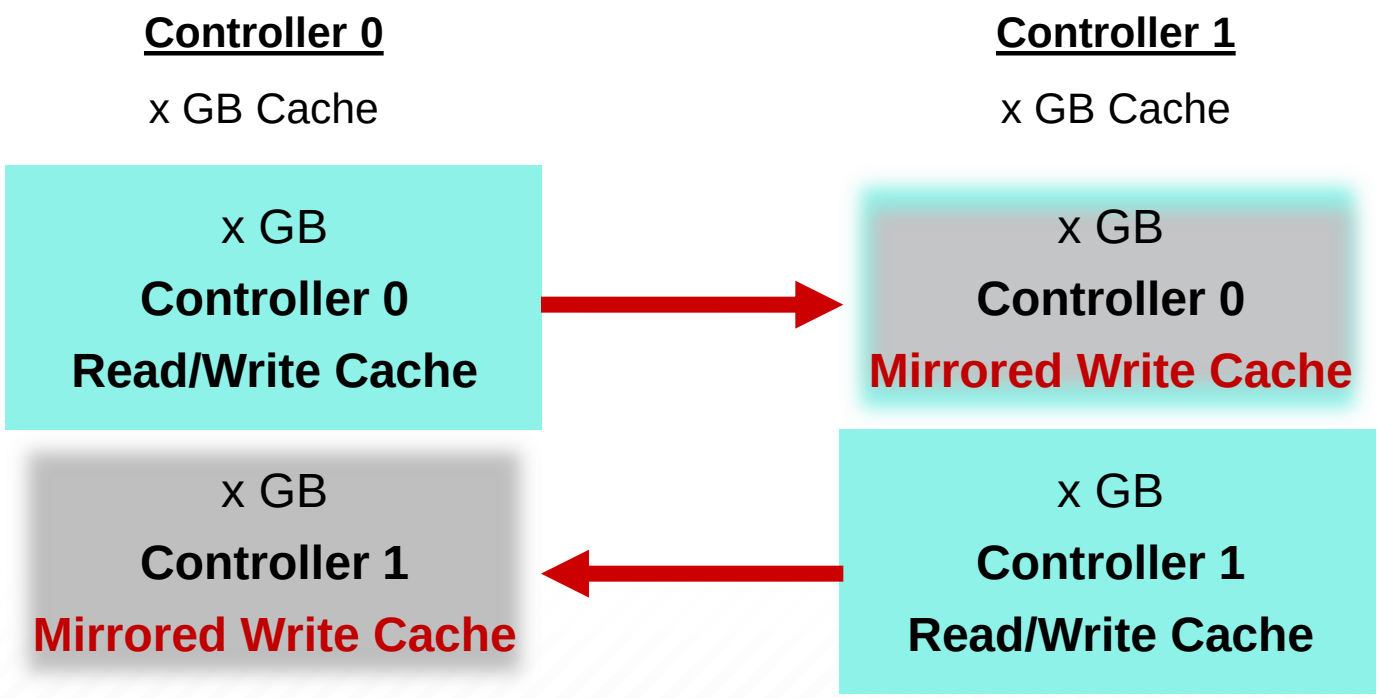

Front-End Architecture

- Interface Board – Ports connecting storage system to servers
- Protocols
  - Fibre Channel (FC) – Defines a multi-layered architecture for moving data; allows data transmission over twisted pair and over fiber optic cables
  - Fibre Channel over Ethernet (FCoE) – Encapsulation of Fibre Channel frames over Ethernet networks
  - iSCSI – Interface connecting storage system to the LAN; allows organizations to utilize their existing TCP/IP network infrastructure without investing in expensive Fibre Channel switches

Front-End Architecture – Other Components

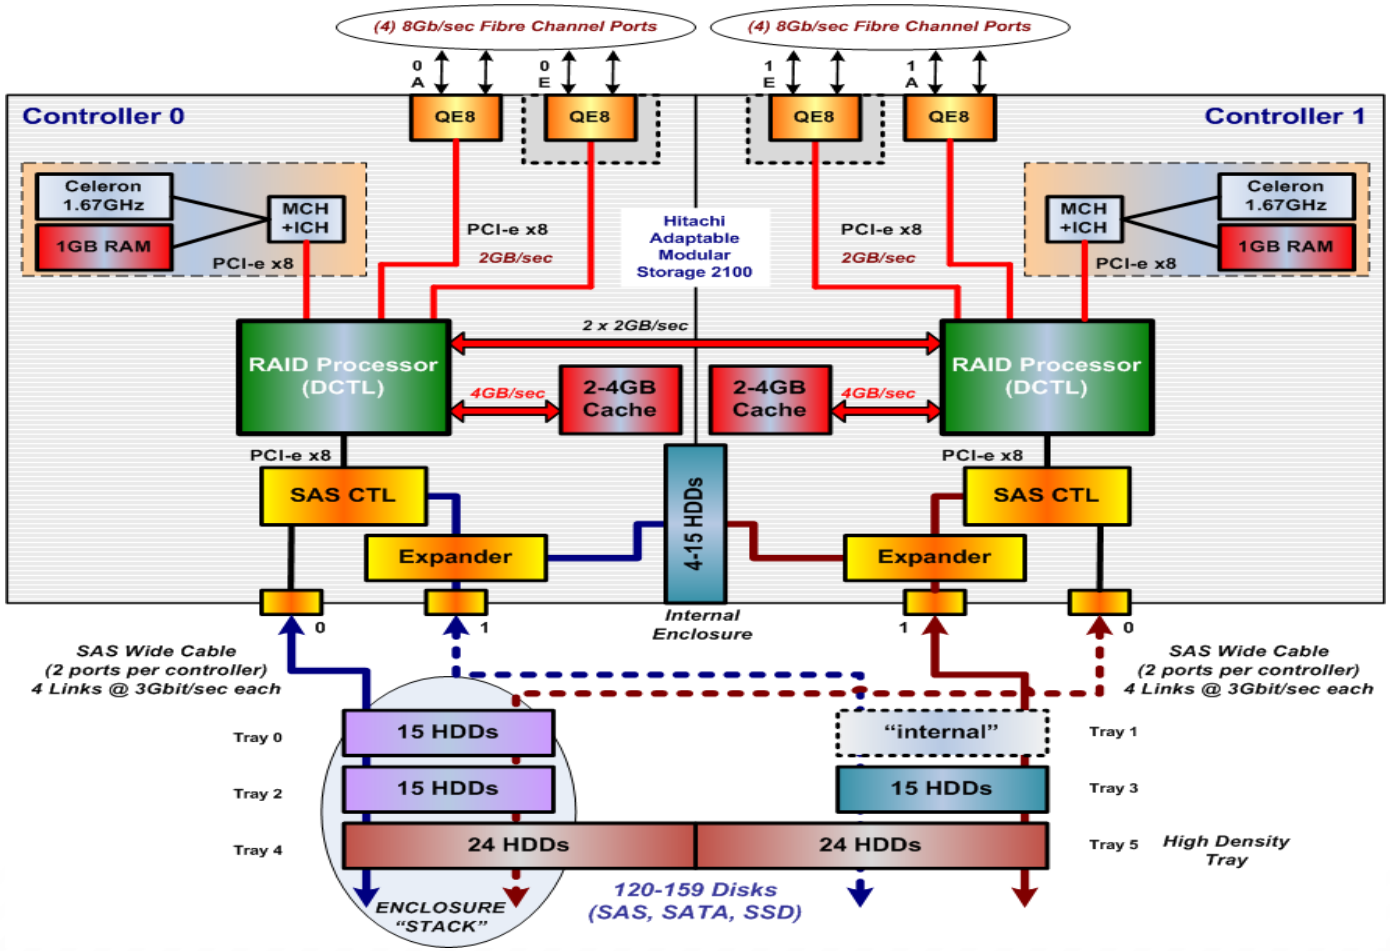

In the frontend, we have QE8 FC port controllers that are part of the interface board. These controllers are mainly responsible for conversion of FC transfer protocol into PCIe bus used for internal interconnection of all components.. Notice the CPU and local RAM memory (not cache).

Active-Passive Architecture

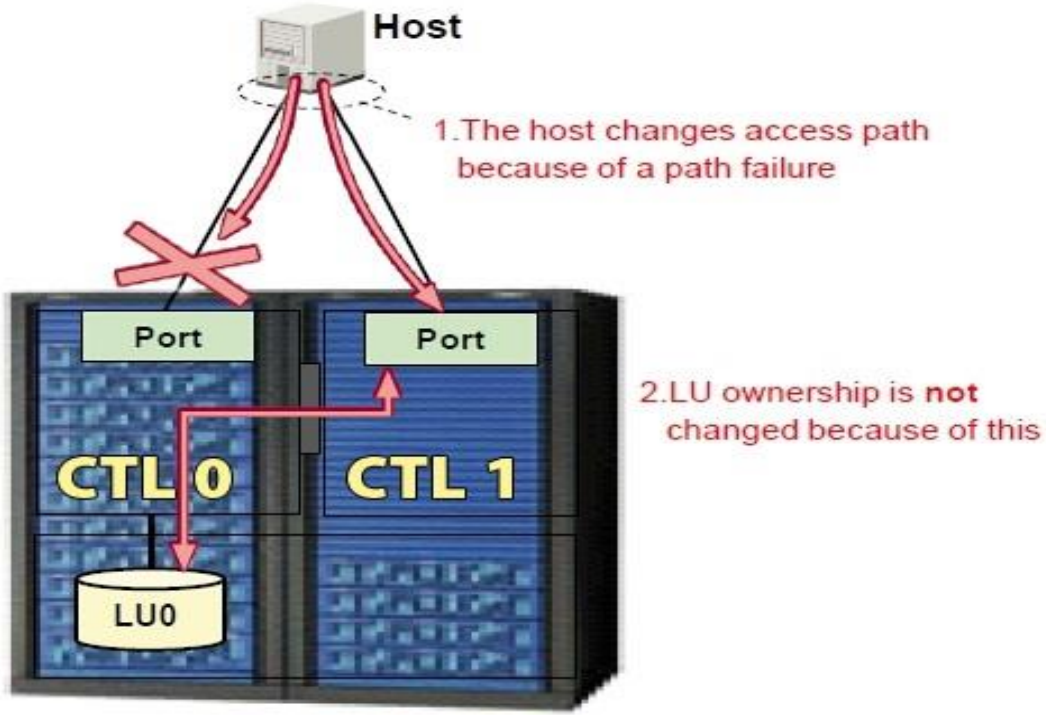

In the event of a path failure, LUN ownership does not change. Data is transferred via the backup path to CTL1 and then internally to CTL0, bypassing CTL0 front end ports. This diminishes performance because of internal communication.

Symmetric Active-Active Controllers

Multiple paths to a single LUN are possible. LUN ownership automatically changes in the event of path failure. Unlike active-passive architecture, active-active architecture offers equal access to the particular LUN via both paths. This means the performance is not influenced by what path is currently used.

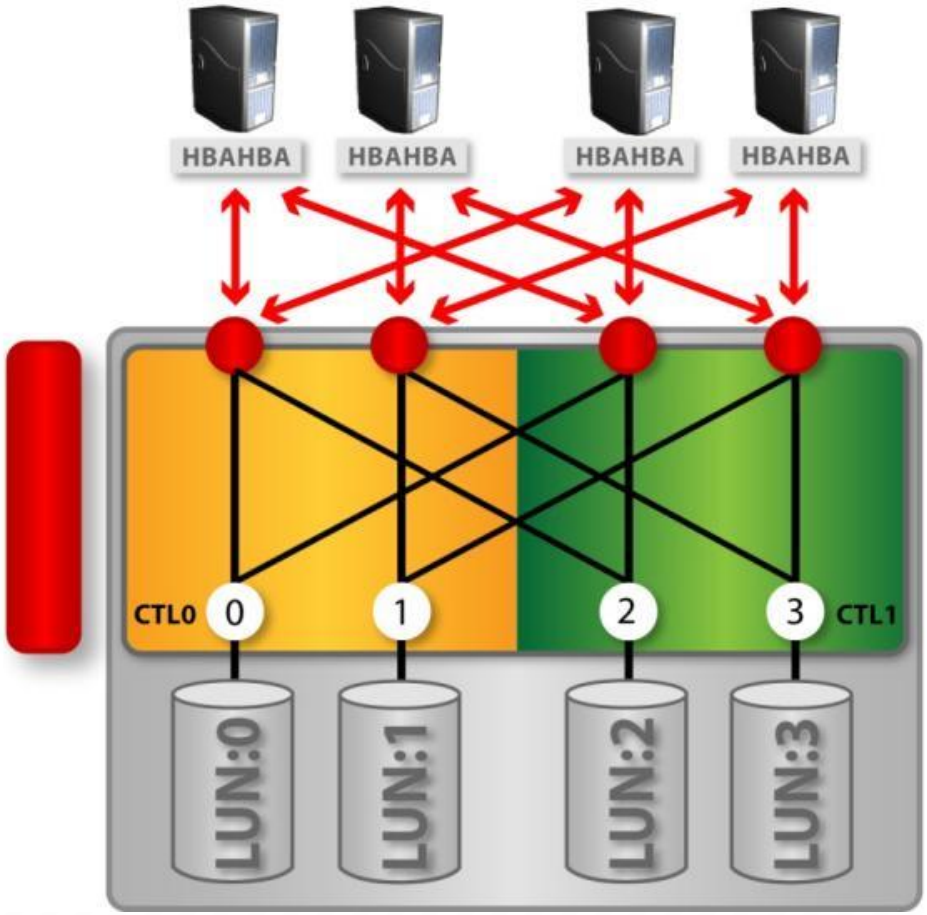

Controller Load Balancing – Active-active Architecture

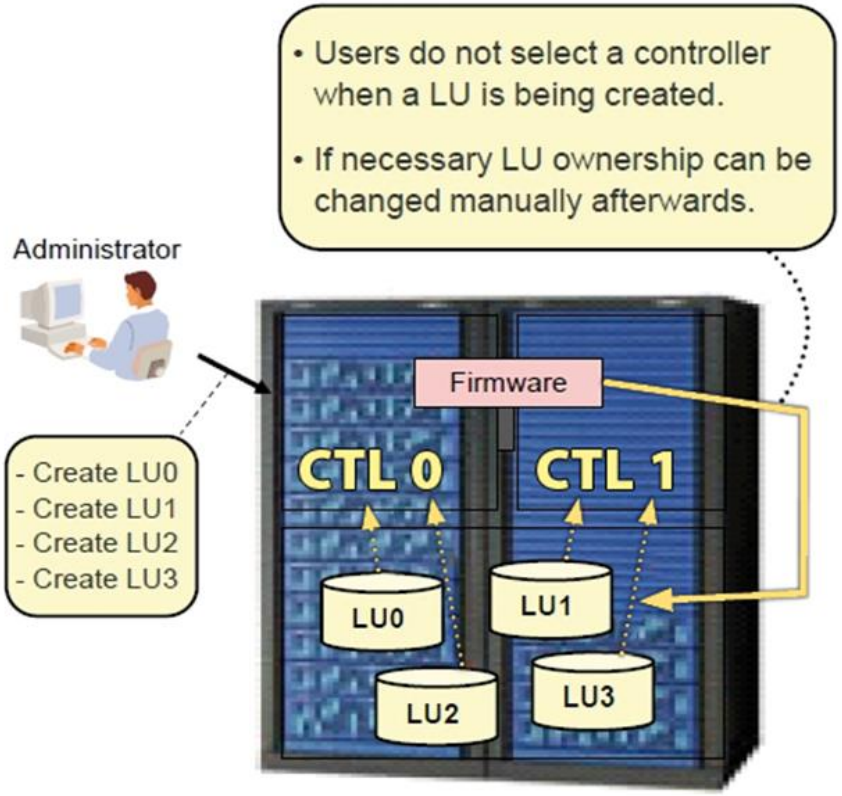

No need to configure LUN ownership manually. Communication goes either via CTL0 or CTL1. In the event of path failure, no bypassing is necessary; therefore there is no communication overhead.

Fibre Channel Ports and their Configuration

Host connected to a Target port

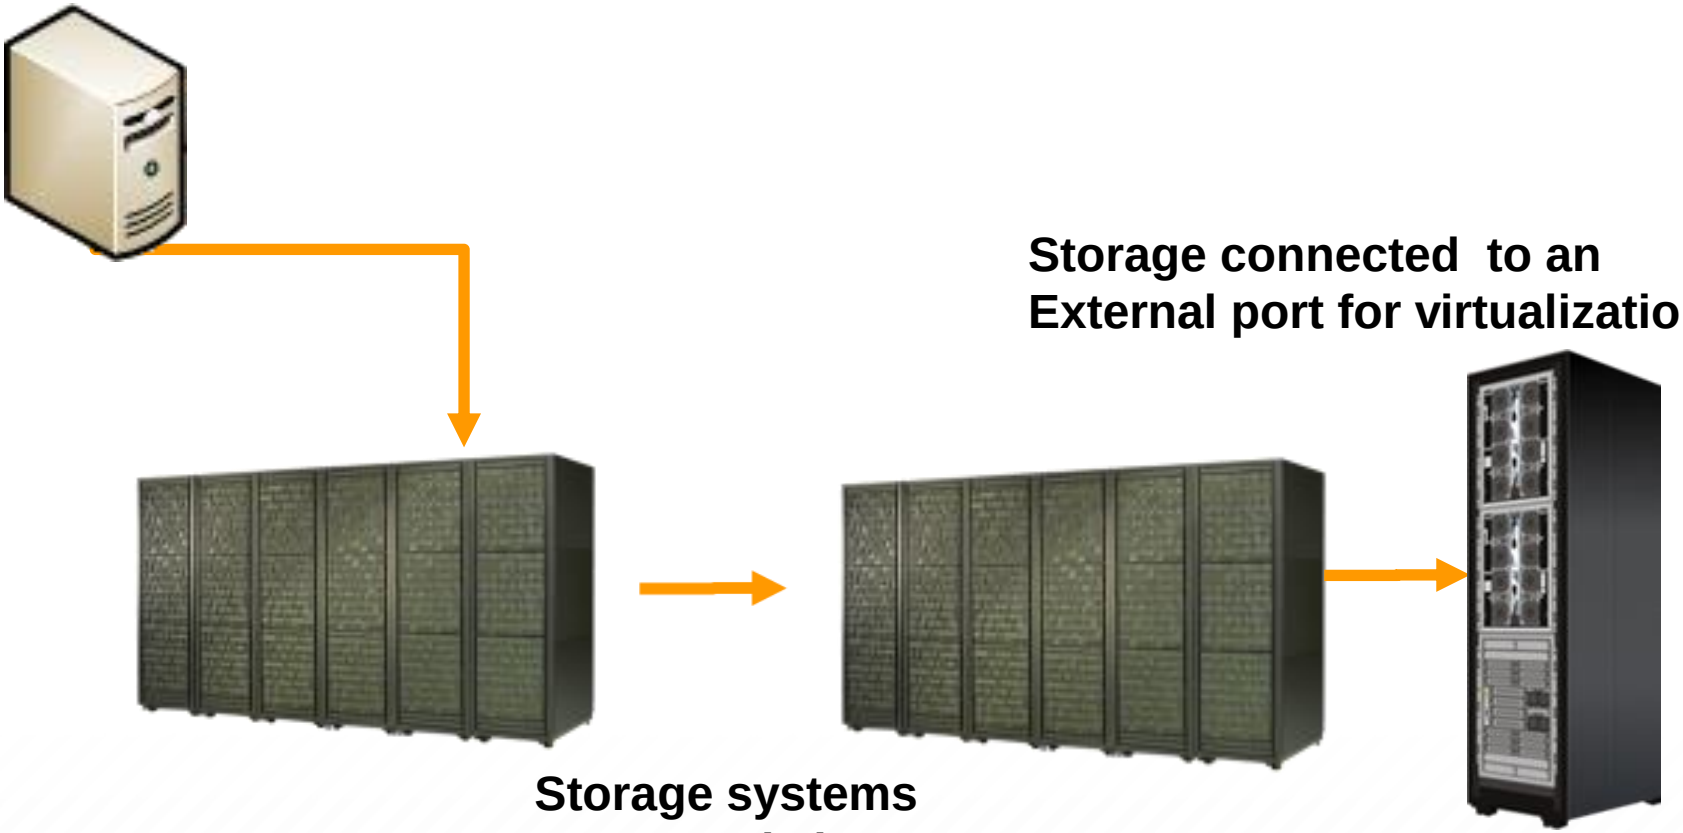

Storage connected to an External port for virtualization

Storage systems connected via an Initiator port for replication

iSCSI Interface

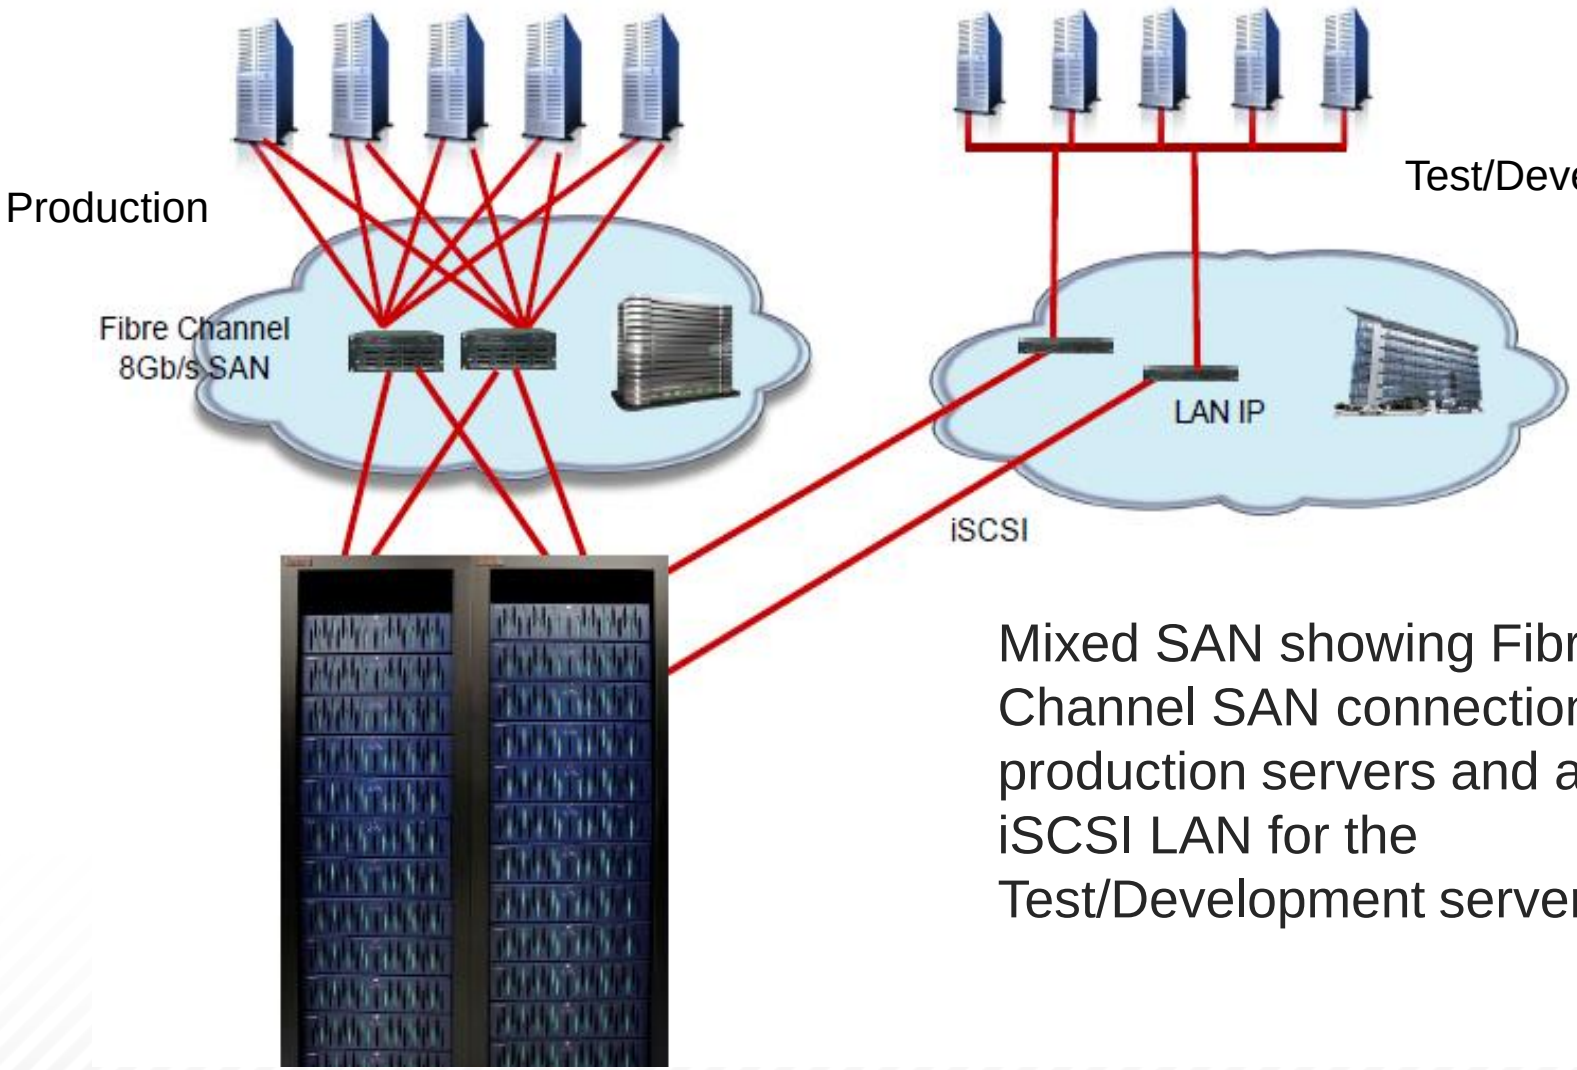

Mixed SAN showing Fibre Channel SAN connection for production servers and an iSCSI LAN for the Test/Development servers

Module Summary

- Upon completion of this module, you should have learned to:
  - Identify and describe the major components of a storage system
  - Explain the different RAID levels and configurations
  - Describe the midrange storage system architecture and its components
  - Explain the factors directly influencing the performance of a storage system
  - Understand and explain the differences between a midrange and an enterprise storage system

Storage Concepts

Business Continuity and Replication

Module Objectives

- Upon completion of this module, you should be able to:
  - Describe basic concepts of business continuity and replication
  - Understand business impact analysis and risk assessment
  - Describe basic concepts of disaster recovery

Module Topics

- Business Continuity concepts
- Business Impact Analysis and Risk Assessment
- Back-up strategies and their implementation
- How Business Continuity and Disaster Recovery connect to IT
- Replication options
- Concepts of clusters and geoclusters

Business Continuity Concepts

Exercise: Data Center Disaster

- Tornado destroys ABC Corporation's data center
  - What needs to be prepared ahead of time?

Set up a Business Continuity team

  - Impacts on the business?

Write your answers on blank lines or send your answers via the CHAT

Exercise: Data Center Disaster

- Tornado destroys ABC Corporation's data center
  - What needs to be prepared ahead of time?

Set up a Business Continuity team

A Disaster Recovery Plan

Another data center location

Suppliers ready to replace damaged equipment

Employees trained in responding to these situations

Staff in place at the alternate site to take over

Failback to main data center

- Impacts on the business?

Financial loss

Damaged reputation and company image

Business Continuity Management: Regulatory Requirements

- Identifies critical business processes and establishes rules and procedures
  - Implementation governed by regulations and standards

Sarbanes Oxley – Corporate reporting and financial results; tells CEO and CFO that they must be able to defend accuracy of their books

Basel II – International banking regulation that deals with amortizing cost of risk into financial markets

British Standard For Business Continuity Management (BS25999) – guidance for determining business processes and their importance

North American Business Continuity Standard – (NFPA1600)

Business Continuity Planning

- The implementation of business continuity concepts with respect to the particular organization. The planning normally includes:
  - Business Impact Analysis: Identifies which business processes, users and applications are critical to the survival of the business
  - Risk Assessment : Determines probability of threats to an organization
  - Policies: Aligns BC policies with the company's business strategy
  - Business Recovery Plan: Defines procedures to be taken when a particular situation occurs
  - Disaster Recovery Plan: Describes how to get the critical applications working again after an incident takes place
  - Testing and Training Schedules: Provides a timeline for business continuity plan testing

Business Impact Analysis – RPO and RTO

- Recovery Point Objective (RPO) – Worst case time between last backup and interruption time
  - Represents how much data must be recovered
  - How much can you afford to lose?
- Recovery Time Objective (RTO) – How long is the customer willing to live with downed systems?
  - Represents outage duration

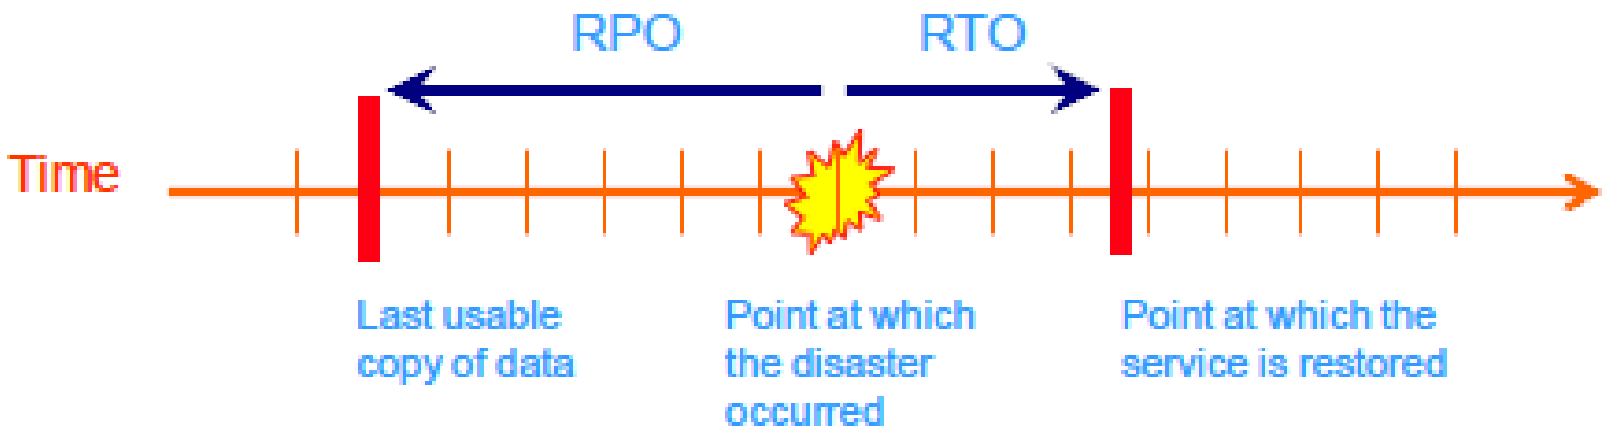

Business Continuity versus Disaster Recovery

Disaster Recovery – Part of business continuity that focuses only on **IT Infrastructure**

Disaster Recovery Plan contains:

- **Basic information** – purpose, area of application, requirements, log of DR plan modifications, members of DR team, their roles and responsibilities
- **Notification/activation phase** – notification procedure, call tree, damage assessment, activation criteria, plan activation
- **Recovery procedures** – succession of recovery procedures according to their importance, logging, escalation
- **Standard operation resumption** – checking whether all systems work properly, termination of DR plan
- **Amendments** – call book, vendor SLA, RTO of processes

Data Backup and Data Replication Concepts

Exercise: Data Backup and Data Replication

- Is data backup different from data replication?

Backup =

Replication =

Write your answers on blank lines or send your answers via the CHAT

Example: Data Backup Configuration

Back-up over LAN is the simplest solution. The back-up server pulls data from production servers and then it sends data to a tape library or NAS device.

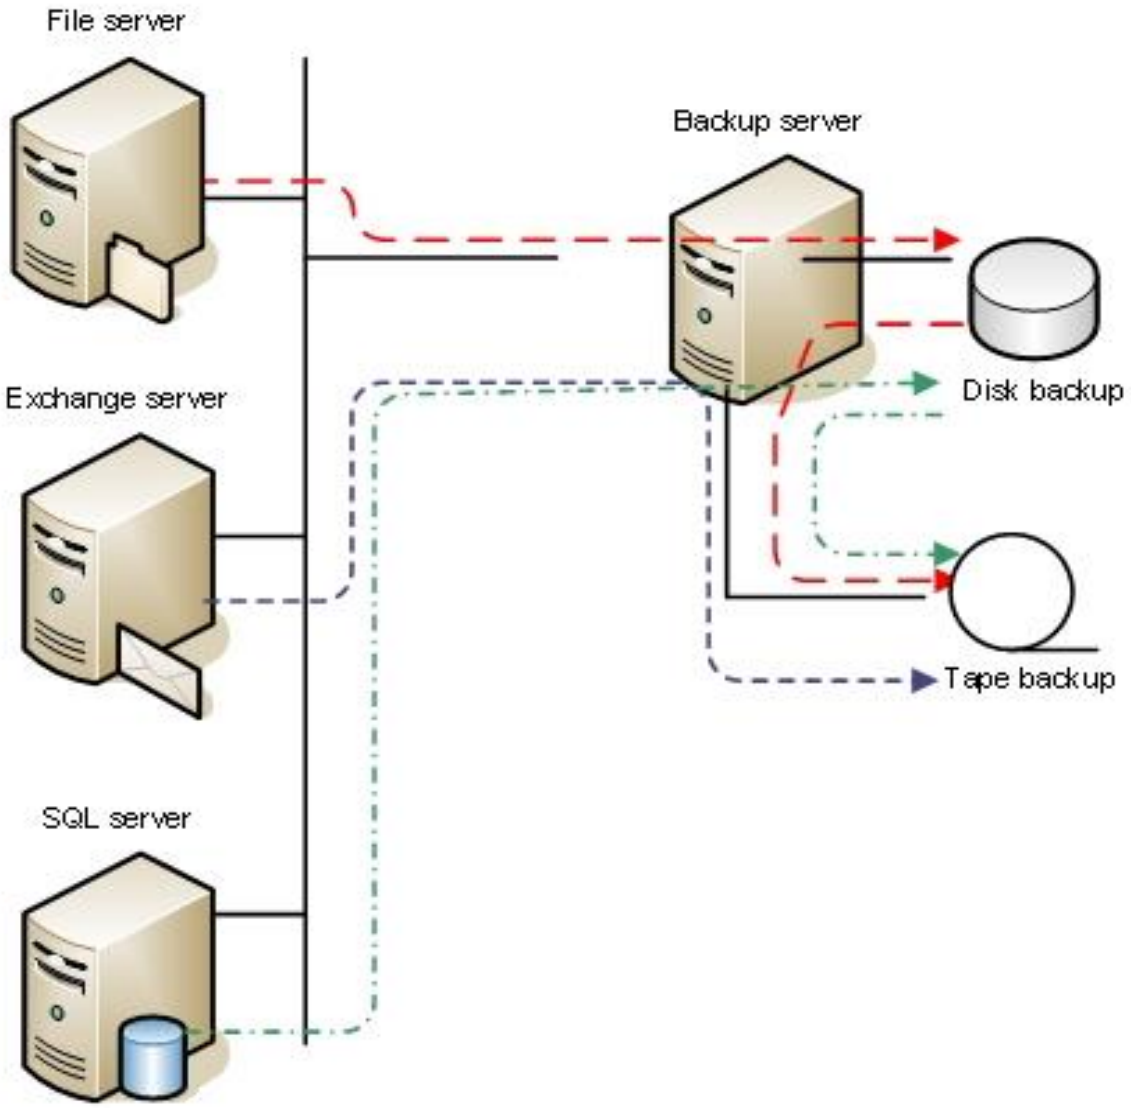

Data Backup Models

- **Full Backup**: Data is stored in exact copies
- **Incremental Backup**: Only the data that was changed or added since last time backup is recorded
- **Differential Backup**: Only the data that differs from the initial full backup is recorded
- **Reverse Delta Backup**: At every scheduled backup, the initial full backup image is synchronized so that it mirrors the current state of data on servers

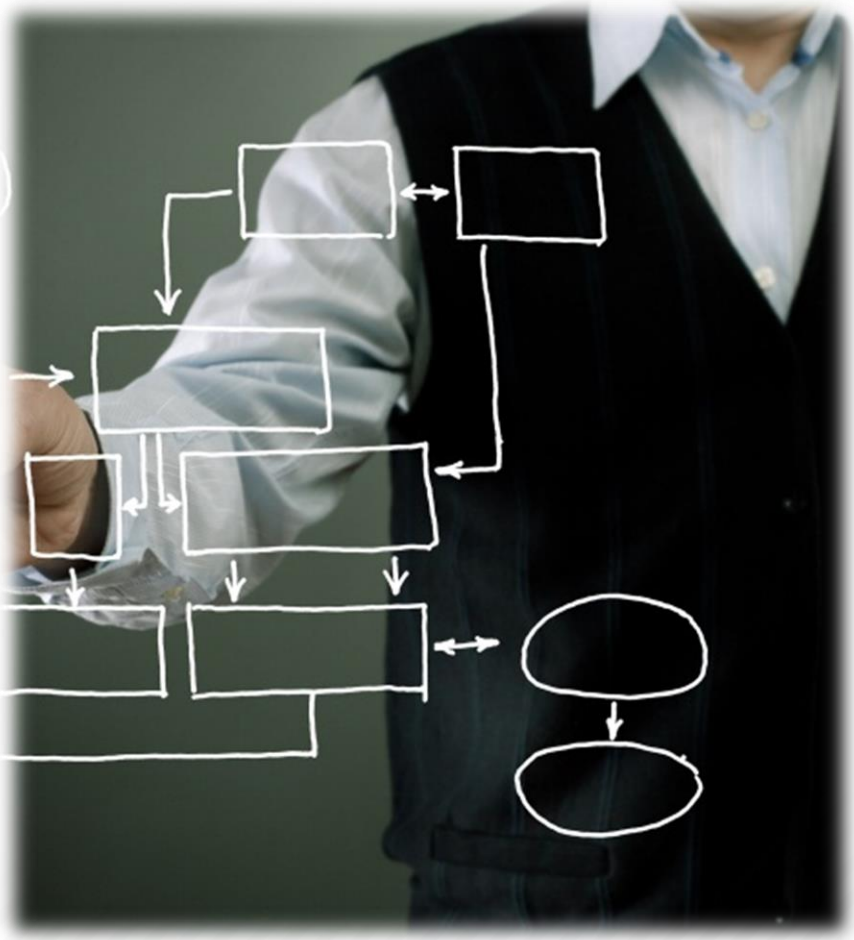

Backup Requirements

To back-up the data from production servers you need:

- **A back-up device** – in enterprise environment it will most probably be a tape library, but it can also be a storage system with a LUN dedicated for back-up.
- **A back-up server** – in most cases you need a back-up server that is communicating directly with the back-up device and that controls back-up from all the servers
- **Back-up software** – software that runs on a back-up server and that allows to make configuration according to your needs
- **Back-up agents** – these agents are small applications installed on all your production servers. They are part of back-up software and they allow communication between a back-up server and production servers.

Backup Optimization

Techniques used to achieve better backup utilization:

- **Compression** – the output archive file is smaller than total of the original files
- **Deduplication** – the technique that eliminates duplicities in data
- **Multiplexing** – the ability of software and equipment to back-up data from several sources simultaneously
- **Staging** – back up to a disk first and then transfer the data from this disk to tape – known as Disk-to-disk-to-tape (D2D2T)

Data Replication Overview

- A volume with source data is called a Primary Volume (P-VOL), and a volume to which the data is copied is a Secondary Volume (S-VOL).
  - In-system – all operations with logical units (LUs) within the same storage system
  - Remote – all operations with LUs across different storage systems

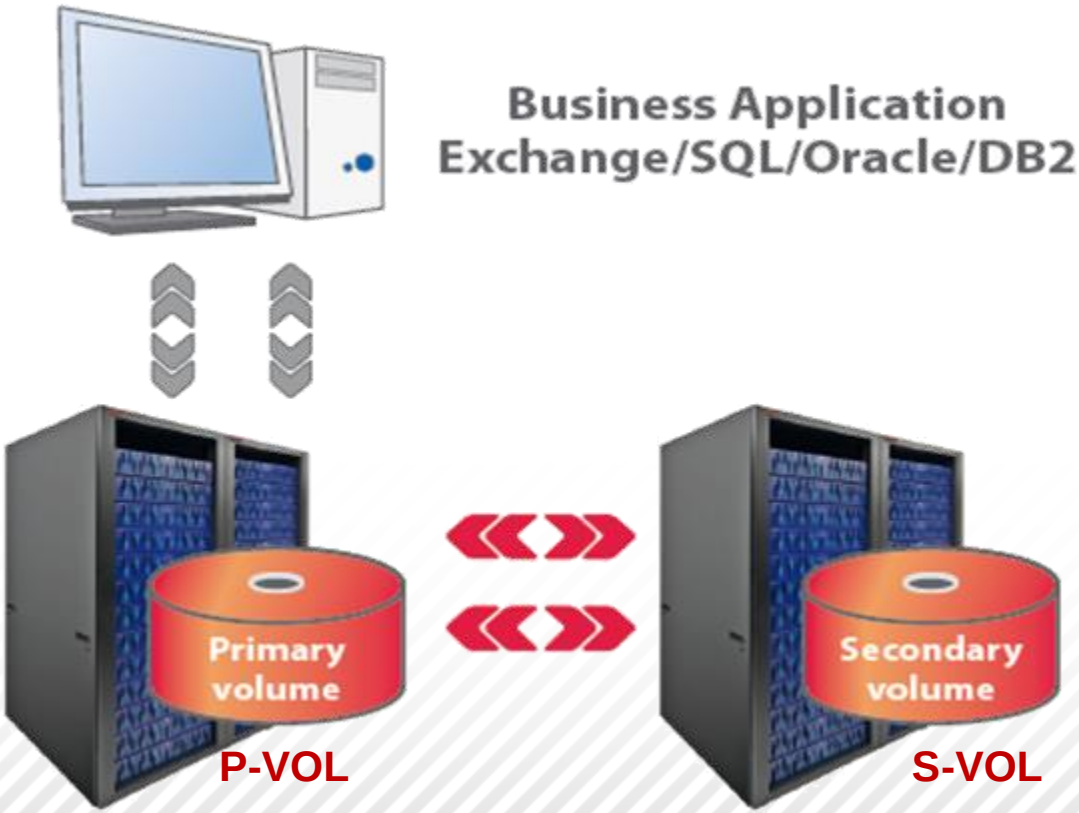

Data Replication Overview

Data Replication

Within and between array heterogeneous replication

- Data replication (or protection) provides operational and disaster recovery
  - Replicates data within or between storage systems without disruption
  - Creates multiple protected copies from each source volume
  - Can run independent of host OS, database, file system
  - Mirrors image of data
  - Offers quick restart and recovery in disaster situations
- Once created, copies can be used for:
  - Data warehousing or data mining applications
  - Backup and recovery
  - Application development

Exercise: Data Backup Or Data Replication?

- In the event of a major disaster, would you use your backup solution or data replication solution to recover your *business critical* online applications?
  - Break up into teams, discuss, and present your reasons for choosing one of the solutions over the other.

If a vILT class, present your answers over the phone

Data Replication – Copy Operations

- Data copy operations:
  - **Initial Copy**
    - All data is copied from P-VOL to S-VOL
    - Copies everything including empty blocks
  - **Update Copy**
    - Only differentials are copied

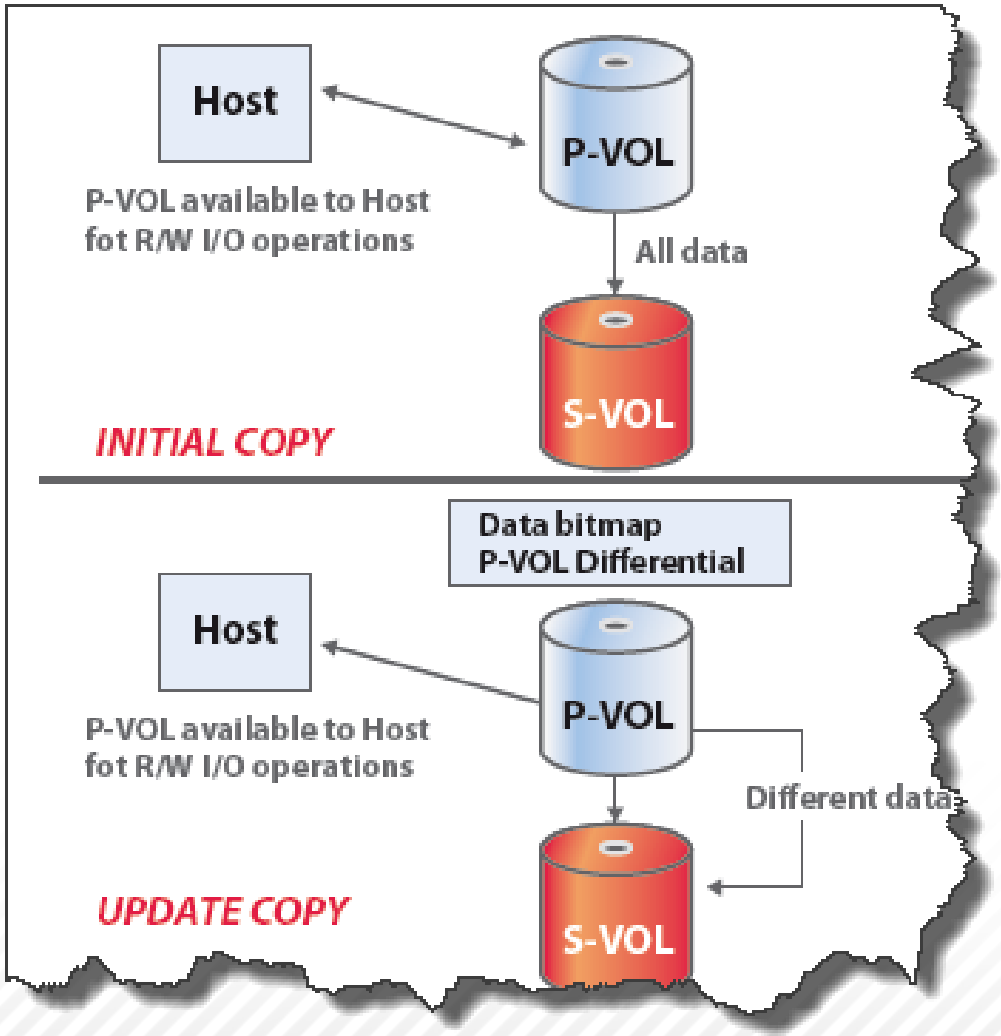

Requirements for All Replication Products

- Any volumes involved in replication operations (source and destination) should be:
  - Same size (in blocks)
  - Must be mapped to a port
    - Source can be online and in use.
    - Destination must not be in use/mounted.
- Intermix of RAID levels and drive type is supported.
- Licensing
  - License is capacity independent.

Data Replication Operations – Establish Pairs

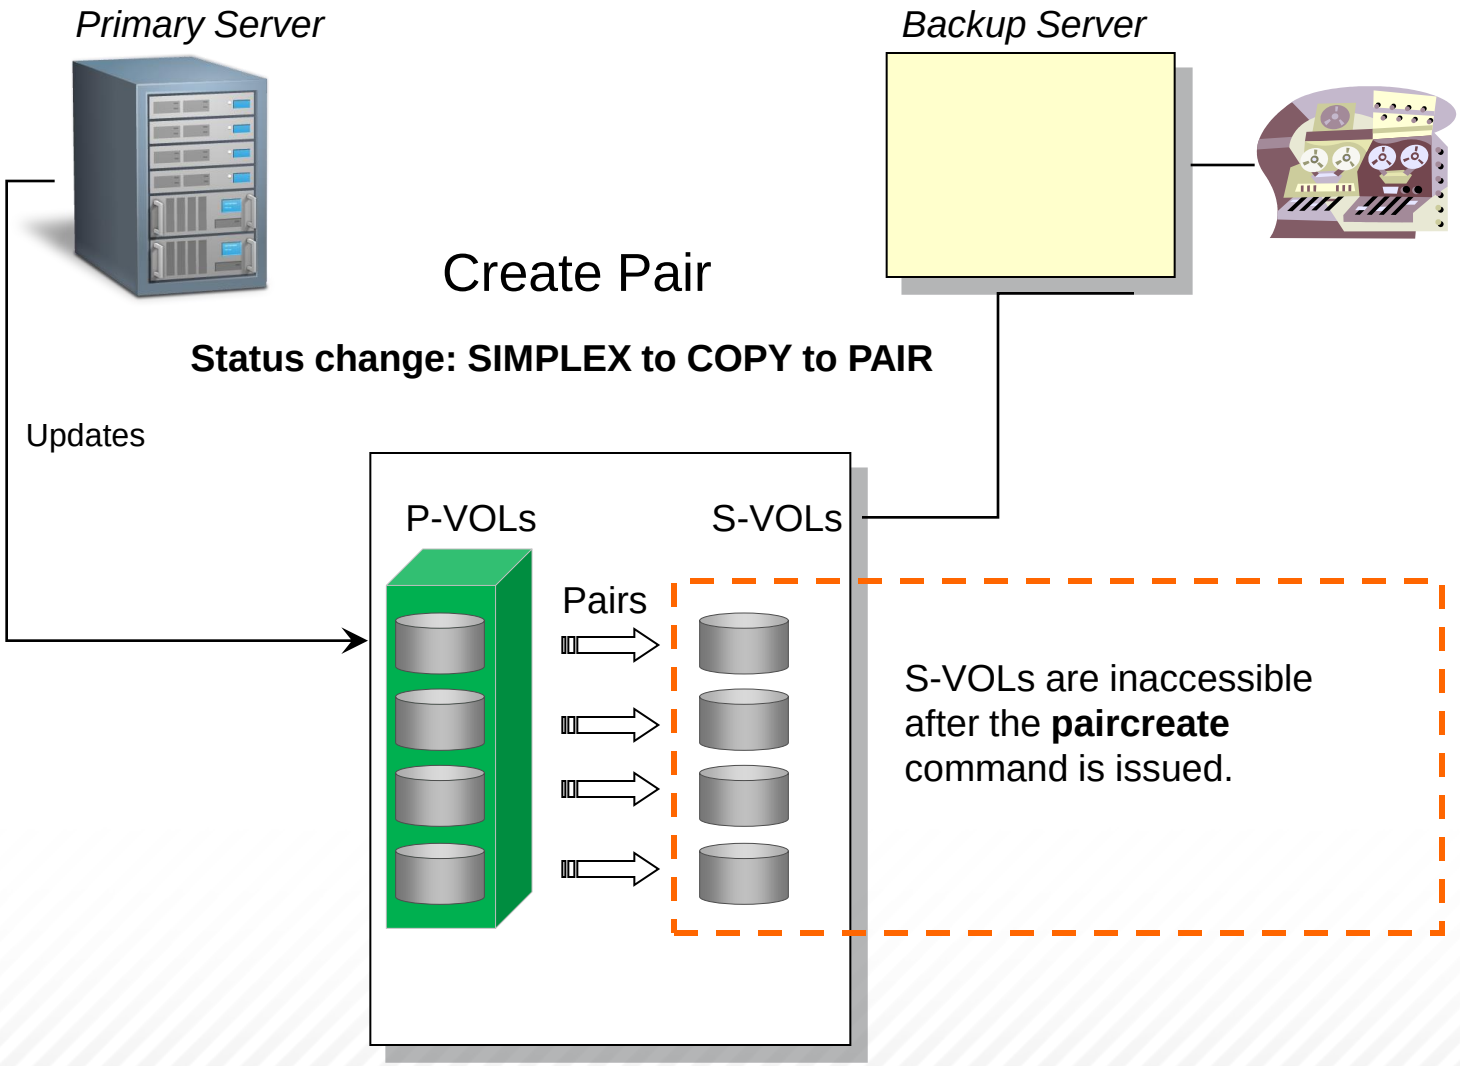

Paircreate

Data Replication Operations – Split Pairs

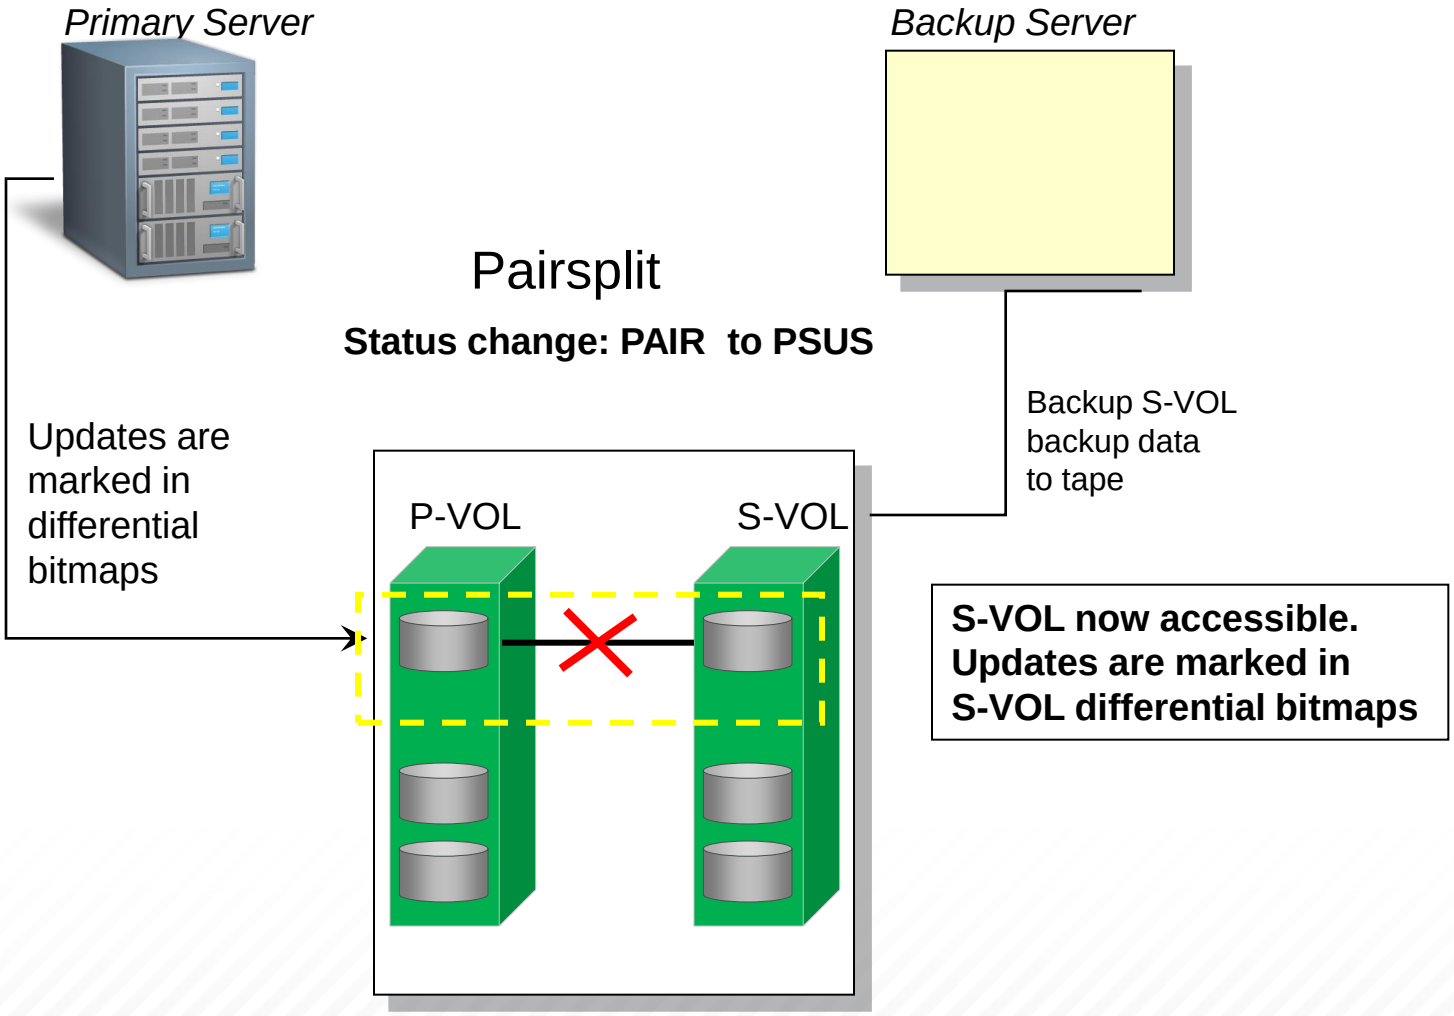

Data Replication Operations – Resynchronizing Pairs

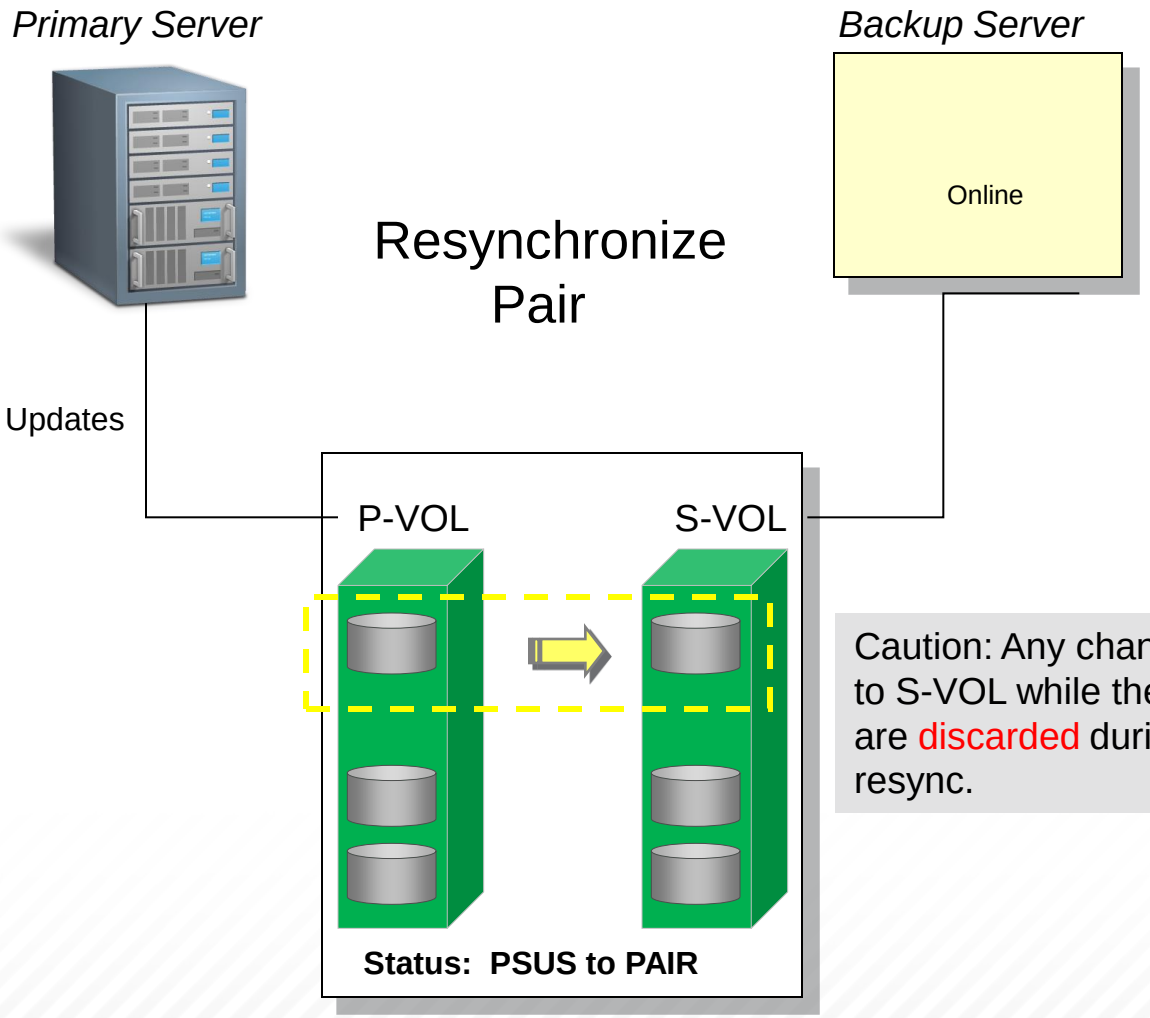

PAIRRESYNC

Normal Resync

Reverse Resync

Caution: Any changes applied to S-VOL while the pair is split are discarded during NORMAL resync.

Data Replication Operations – Quick Restore

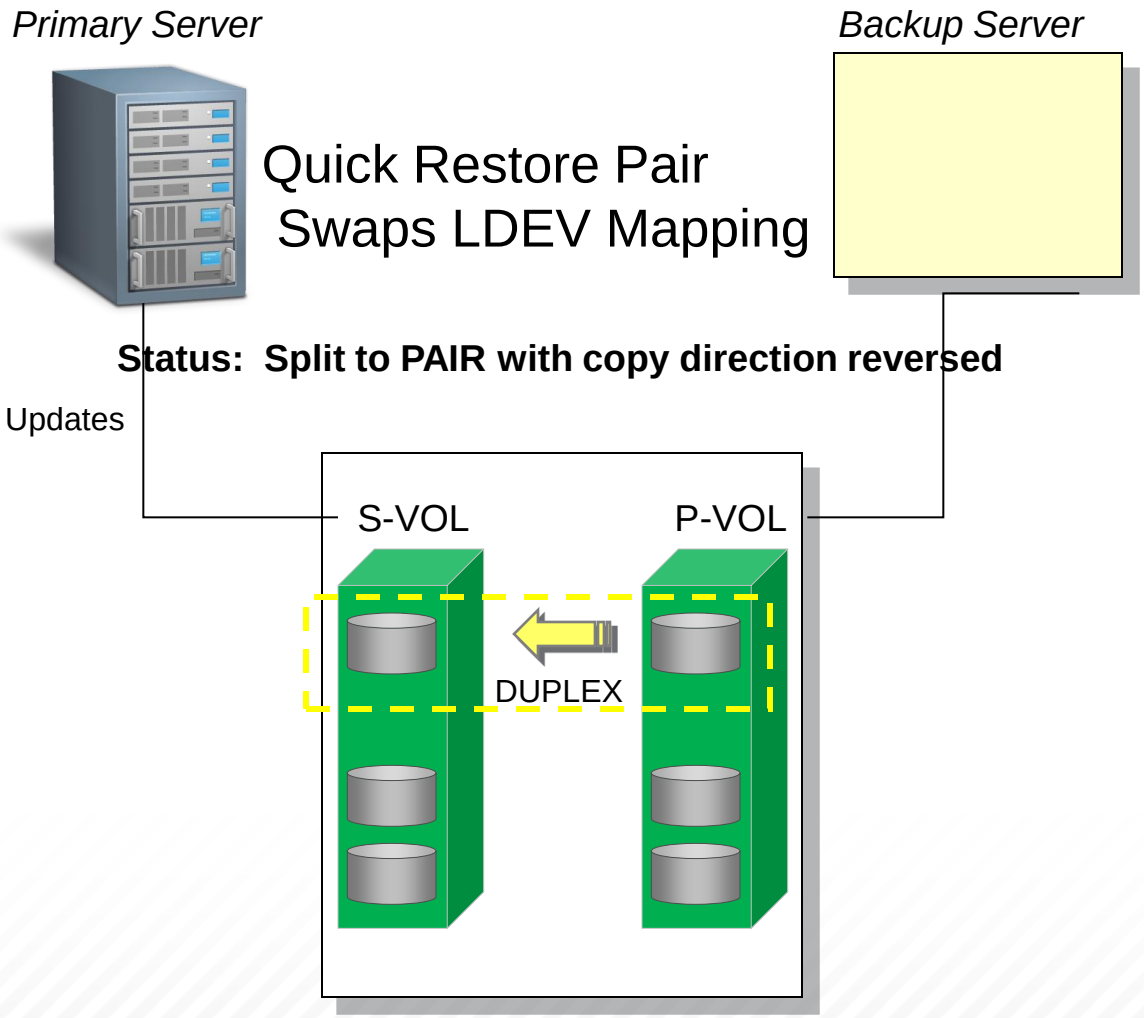

Quick Restore

In-System Replication

In-System Replication

- In-system hardware-based copy facility that provides:
  - Full volume point-in-time copies
  - Host Independent
  - Nondisruptive replication
  - Clone images are RAID protected copies

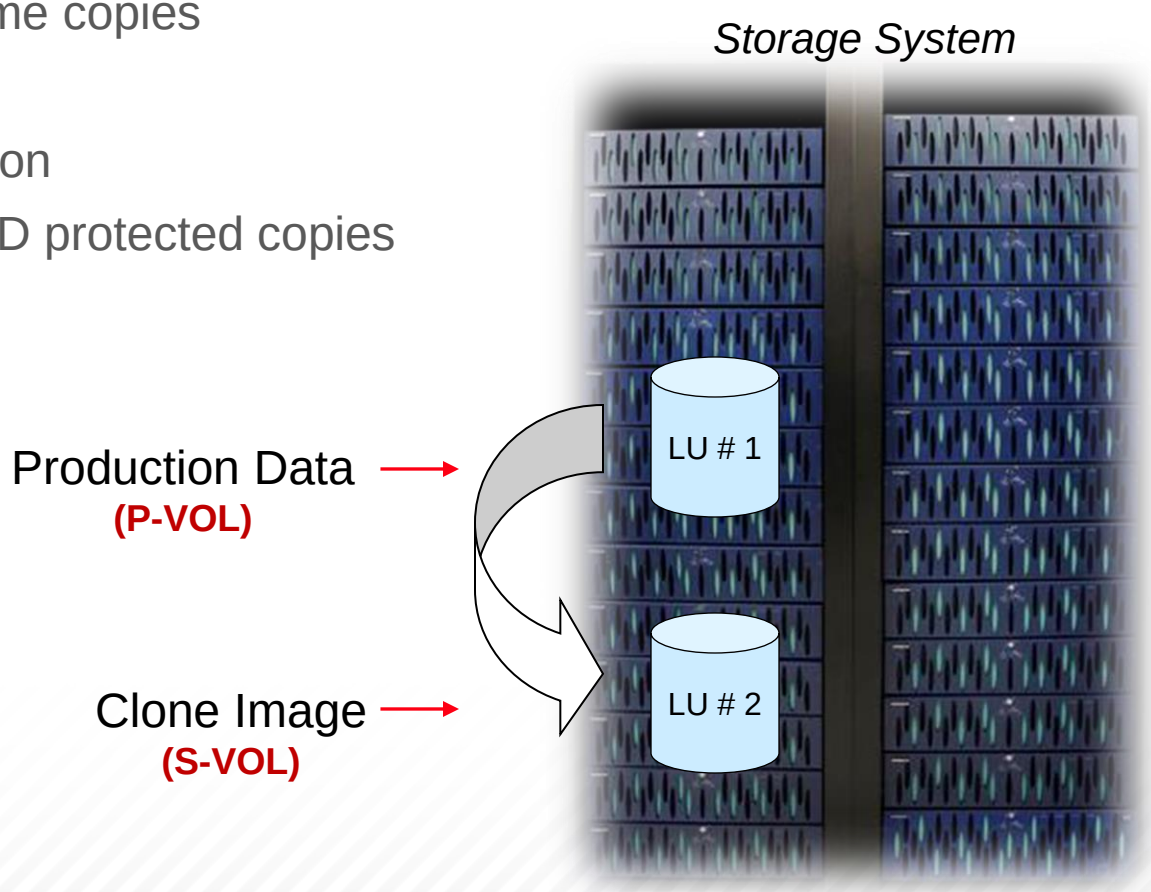

In-System Replication: Copy-On-Write Snapshots

- Snapshots
  - Create a point-in-time (PiT) copy or “snapshot” of the data
  - Uses less space than full copies or clones Frequent, cost effective, point-in-time copies
  - Multiple copies of a primary volume
  - Immediate read/write access to virtual copy
  - Fast restore from any virtual copy

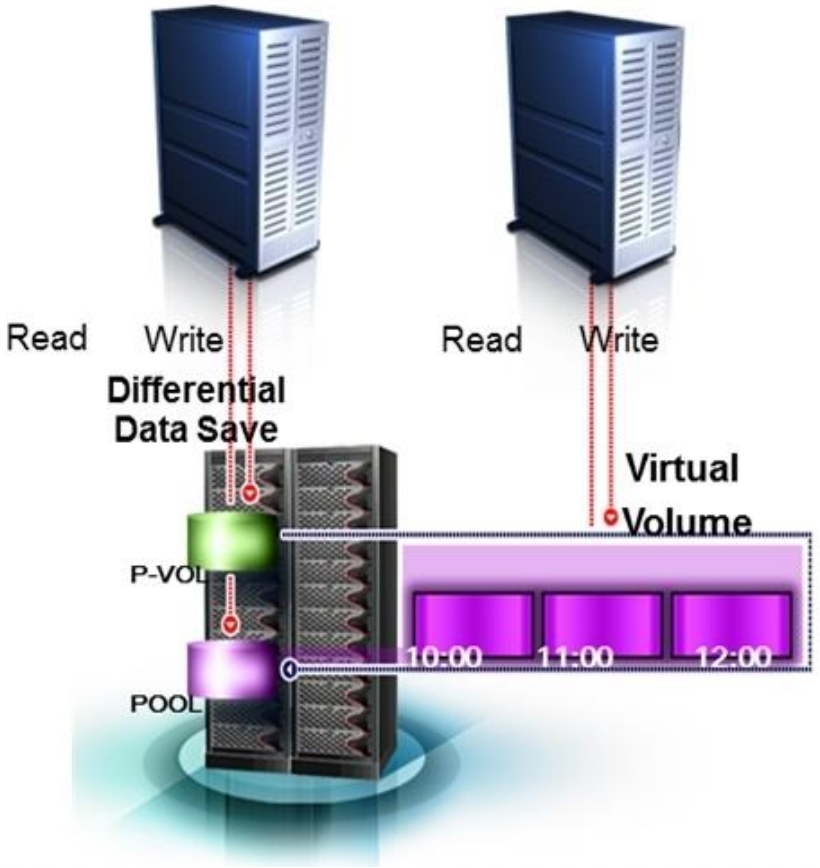

Copy on write snapshot. Notice that both P-VOL and V-VOL are accessible for I/O operations. Snapshots can be created instantly

In-System Replication: Copy-On-Write Snapshots

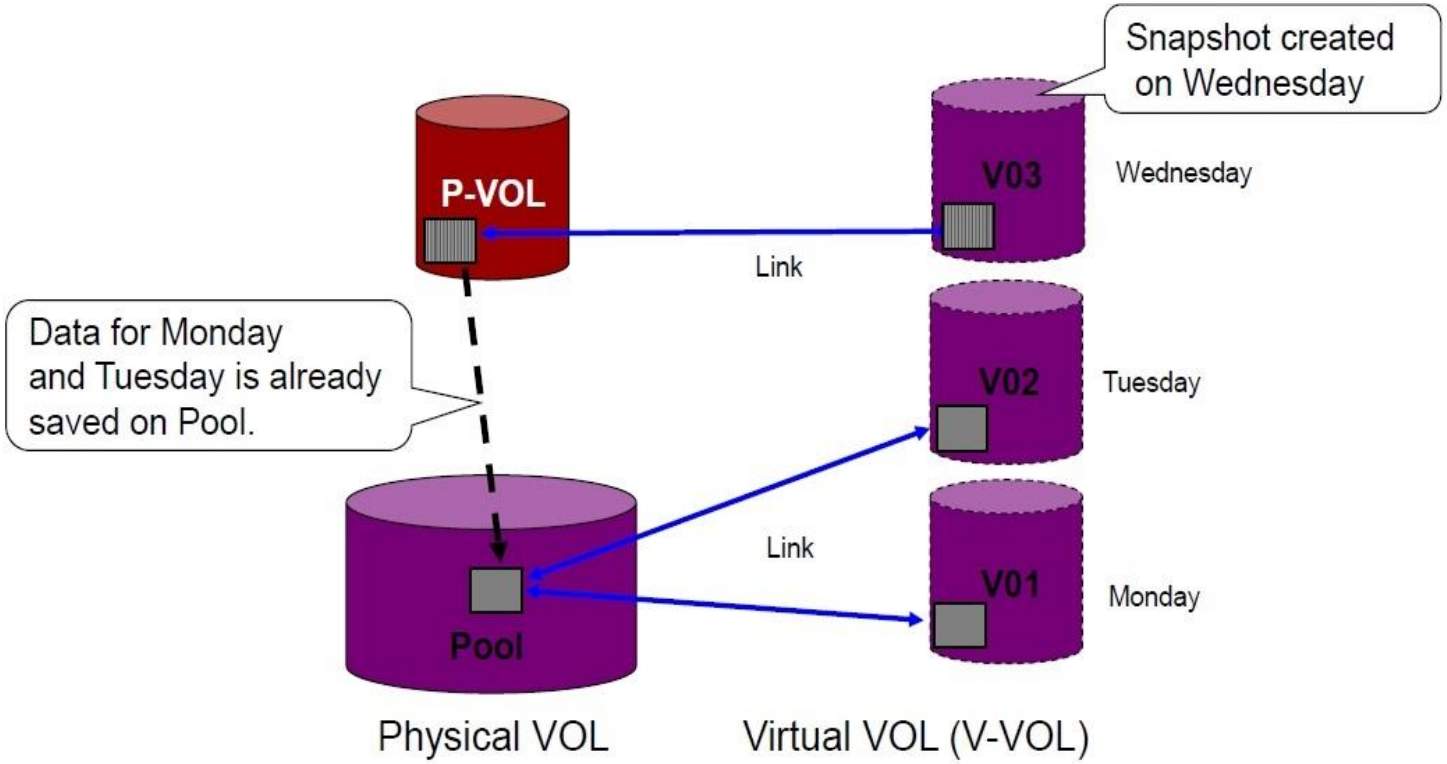

- Copy-on Write virtual volume (V-VOL) maintains a view of the primary volume (P-VOL) at a particular point in time
- V-VOL is a composite of original data in the P-VOL and change data in the pool
- V-VOL presents as a full volume copy to any secondary host
- Since V-VOL does not copy all data, it can be created or deleted almost instantly

Remote Replication

Remote Replication Scheme

Remote replication scheme; DWDM, ATM or IP connections to remote site are possible.

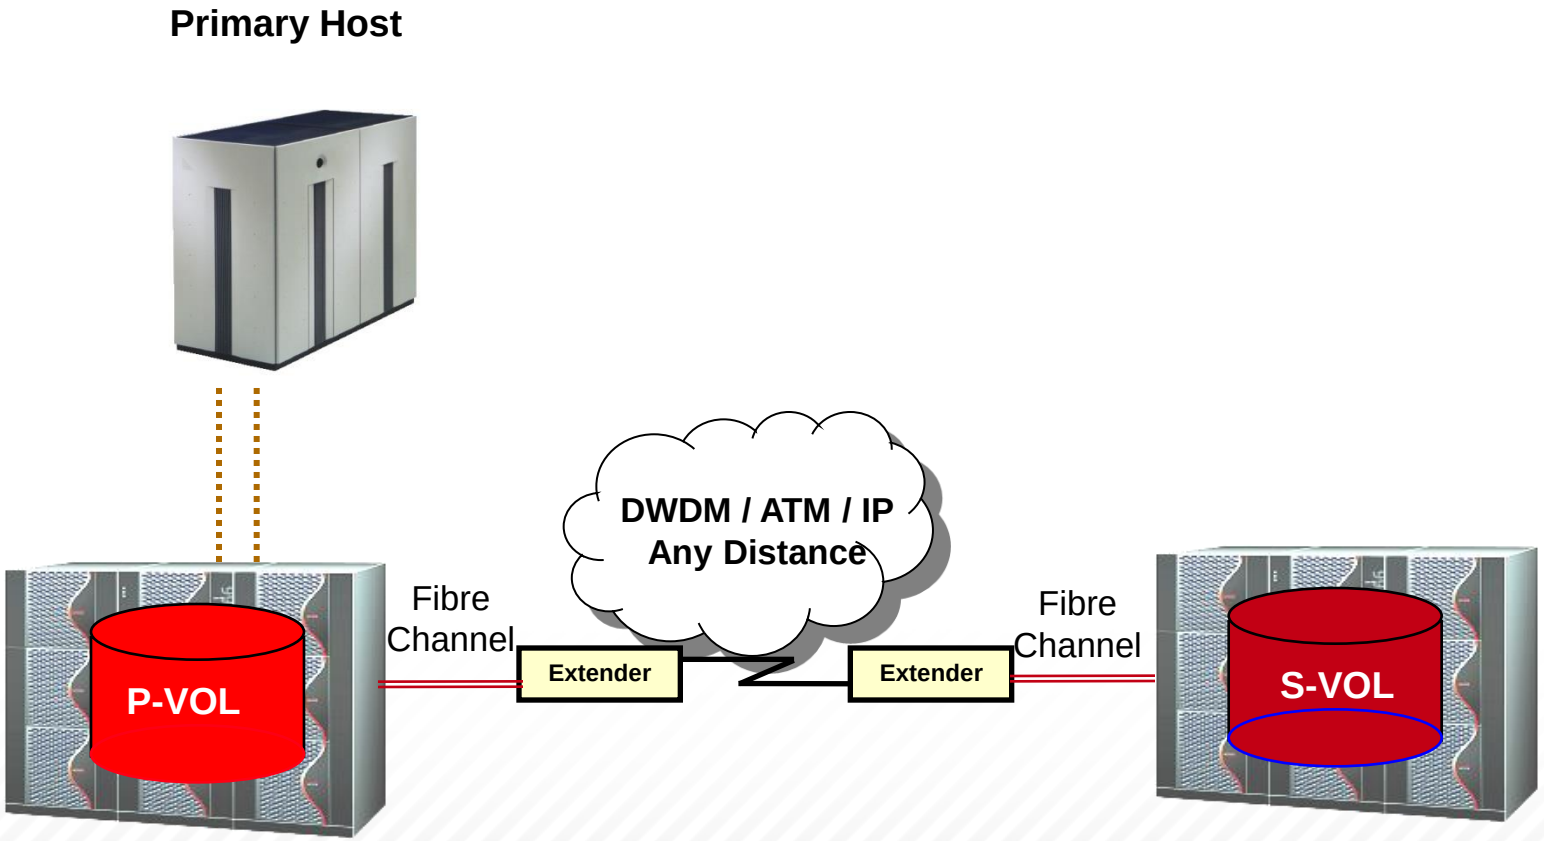

Synchronous Remote Replication

- The remote I/O is not posted “complete” to the application until it is written to a remote system
- The remote copy is always a “mirror” image
- Provides fast recovery with no data loss
- Limited distance – response-time impact

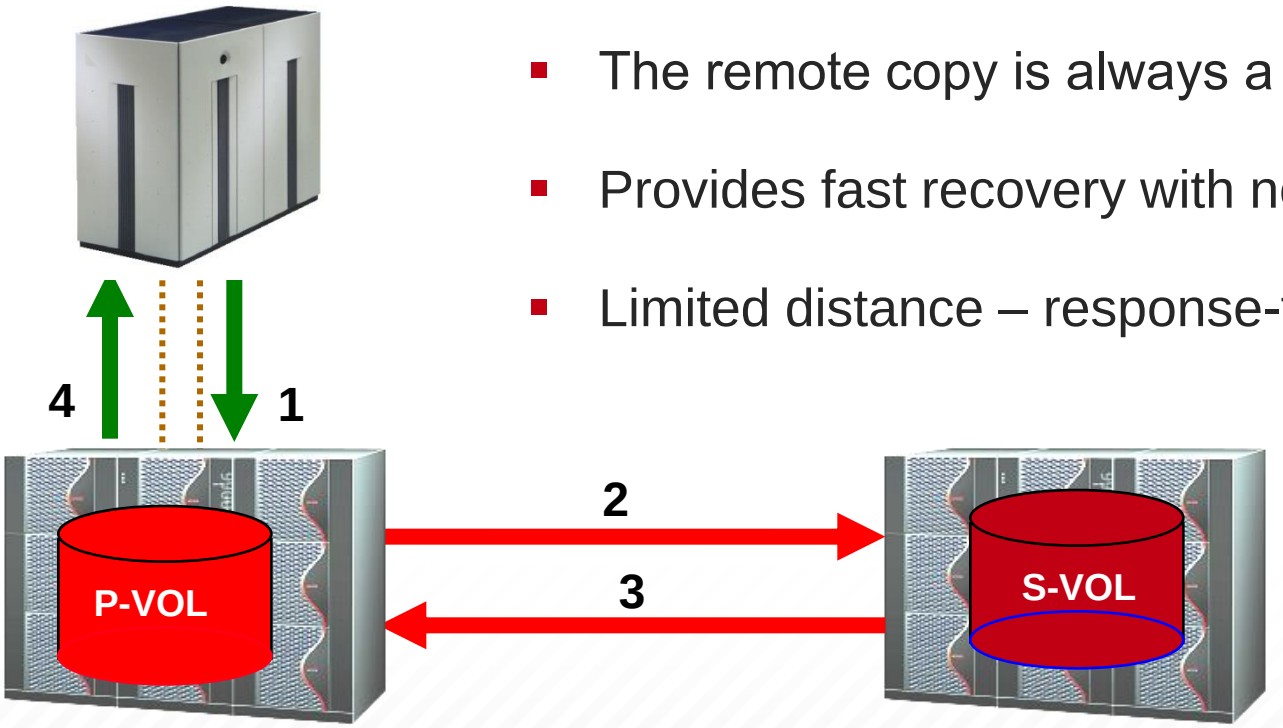

Synchronous Replication

Provides a remote “mirror” of any customer data

- The remote copy is always identical to the local copy
- Allows very fast restart/recovery with no data loss
- No dependence on host operating system, database, or file system
- Distance limit is variable, but typically less than 100 kilometers
- Impacts application response time
- Distance depends on application read/write activity, network bandwidth, response-time tolerance and other factors

Asynchronous Remote Replication

- The local I/O is disconnected from the remote I/O
- Very little impact to response time over any distance
- Data integrity and update sequence maintained over any distance
- Fast restart/recovery

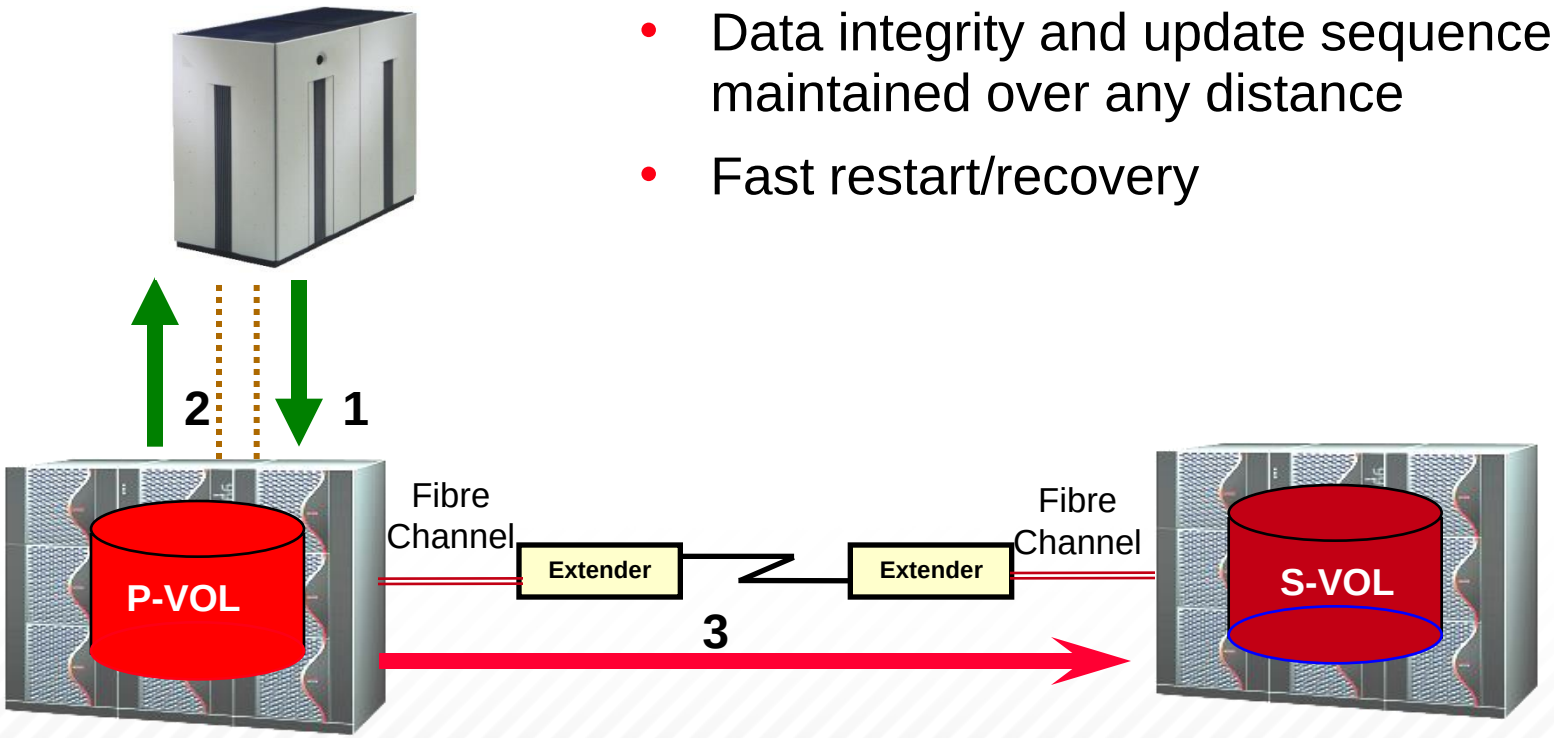

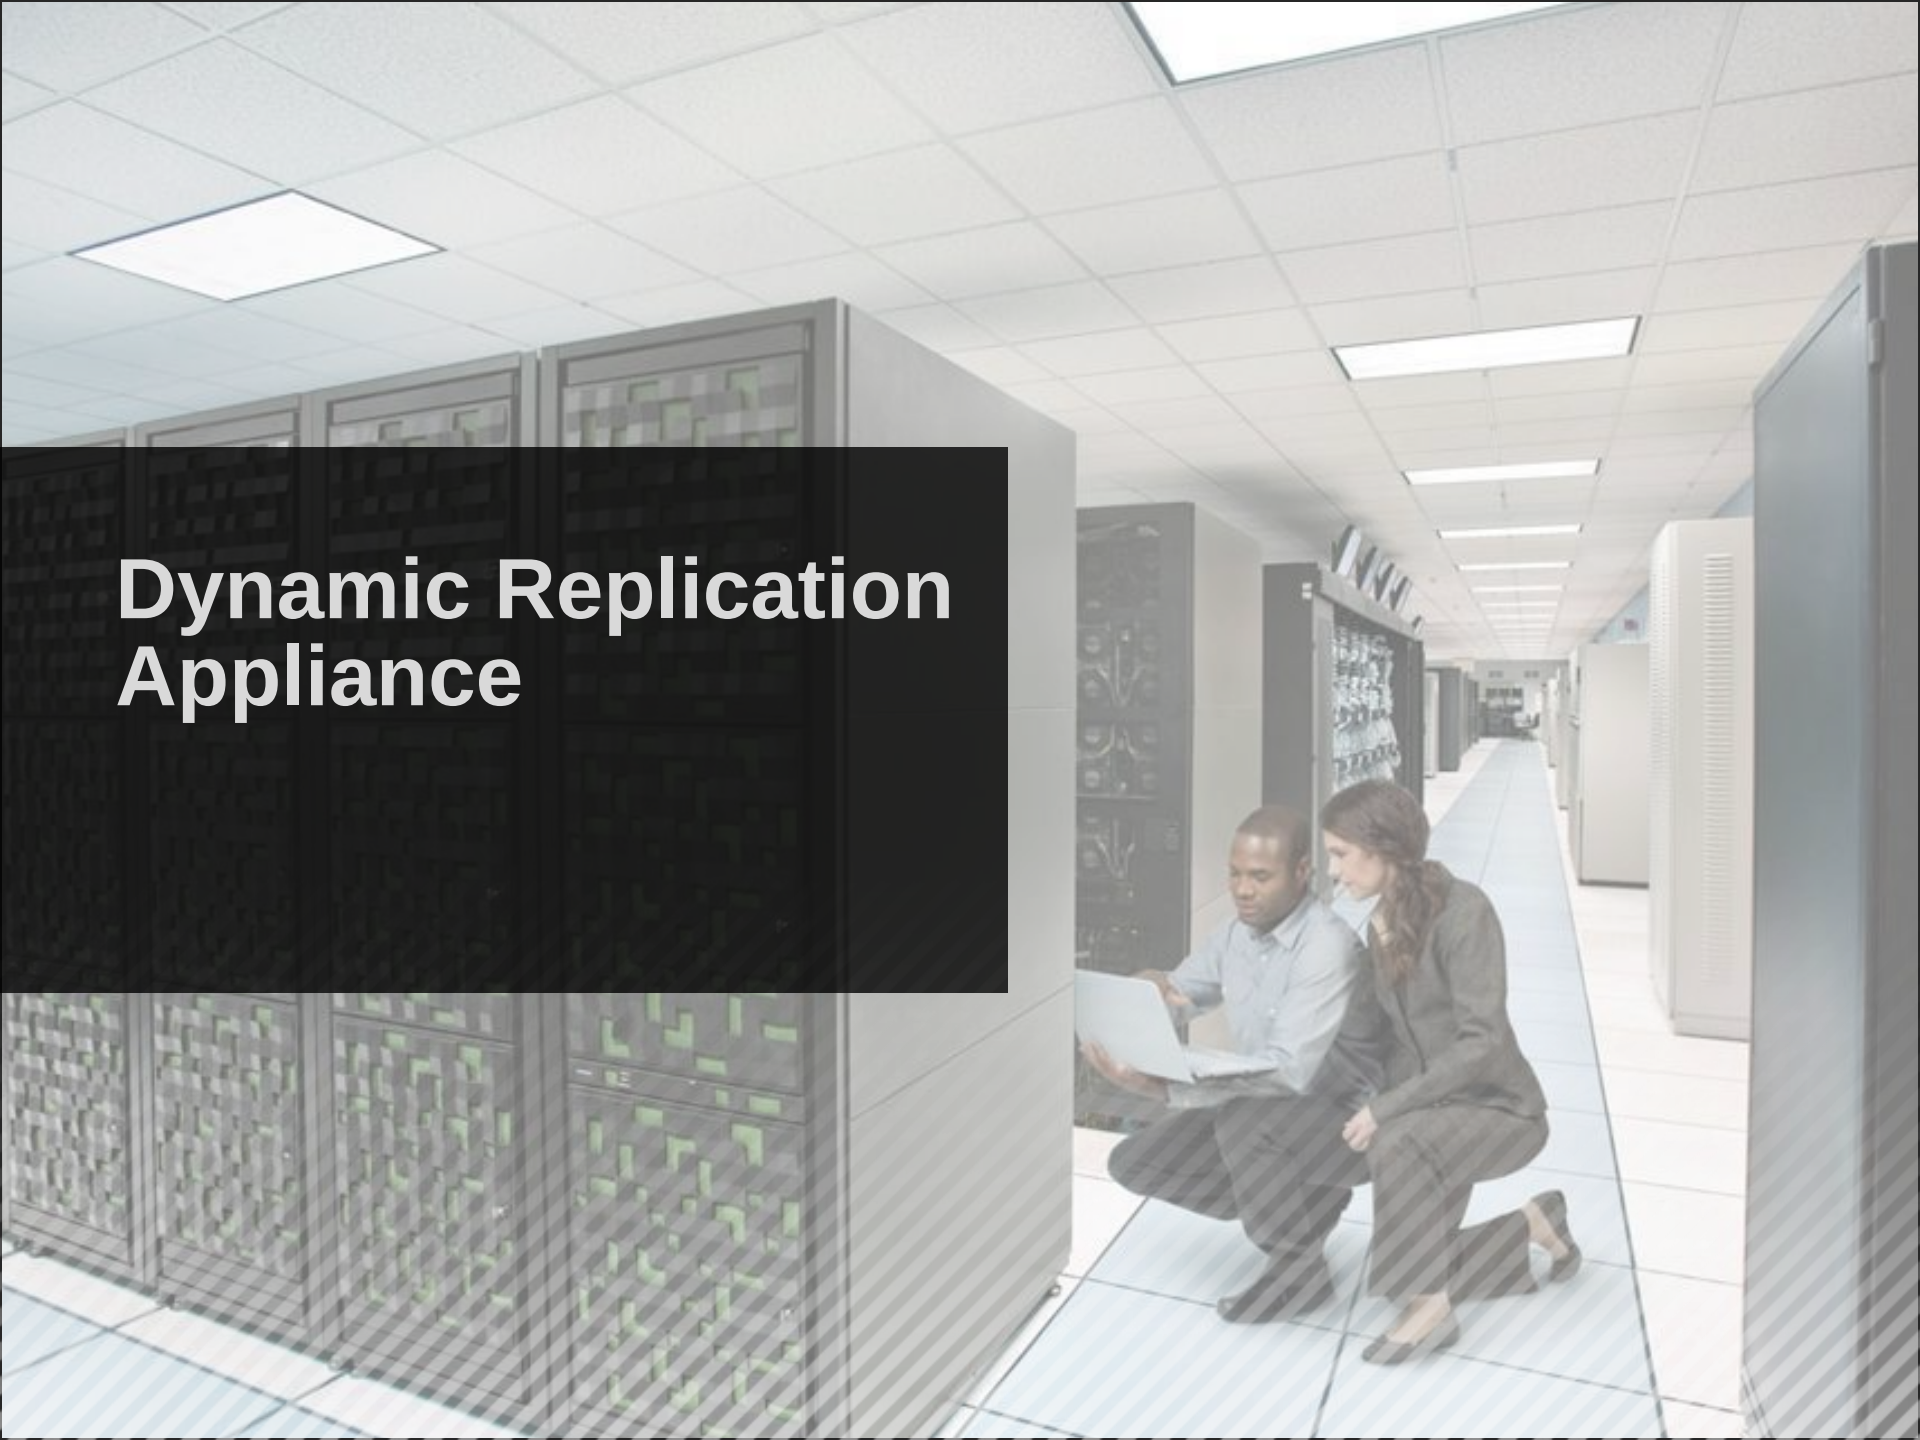

Dynamic Replication Appliance

Dynamic Replication Appliance

A possible implementation of a Replication Appliance. Data is collected from servers over LAN and then it is send to a storage system. Each server is running an agent that splits the data.

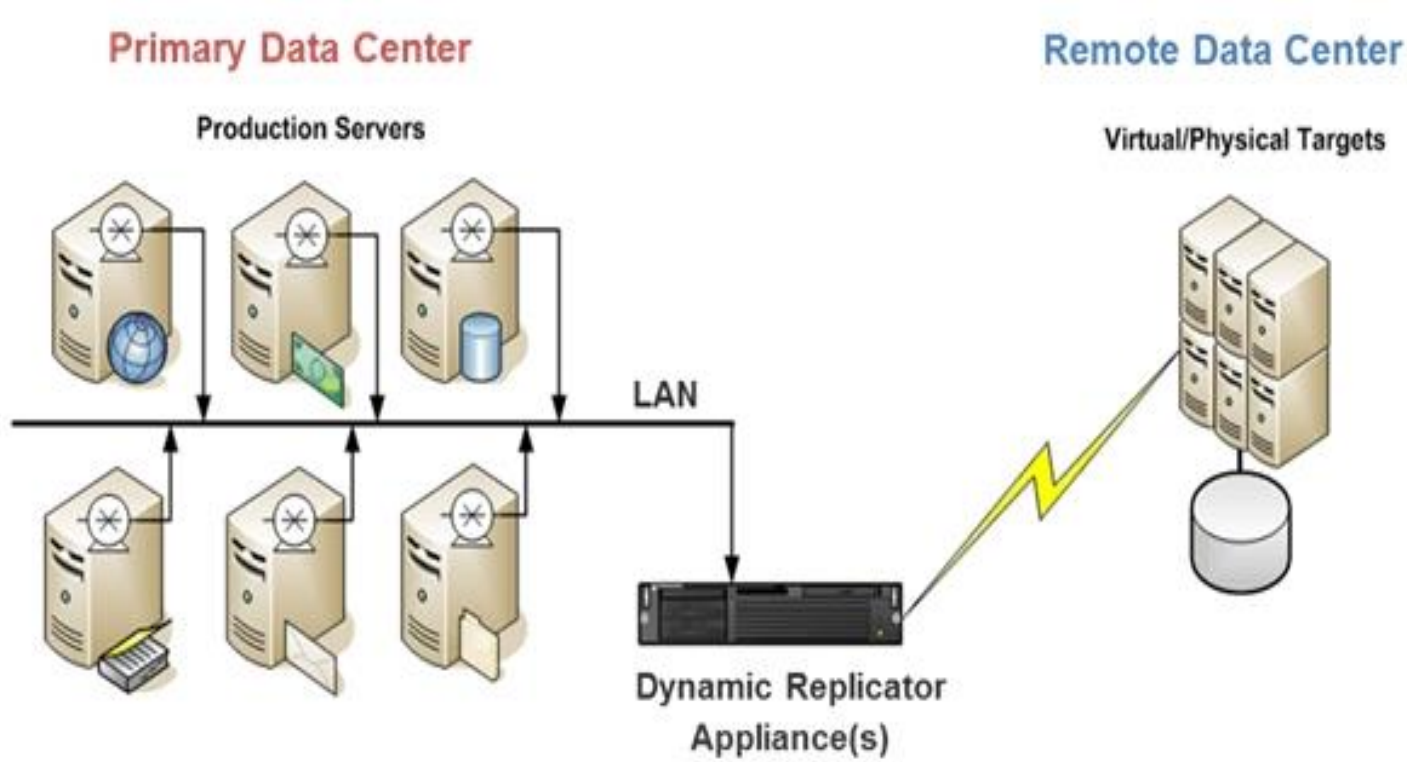

Remote Replication and Geoclusters

Geocluster interconnection scheme. Both sites (local and remote) are equipped with the same nodes. Data from the Disk Array A are synchronously replicated to Disk Array B over iFCP, FCIP or Dark Fiber technology, both SANs are interconnected. Servers in both locations are also interconnected, usually using TCP/IP protocol.

.

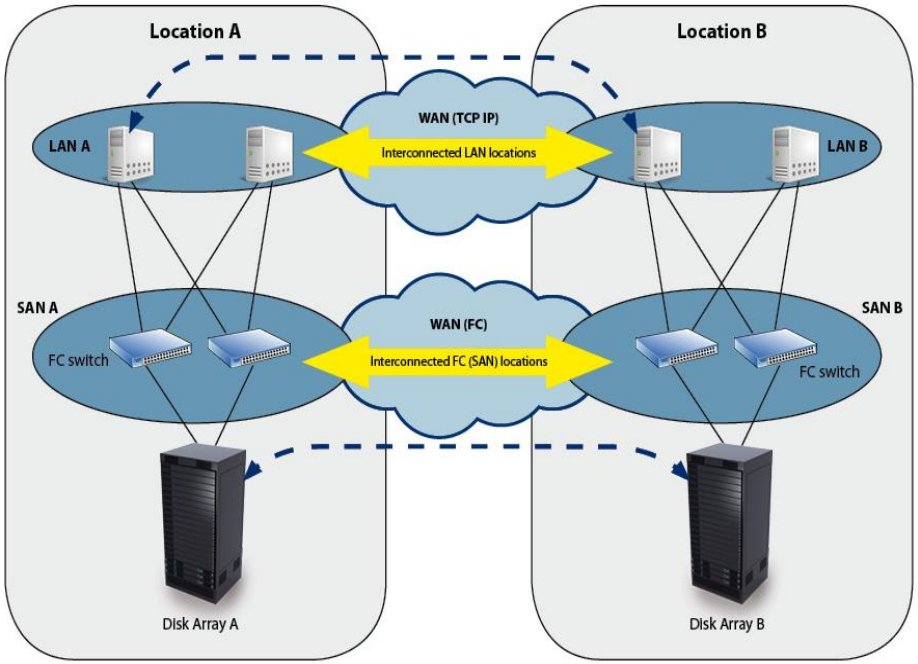

Three Data Center Multi-target Replication

Three data center multi-target replication. The maximum possible data protection is ensured by using two remote sites for data replication.

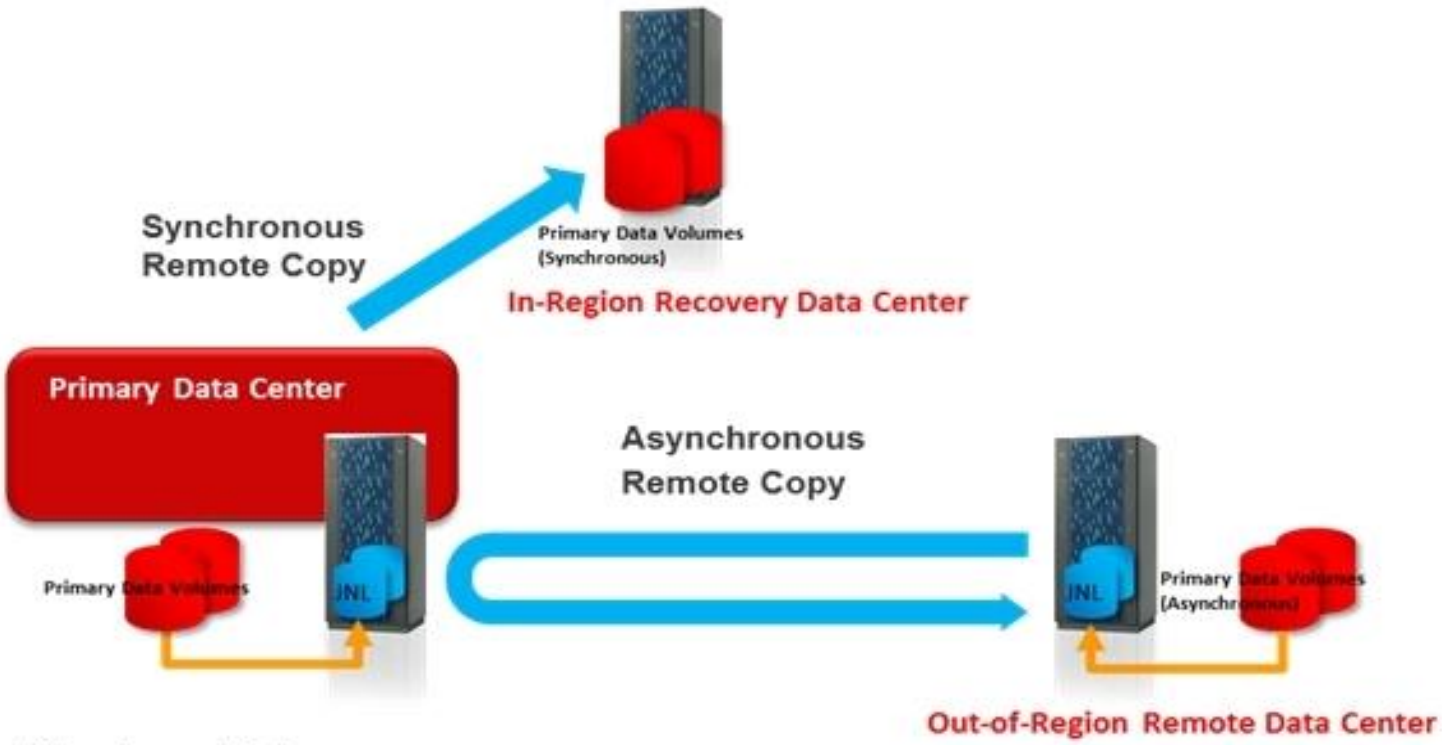

JNL = Journal Volume

Primary JNL – Data is stored before transmission.

Secondary/Remote JNL – Correct record sequence is re-established after transmission.

Diversity in Data Protection Requirements

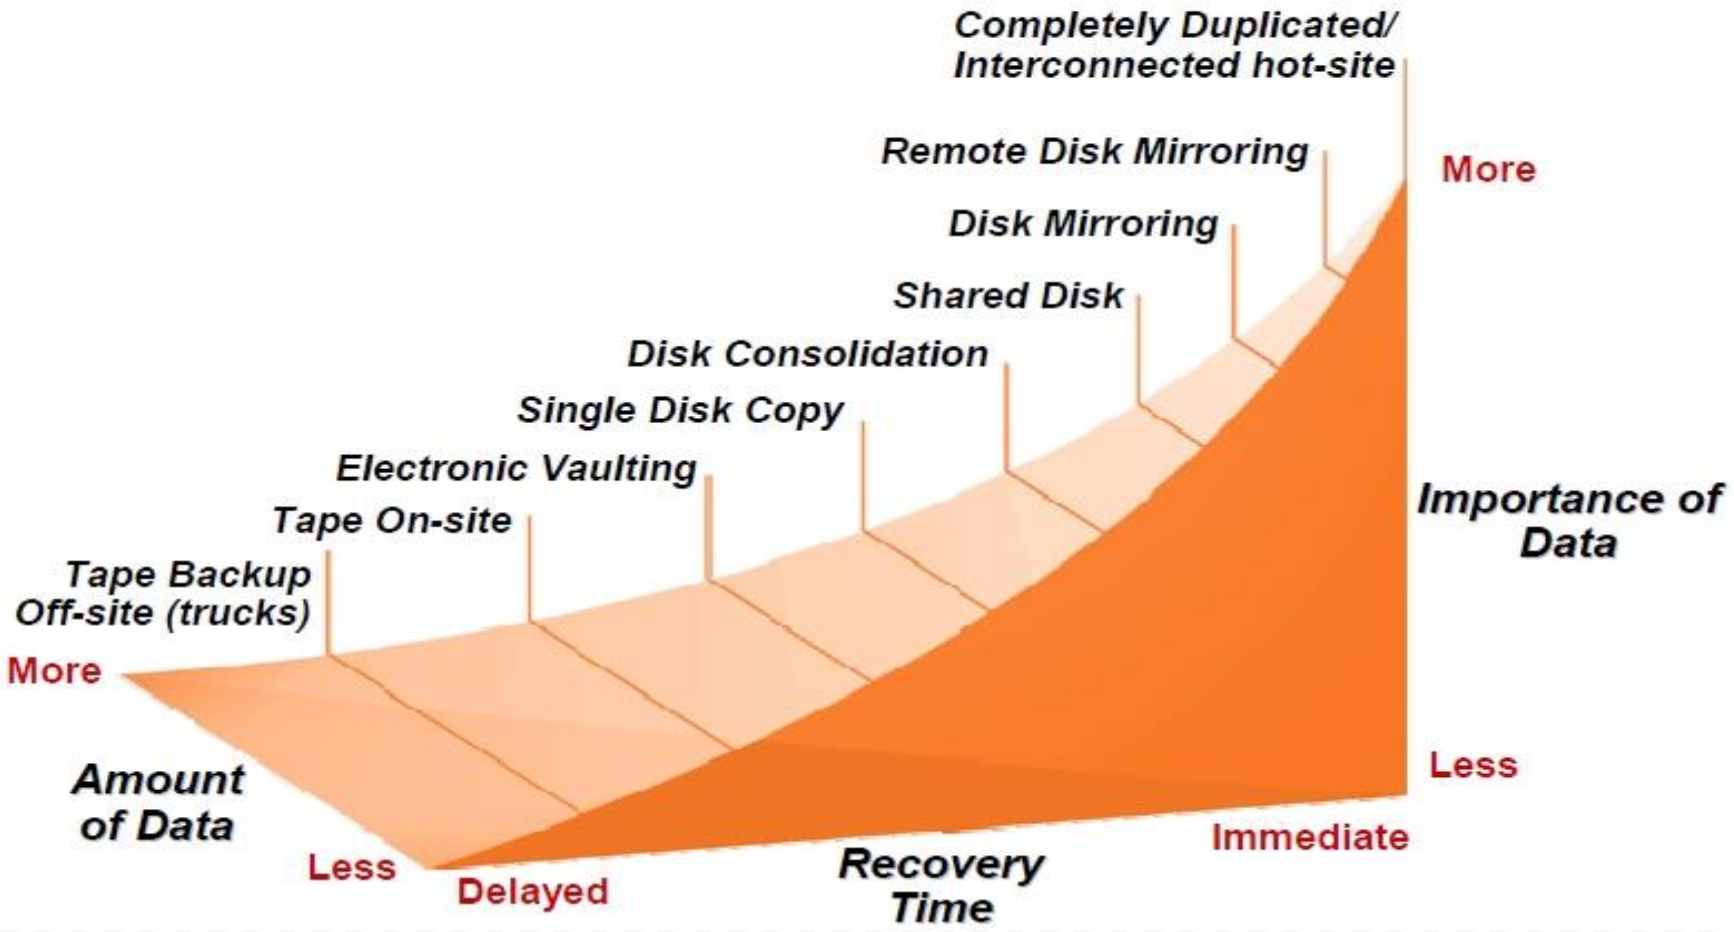

Solution Area of Cost, Performance and Distance

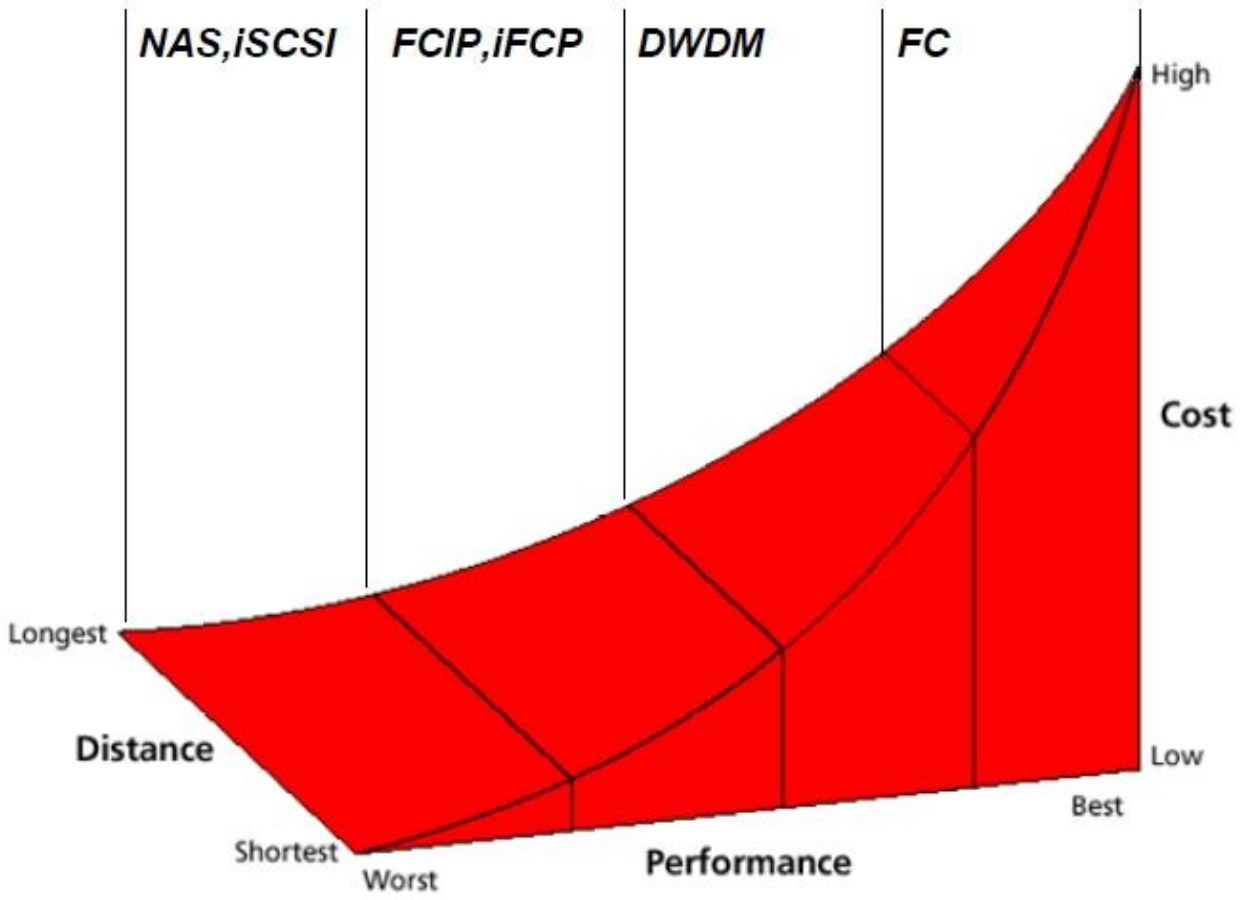

Exercise: Replication Scenario

- Scenario:
  - Financial services business with two data centers 300 miles apart
- Your task:
  - Describe a disaster recovery strategy for this business using data replication.

If a vILT class, use your drawing tools to show your configuration on the slide

Module Summary

- Upon completion of this module, you should have learned to:
  - Describe basic concepts of business continuity and replication
  - Understand business impact analysis and risk assessment
  - Describe basic concepts of disaster recovery

Storage Concepts

Virtualization of Storage Systems

Module Objectives

- Upon completion of this module, you should be able to:
  - Understand and explain virtualization concepts and its benefits
  - Explain the difference between *fat* and *thin* provisioning
  - Describe SAN virtualization concepts

Module Topics

- Virtualization concepts, features, and benefits
- Different types of virtualization
- “Fat” and “Thin” provisioning concepts and features

Virtualization Concepts

What Is Virtualization?

- Definition (*source: www.wikipedia.org*)
  - **Virtualization** is the abstraction of computer resources.
  - Hides the physical characteristics of computing resources from the way in which other systems, applications, or end users interact with those resources
  - Single physical resources appear to function as multiple logical resources, or multiple physical resources appear as a single logical resource.
- Virtualization has moved “out of the box” and into the infrastructure, or *cloud*, and virtualization solutions are available at these layers:
  - Virtualization of applications
  - Virtualization of computers
  - Virtualization of networks
  - Virtualization of storage

Server Virtualization

The traditional architecture model requires one physical server per operating system and application. A virtualized server is able to run several virtual machines that all share the physical hardware.

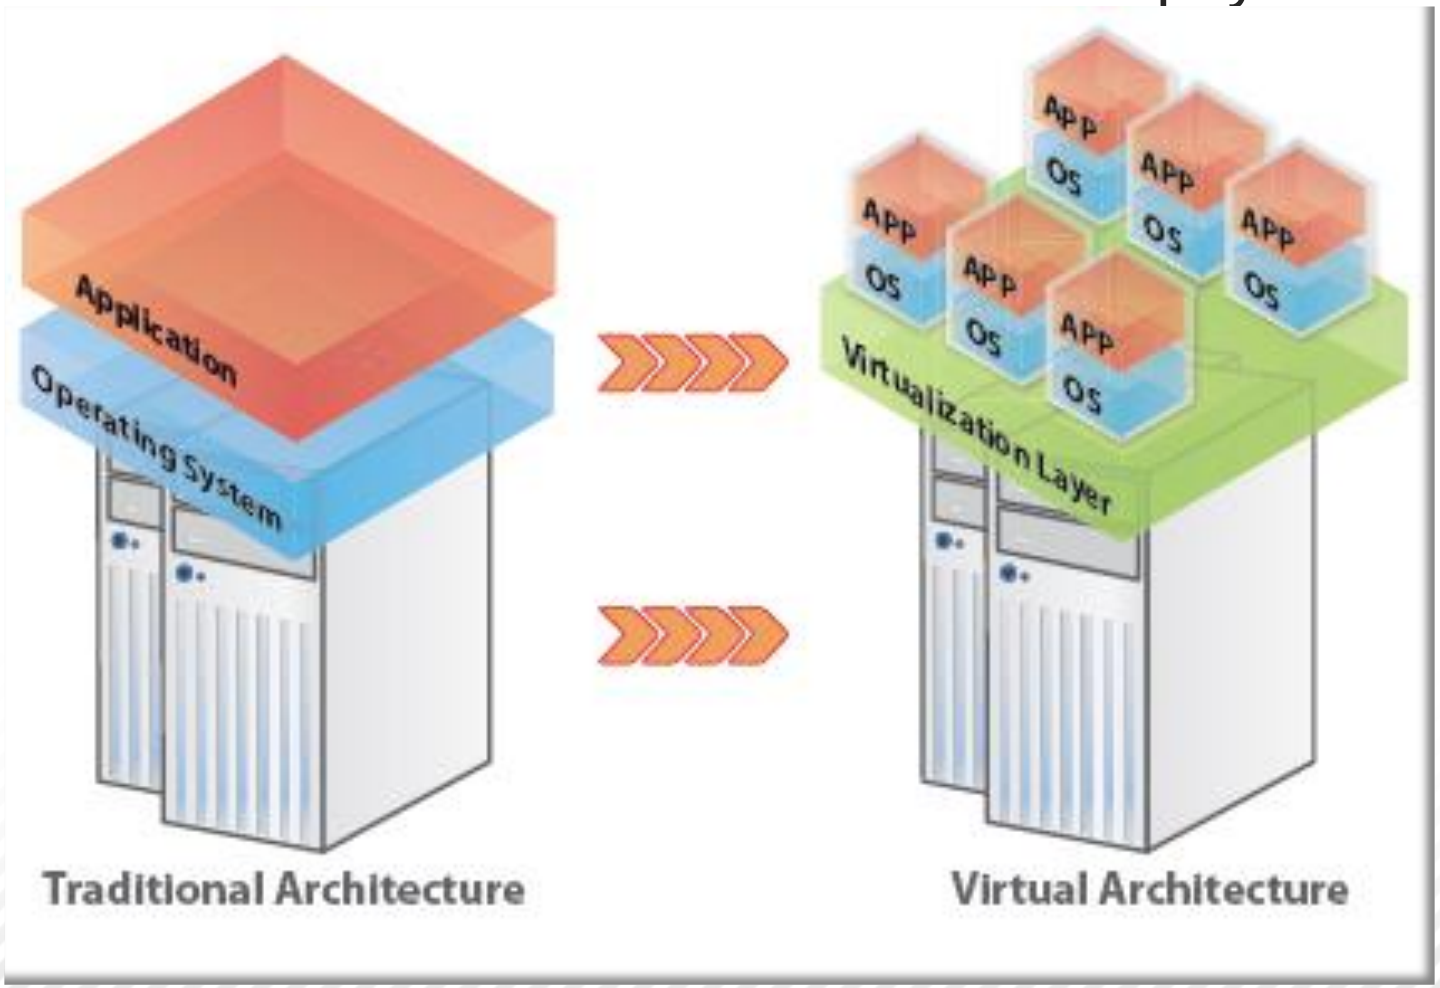

Virtualization Elements

Elements of Virtualization

- Virtualization Areas

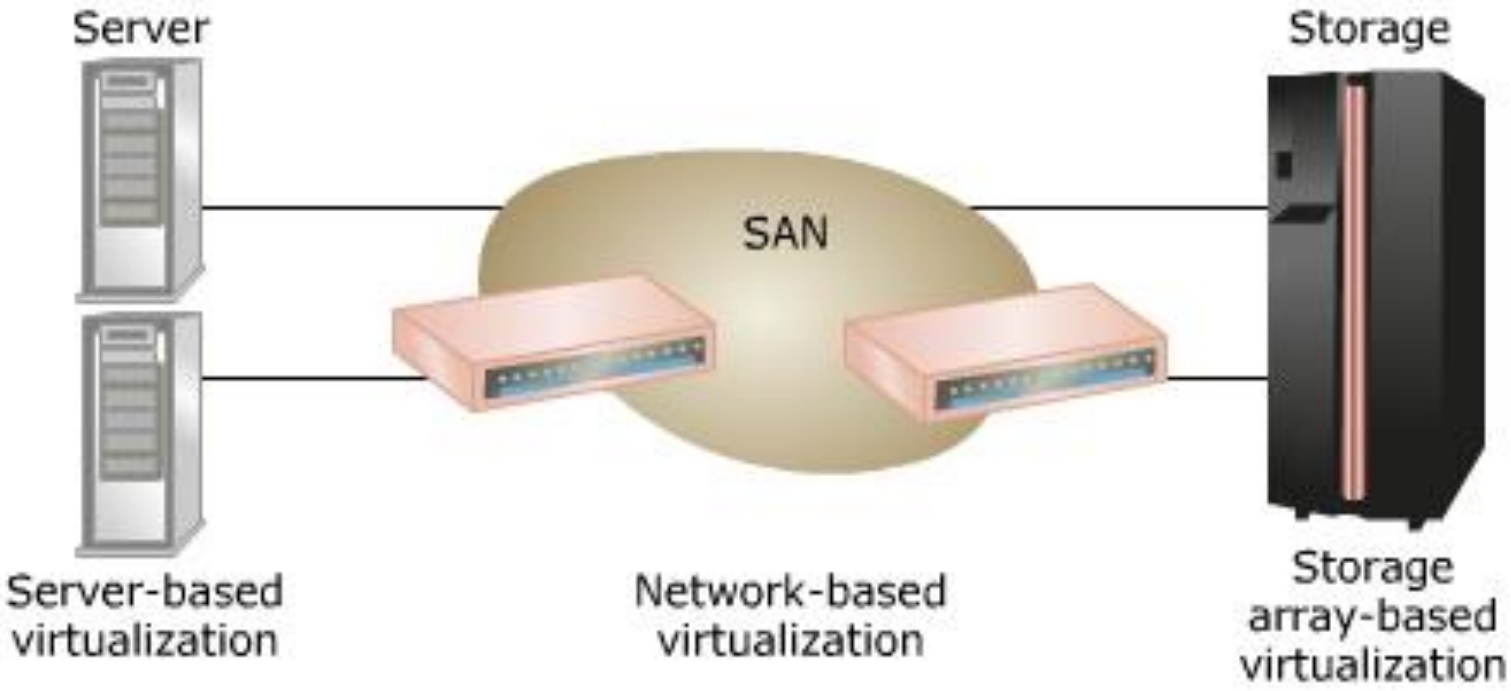

Server Virtualization

- Server virtualization is increasingly used; it provides better utilization of server resources
- One-to-many virtualization – Makes one physical server look like many servers; allows multiple operating systems on one physical server

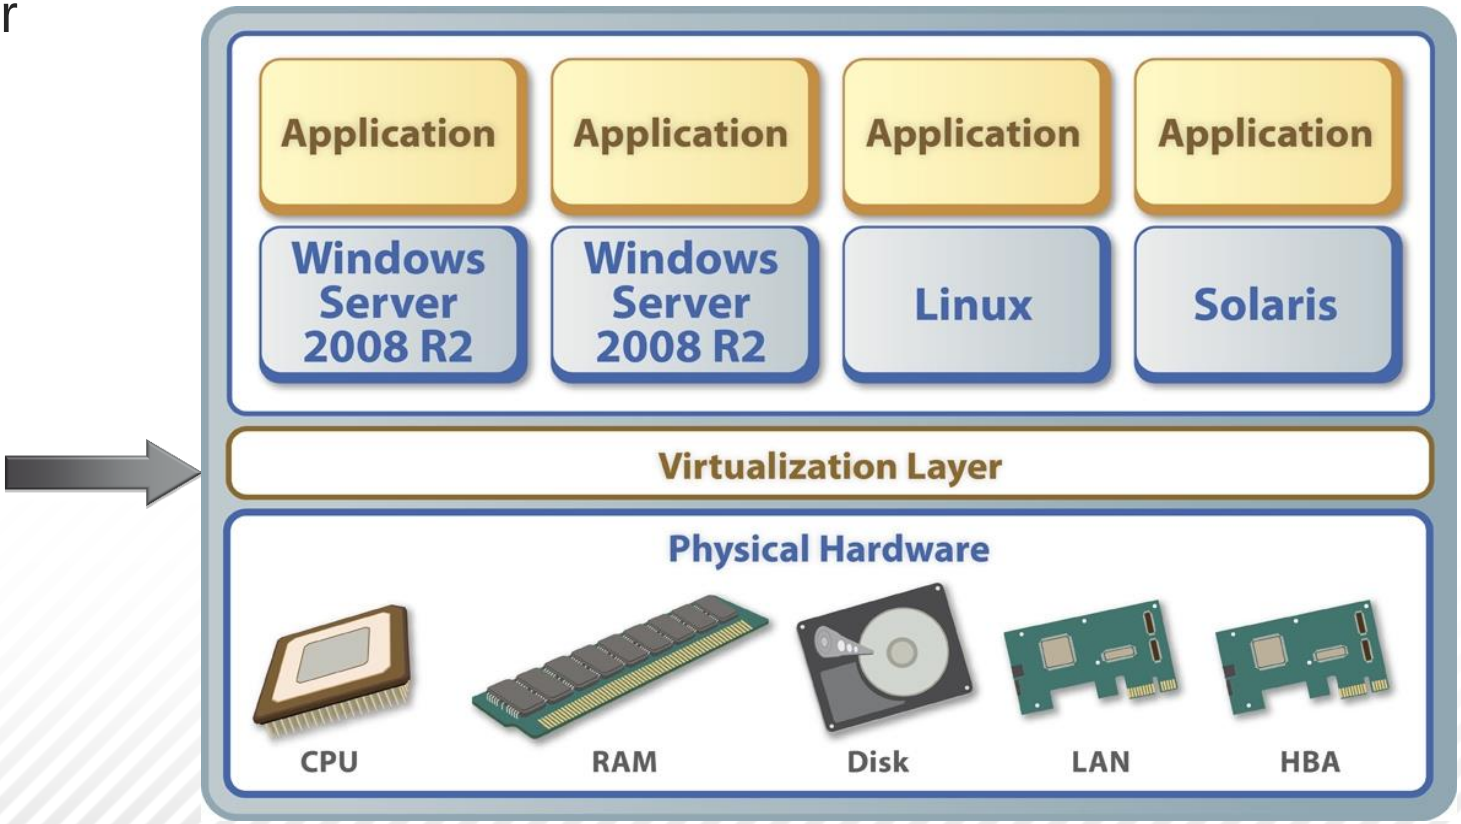

Layers of Virtualization

- Users can access the application from s a virtual desktop.
- Applications can run on several virtual machines
- The server can be virtualized
- The host can access a virtualized volume

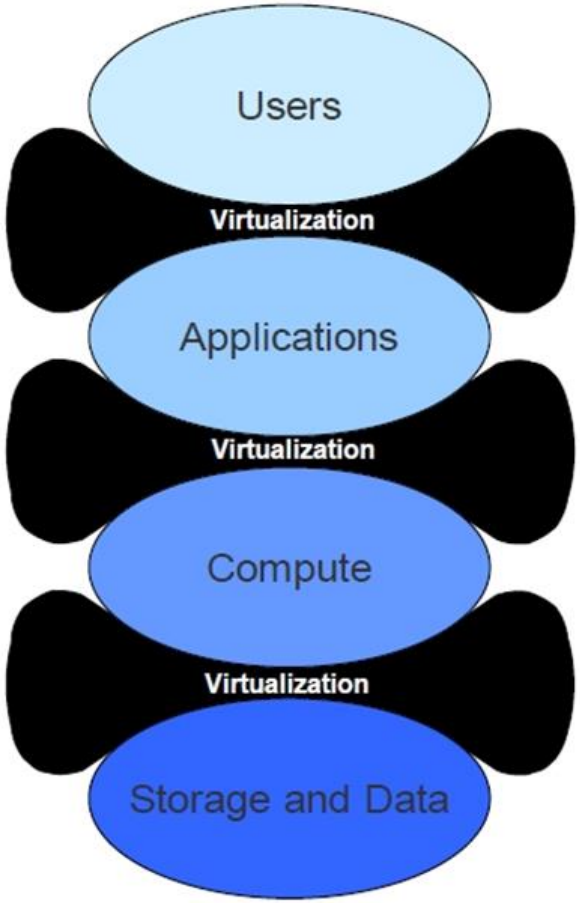

VMware® Based Server Virtualization

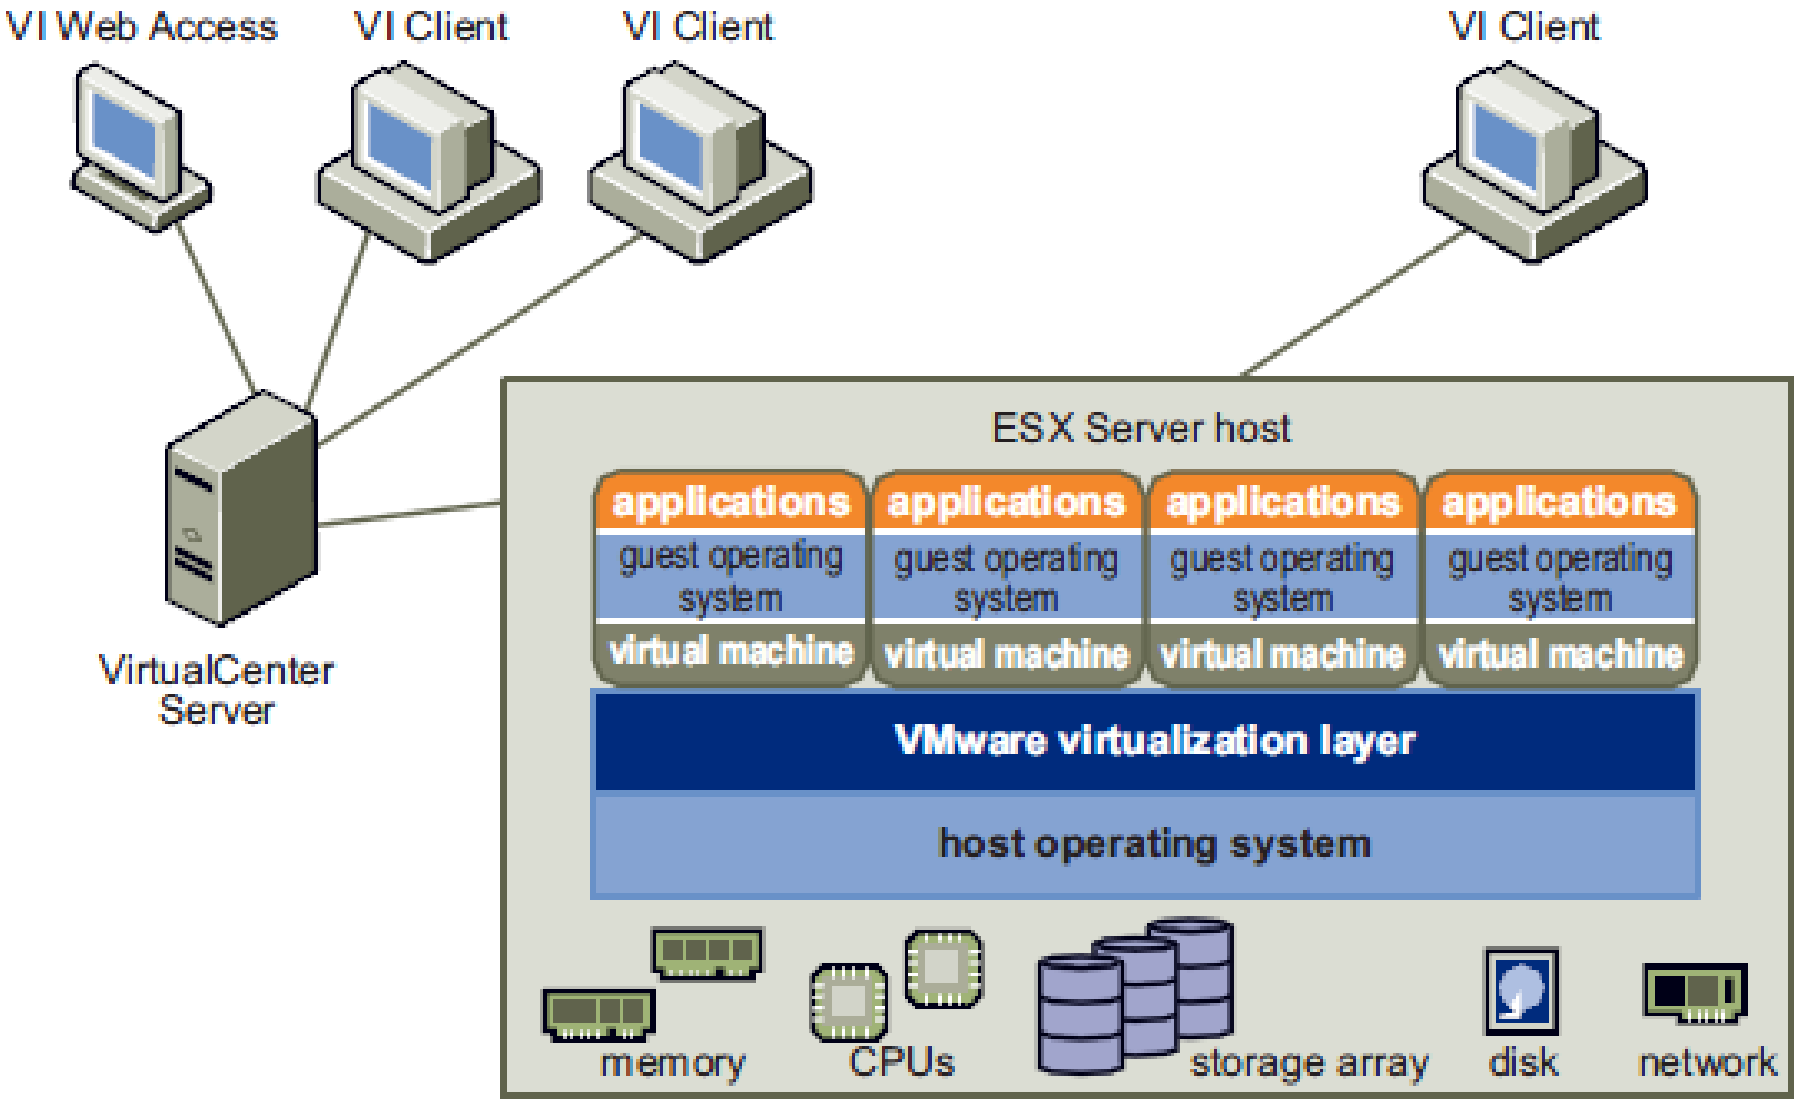

Hyper-V™ Based Server Virtualization

The hypervisor virtualization layer is thin and optimized for direct access to hardware resources. Virtual machines in child partitions access the virtualization layer through the VMBus interface. Device drivers for virtual machines are loaded from the parent partition with the original instance of Windows 2008 Server. Virtual machine configuration and management are also done in the parent partition operating system.

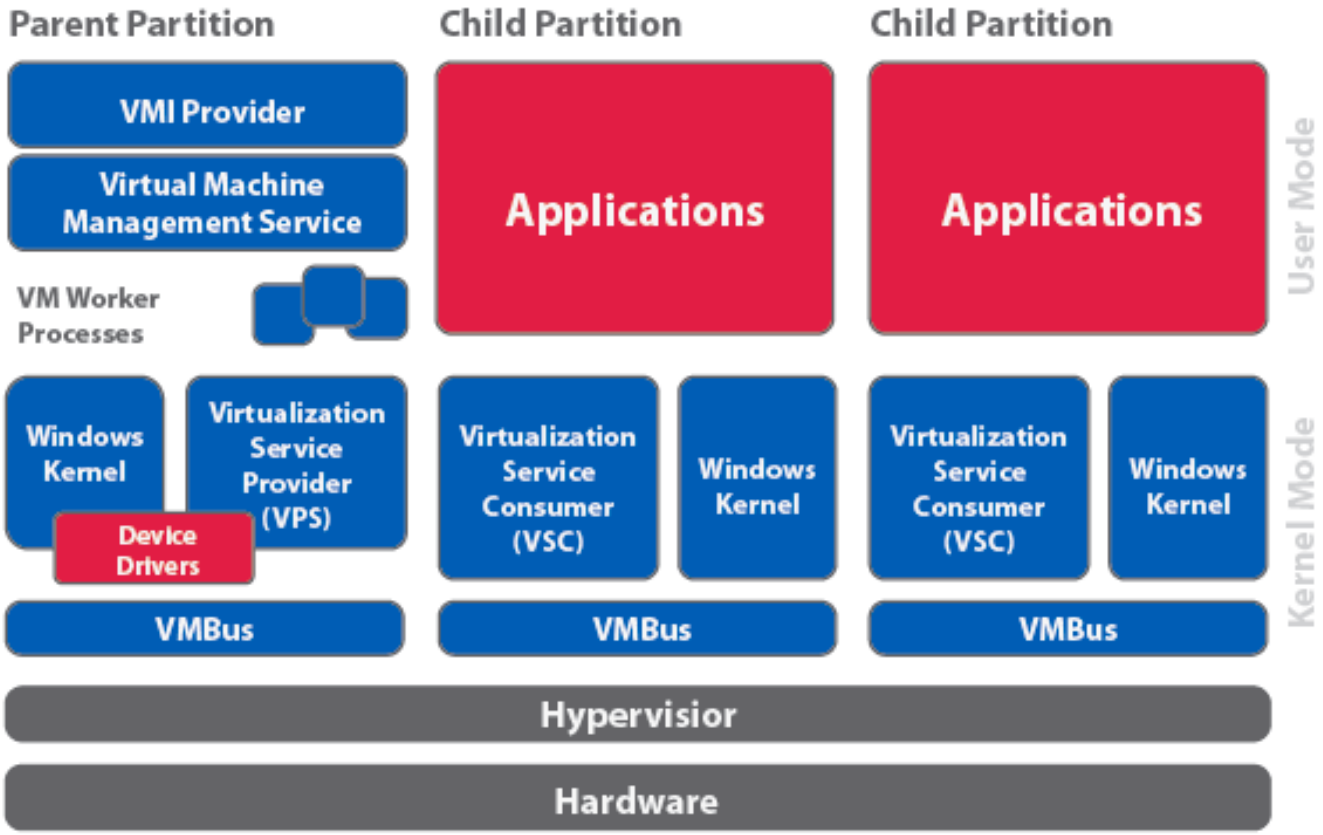

Virtualization with Blade Servers

Blade servers are installed in a blade server chassis. The chassis is then placed in a standard rack. These blade servers offer logical partitioning, which is a highly sophisticated form of server virtualization.

Storage System Virtualization

Every storage system offers RAID functionality but additionally may offer other virtualization capabilities such as:

- Cache Partitions
- Virtual Ports
- Storage virtualization (of other storage systems)
- Thin provisioning and automatic tiering

Virtualization Benefits

Migration – VMs and LUNs can be easily transferred from one physical device to another.

Backup – encapsulation simplifies backup of the VM

Hardware platform independence – physical servers can be of different configurations

Enhanced utilization – allows effective use of resources

Lower power consumption

Lower RPO and RTO

Physical resources can be added without disruption

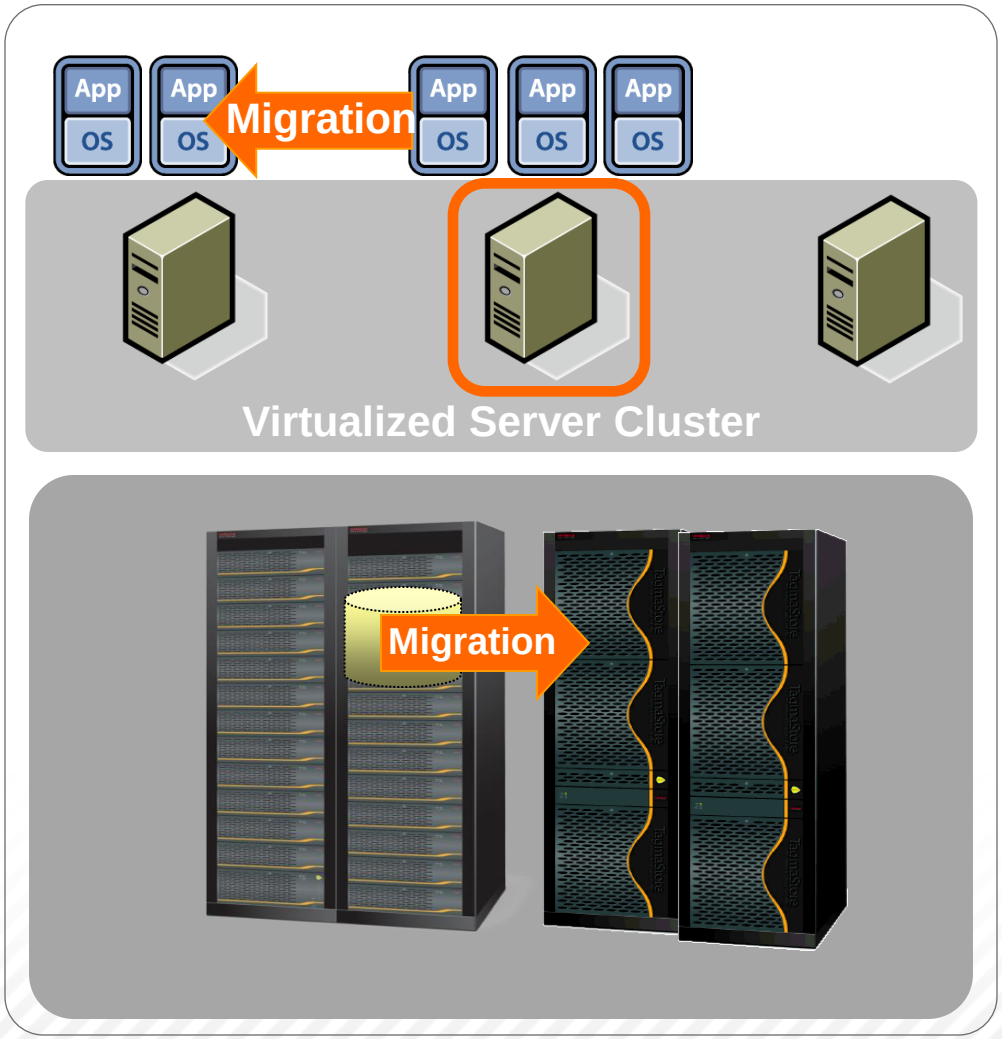

Virtualization – Thin Provisioning and Automated Tiering

Comparison of Fat and Thin Provisioning

- To avoid future service interruptions, today it is common to over-allocate storage by approximately 50% - 75%.

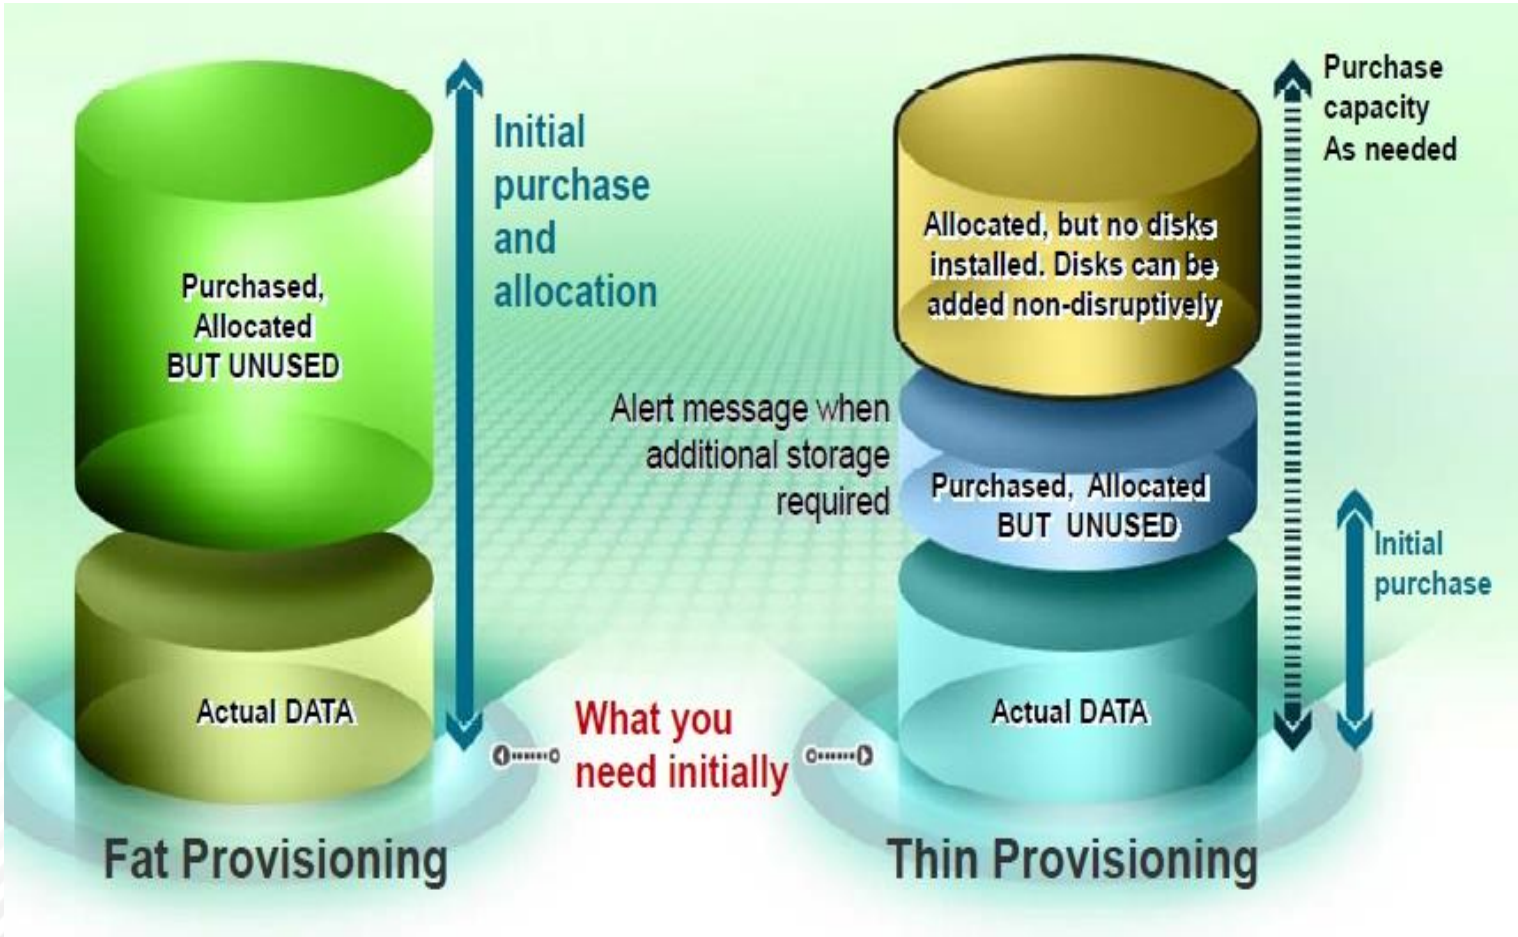

Thin Provisioning

Parity groups are added to a thin provisioning pool. Virtual volumes are mapped to servers. Virtual volumes do not contain any actual data. Data is stored in the storage pool. Virtual volumes contain pointers that point to the location of data in the pool.

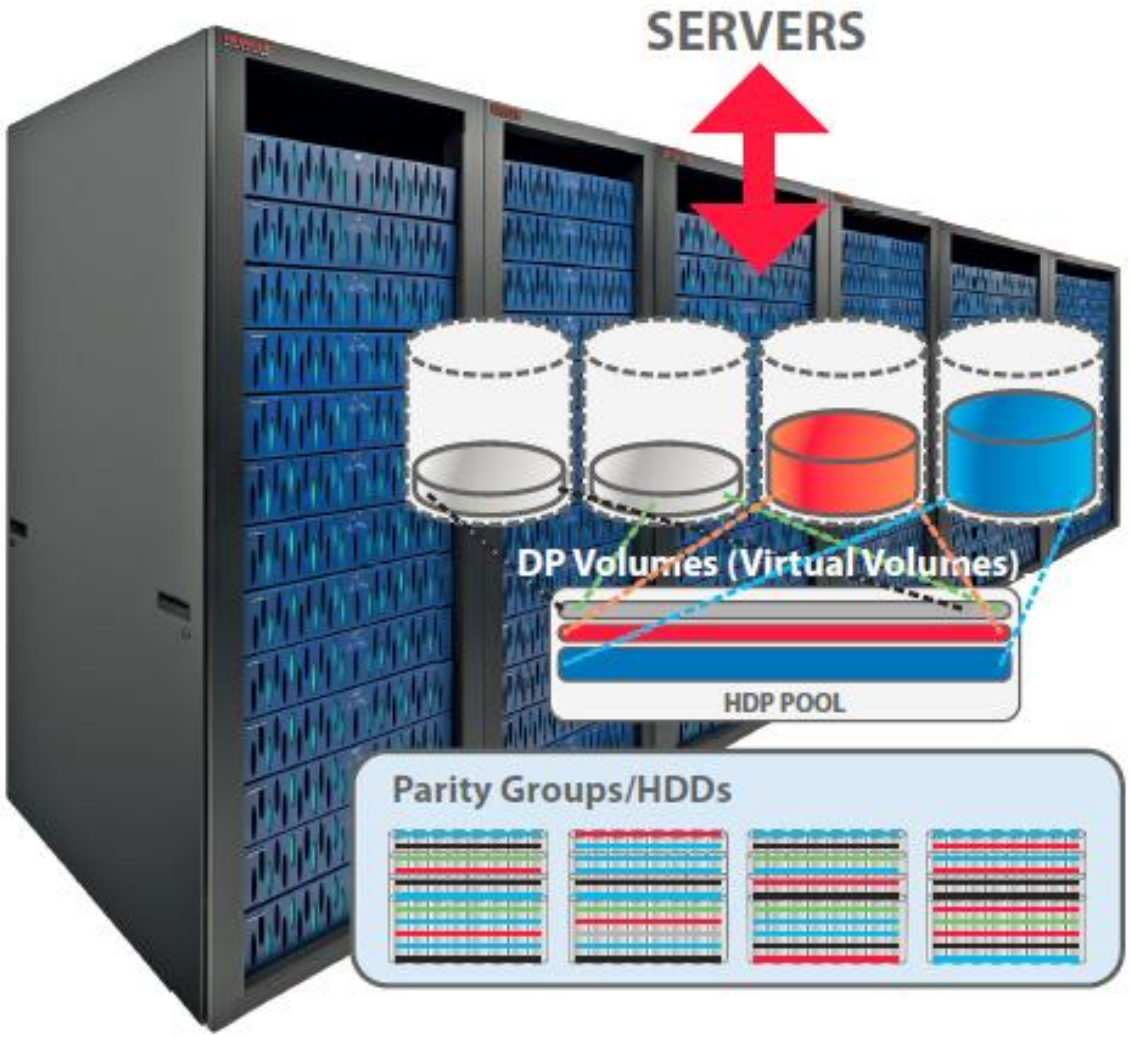

Thin Provisioning Benefits

- Increased physical disk striping for better performance
- Reduces the need for performance expertise
- Simplifies storage capacity planning and administration
- Increases storage utilization
- Eliminates downtime for application storage capacity expansion
- Improves application uptime and SLAs

Zero Page Reclaim

The storage system scans the storage pool for used data blocks that contain only zeros. These blocks are then erased and freed automatically.

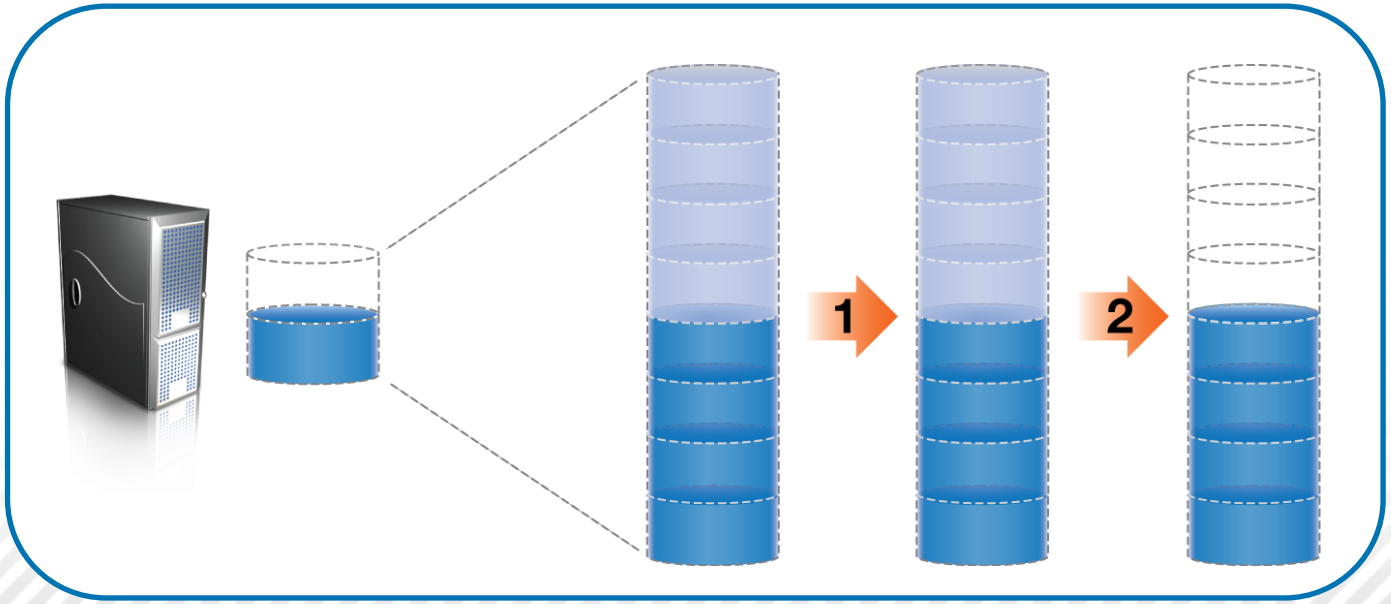

Thin Provisioning

An example of how thin provisioning can help you save the cost of buying all the capacity in advance.

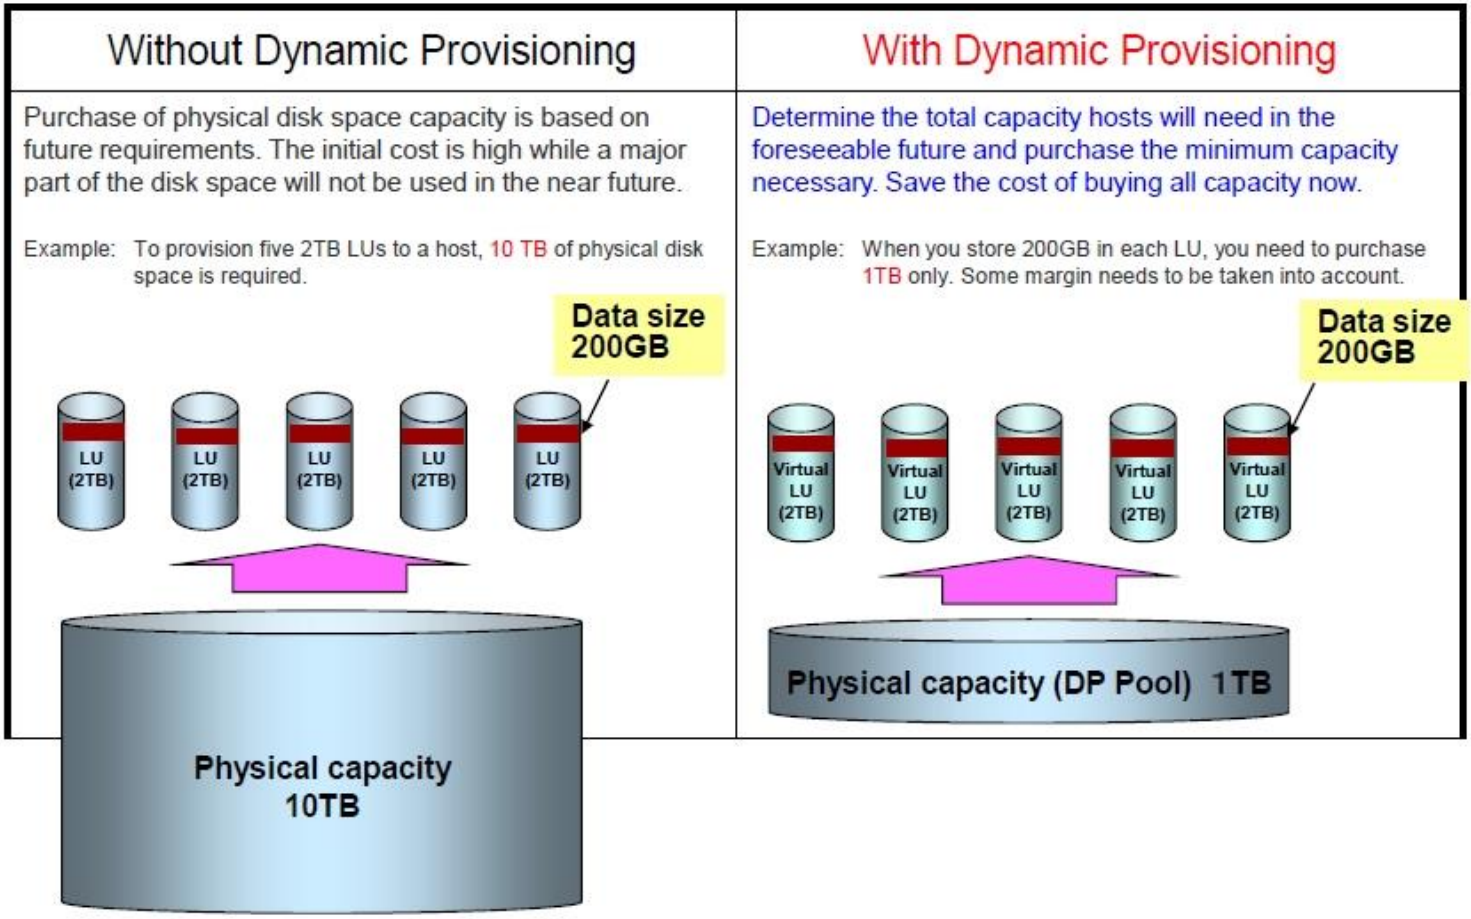

Exercise: Virtualization

1. Definition of Virtualization
   - **Virtualization** is the _____ of computer resources.
   - Hides the _____ of computing resources from the way in which other systems, applications, or end users interact with those resources.
   - Single physical resources appear to function as _____ .
   - Or multiple physical resources appear as _____.

2. Identify some of the common objects that can be virtualized:
   - _____
   - _____

If a vILT class, write your answers on blank lines

Virtualization in Storage System Controller

Virtualization in Storage System Controller

Controller based virtualization of external storage. Hitachi Virtual Storage Platform (VSP) is an example of an enterprise-level storage system that supports virtualization of external storage.

Before virtualization the data center consists of various heterogeneous storage systems in a SAN.

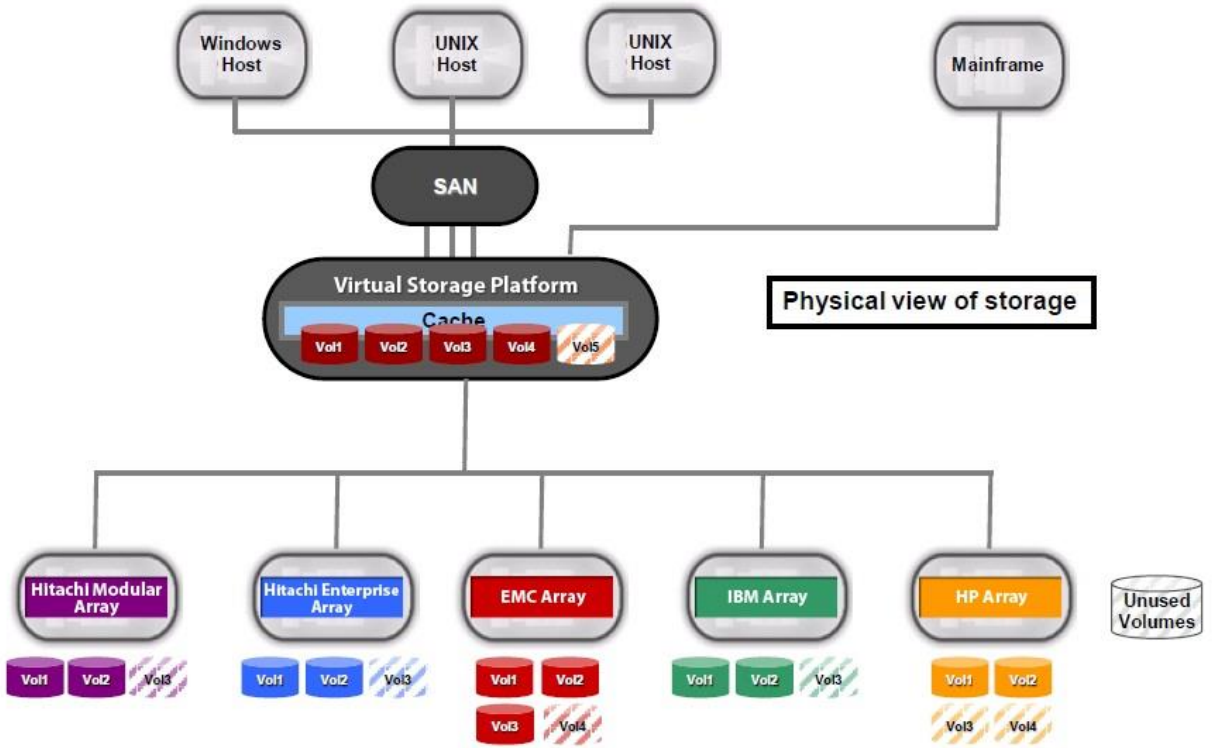

Virtualization in Storage System Controller

After virtualization the VSP storage system provides access to virtualized volumes from the other storage systems

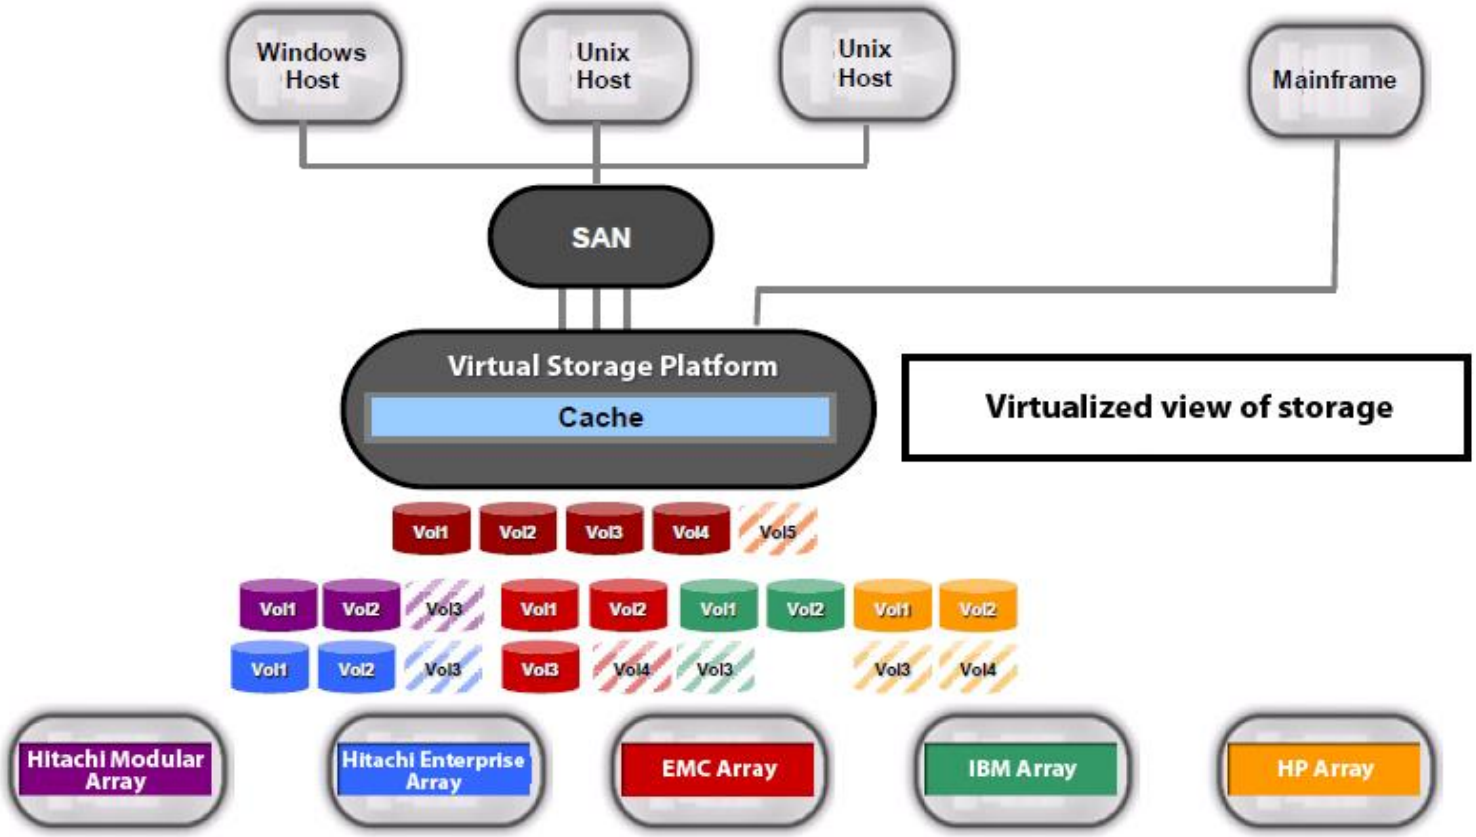

Virtualization – Logical Partitioning

Storage Virtualization – Logical Partitioning

In this figure we see one storage system (VSP) with three external storage systems that create a virtualized storage pool. The VSP is then virtualized to provide two logical partitions — private virtual storage machines. Hosts are then able to access and use only the resources (cache, ports and disks) assigned to the respective partition.

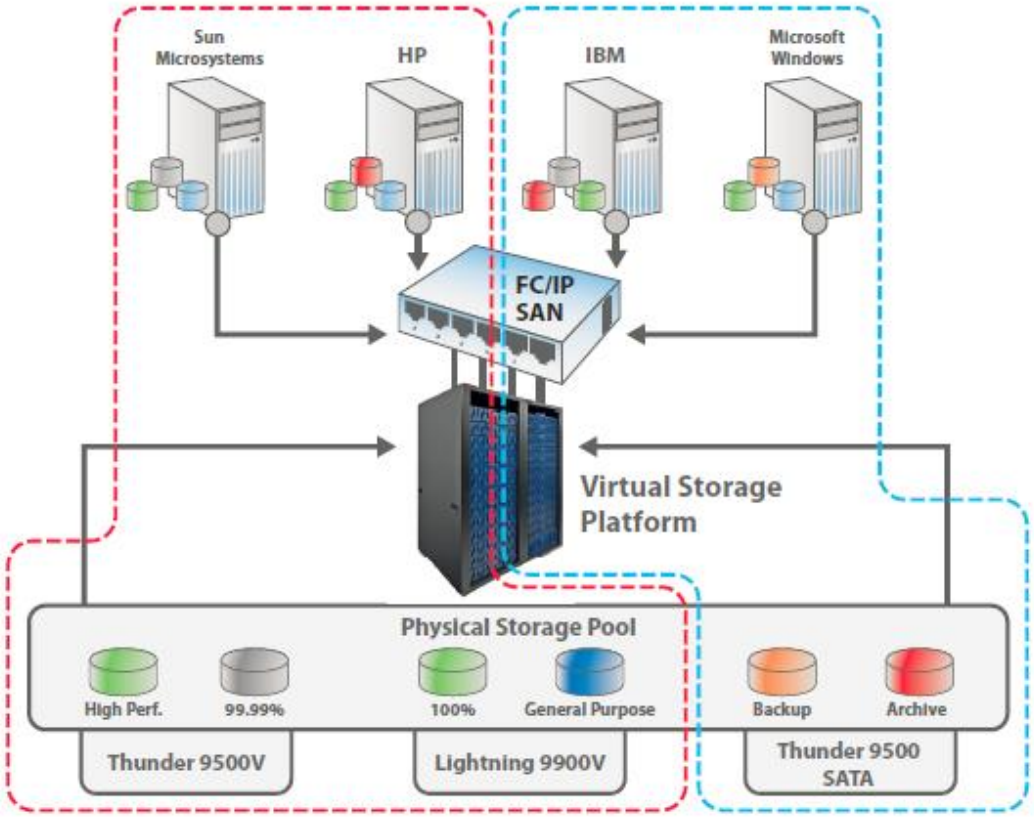

Exercise: Benefits

- From your point of view, what do you think are the benefits of partitioning storage?

If a vILT class, send your answers via the CHAT or over the phone

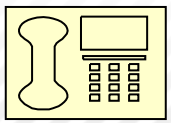

Dynamic (Automated) Tiering

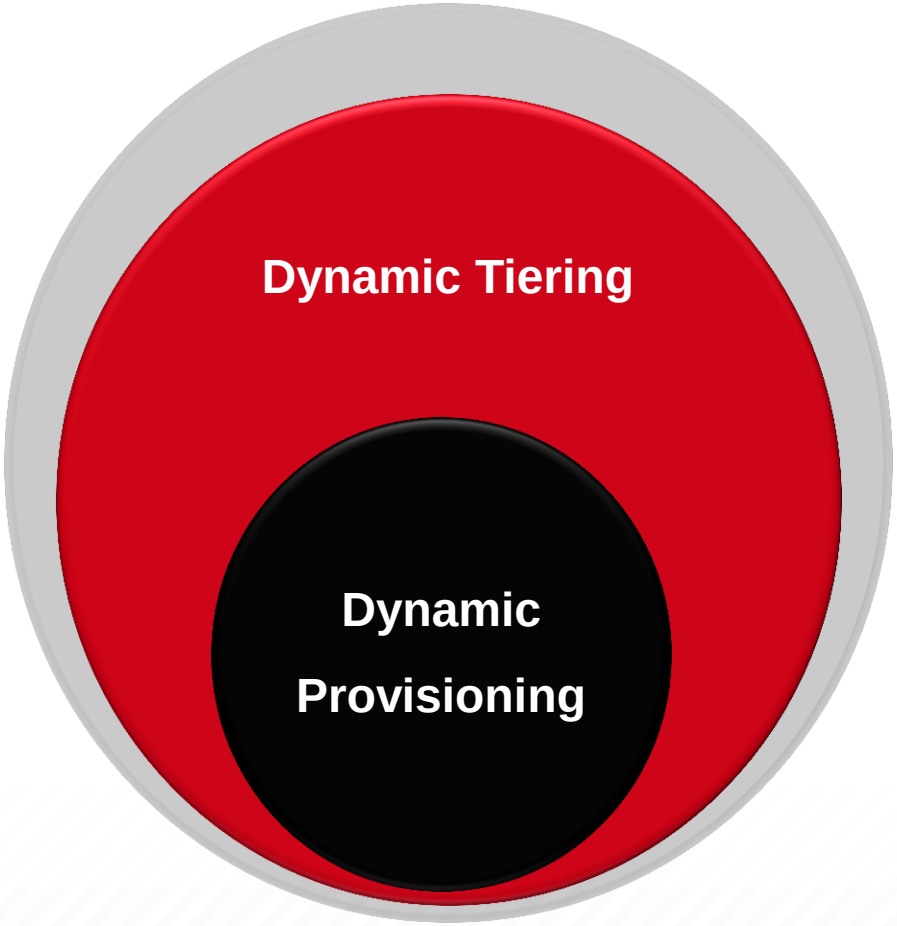

All the benefits of Dynamic Provisioning

+

- Further simplified management
- Further reduced OPEX
- Better Performance

Dynamic Tiering in Hitachi Virtual Storage Platform

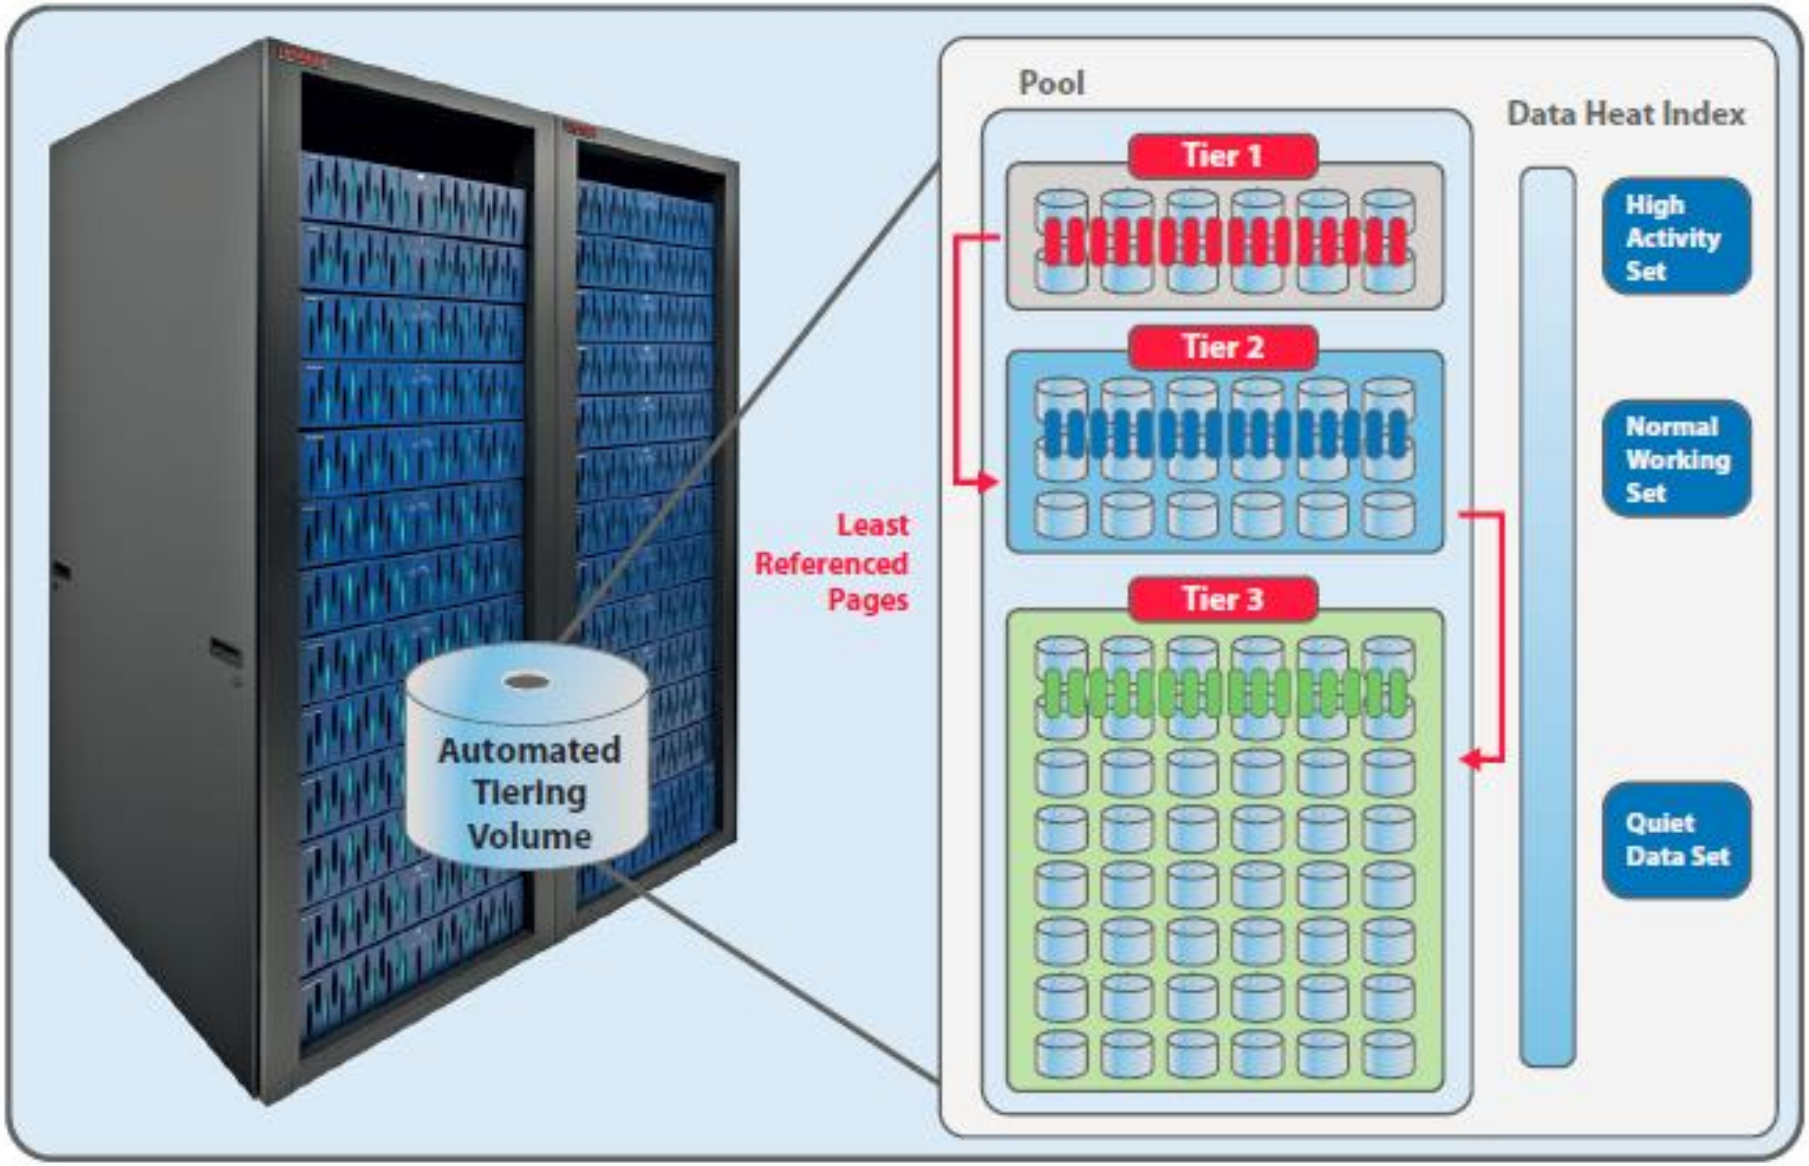

The Future of Virtualization

Storage systems are heading towards fully virtualized solutions.

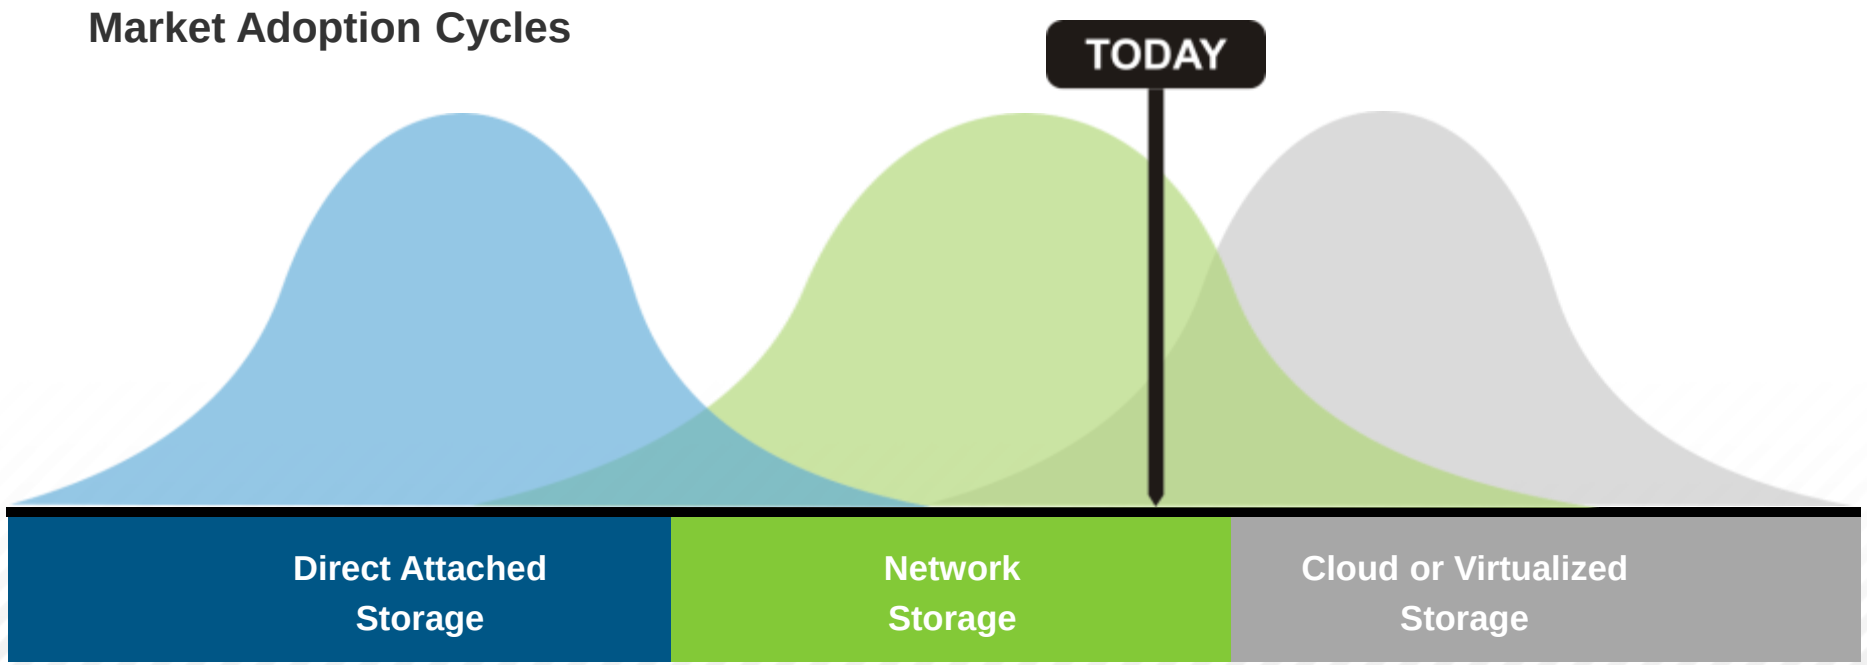

Cloud Storage

"Cloud is a way of using technology, not a technology in itself – it is a self-service, on-demand pay-per-use model. Consolidation, virtualization and automation strategies will be the catalysts behind cloud adoption."
– The 451 Group

Key characteristics of the cloud are:

The ability to scale and provision dynamically in a cost efficient way

The ability to make the most of new and existing infrastructure without having to manage the complexity of the underlying technology

Cloud architecture can be:

- **Private: Hosted within an organization's firewall**
- **Public: Hosted on the internet**
- **Hybrid: A combination of private and public**

Module Summary

- Upon completion of this module, you should have learned to :
  - Understand and explain virtualization concepts and its benefits
  - Explain the difference between *fat* and *thin* provisioning
  - Describe SAN virtualization concepts

Storage Concepts

Archiving and File and Content Management

Module Objectives

- Upon completion of this module, you should be able to:
  - Explain the basic concepts and features of archiving
  - Describe what is meant by Fixed Content
  - Understand the differences between archiving and backup

Module Topics

- Fixed Content and its characteristics
- Components of a digital archive
- Content management

Introduction to Archiving

What is **Fixed Content?**

- Data objects that have a **long-term value**, **do not change over time**, and are easily accessible and secure

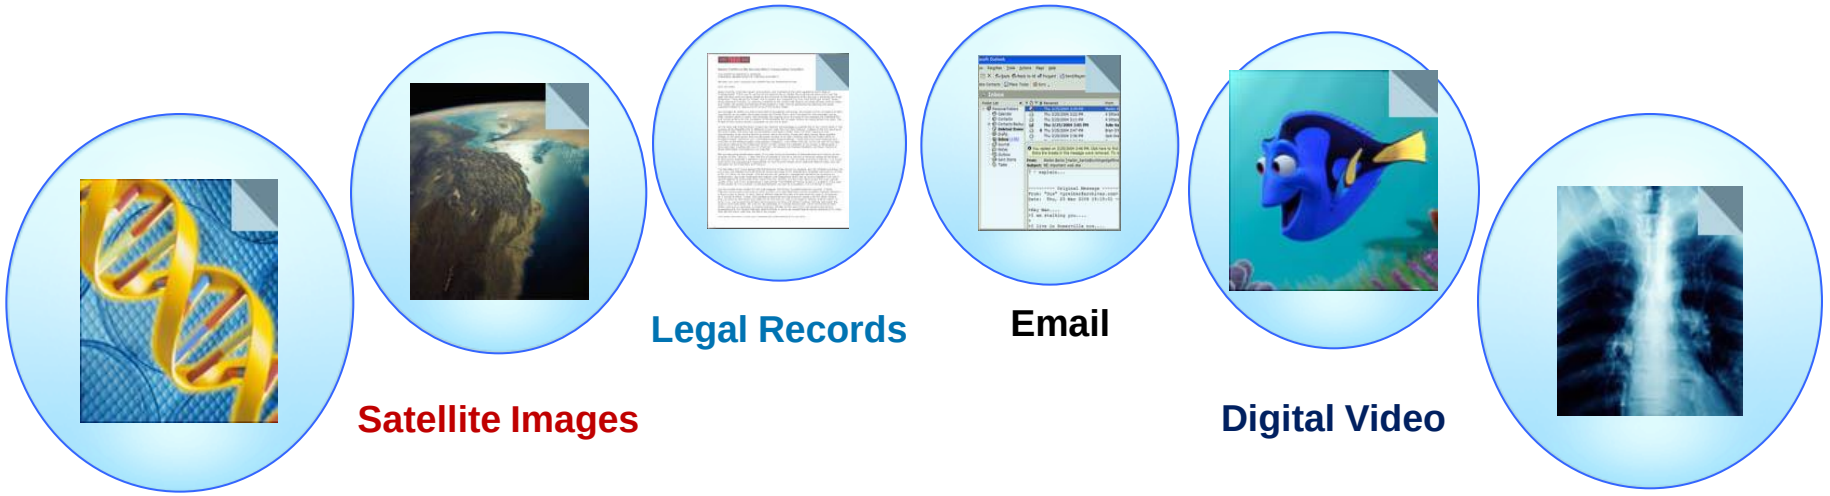

Motivation for Data Archiving

There are several reasons for implementing data archiving policies and technical solutions. Some of the most prominent reasons are:

- Effective utilization of high performance tiers
- Cheaper storage for fixed content
- Data retention regulation
- Simplified content management
- Indexing and searching capabilities of a digital archive

Archive Solutions?

Seeking an archiving solution. The storage systems we have discussed up until now are not very suitable for fixed content storage and archiving.

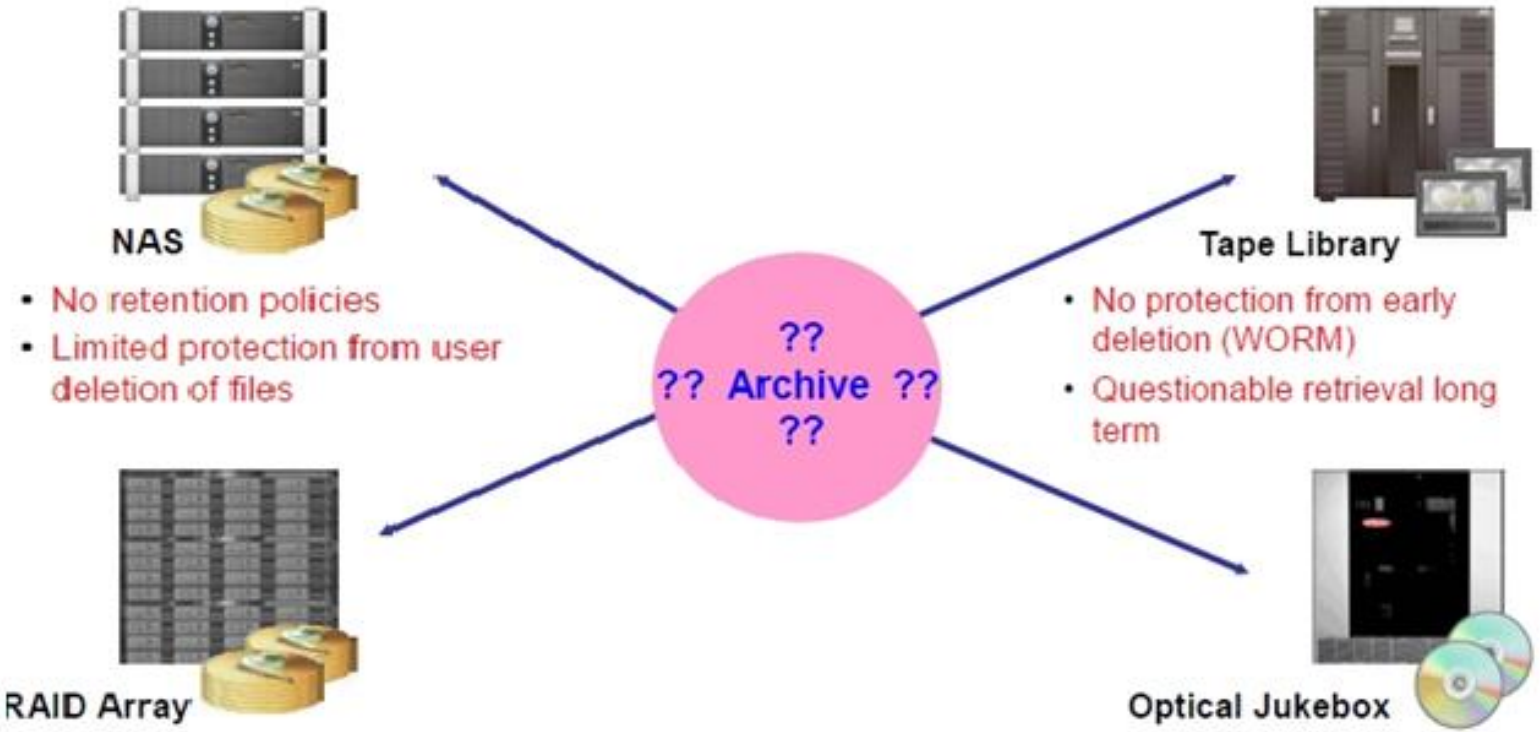

Legal Requirements of Data Retention Periods

Many organizations have increasing regulations, especially in the pharmaceutical industry, the food processing industry, healthcare, financial services and auditing. The Sarbanes-Oxley Act very strictly regulates the length of retention of financial records and accounting in companies.

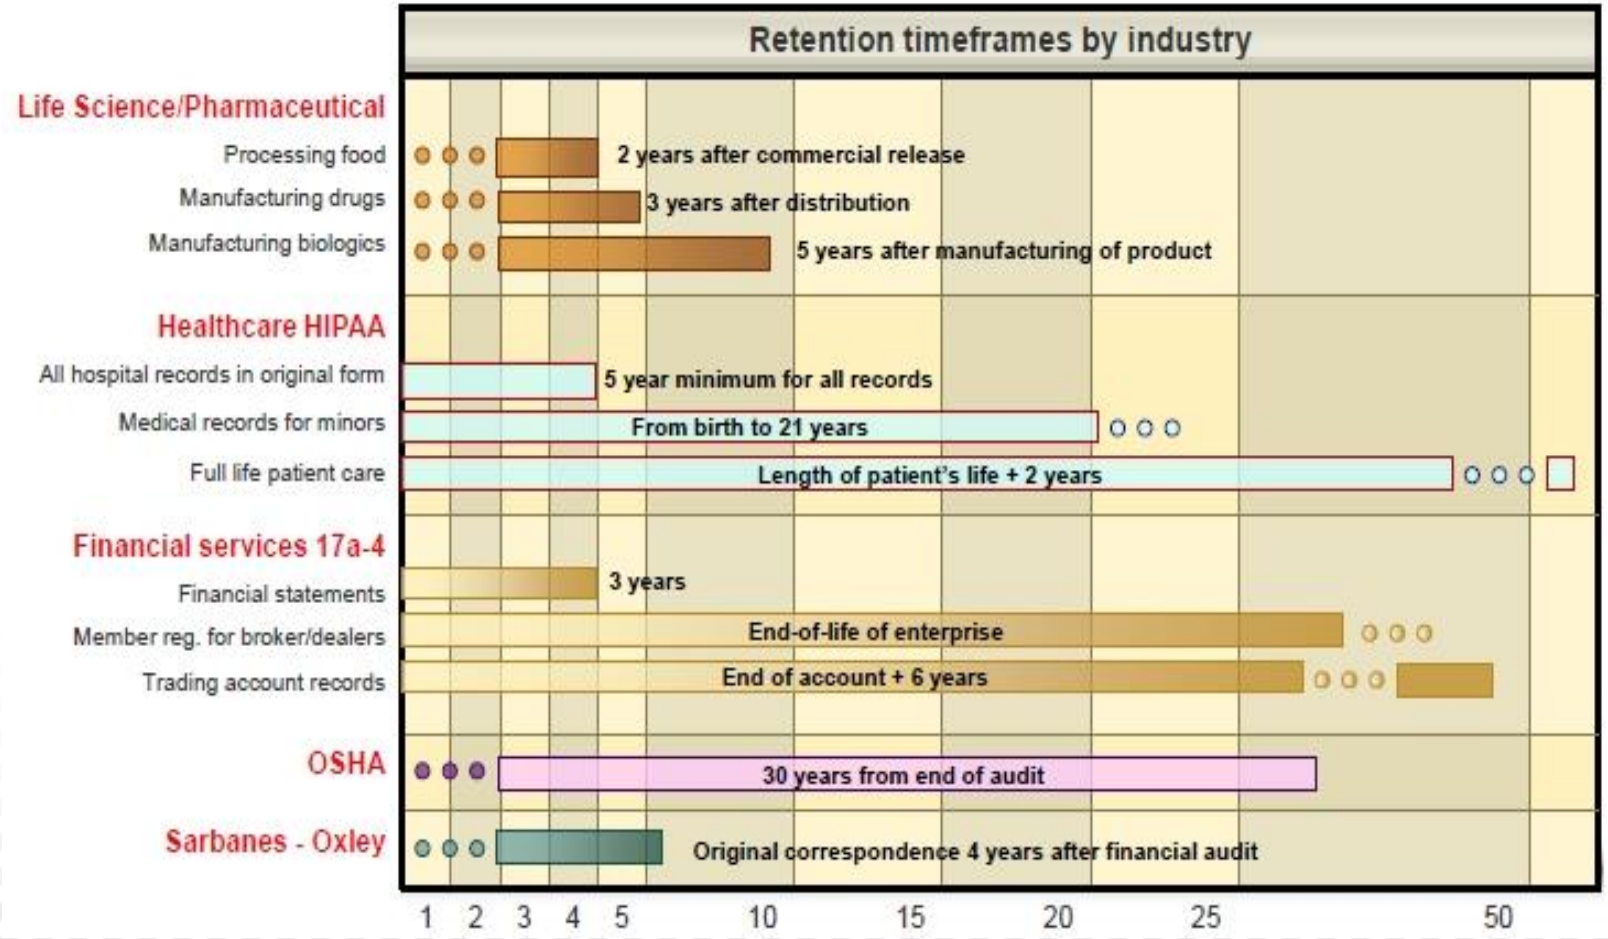

The Need for a Better Archiving Solution

An example of a decentralized and fragmented archiving solution. Disparate storage systems do not provide a common search engine, and they are not very scalable. A digital archive can solve this problem.

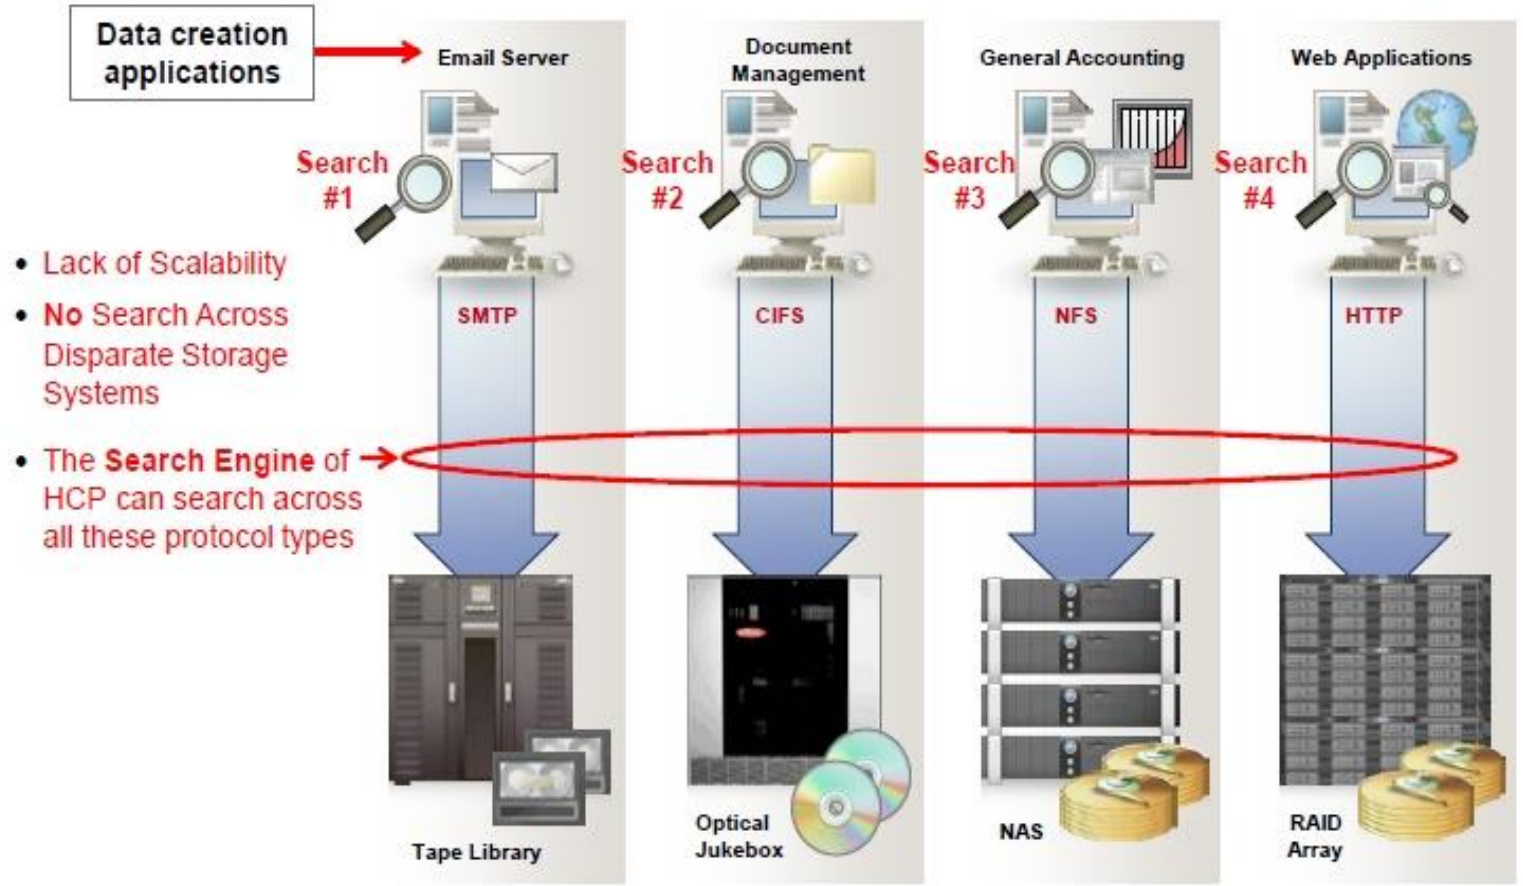

A Digital Archive

A digital archive works on the object level. Each object contains fixed content data, metadata and description of policies.

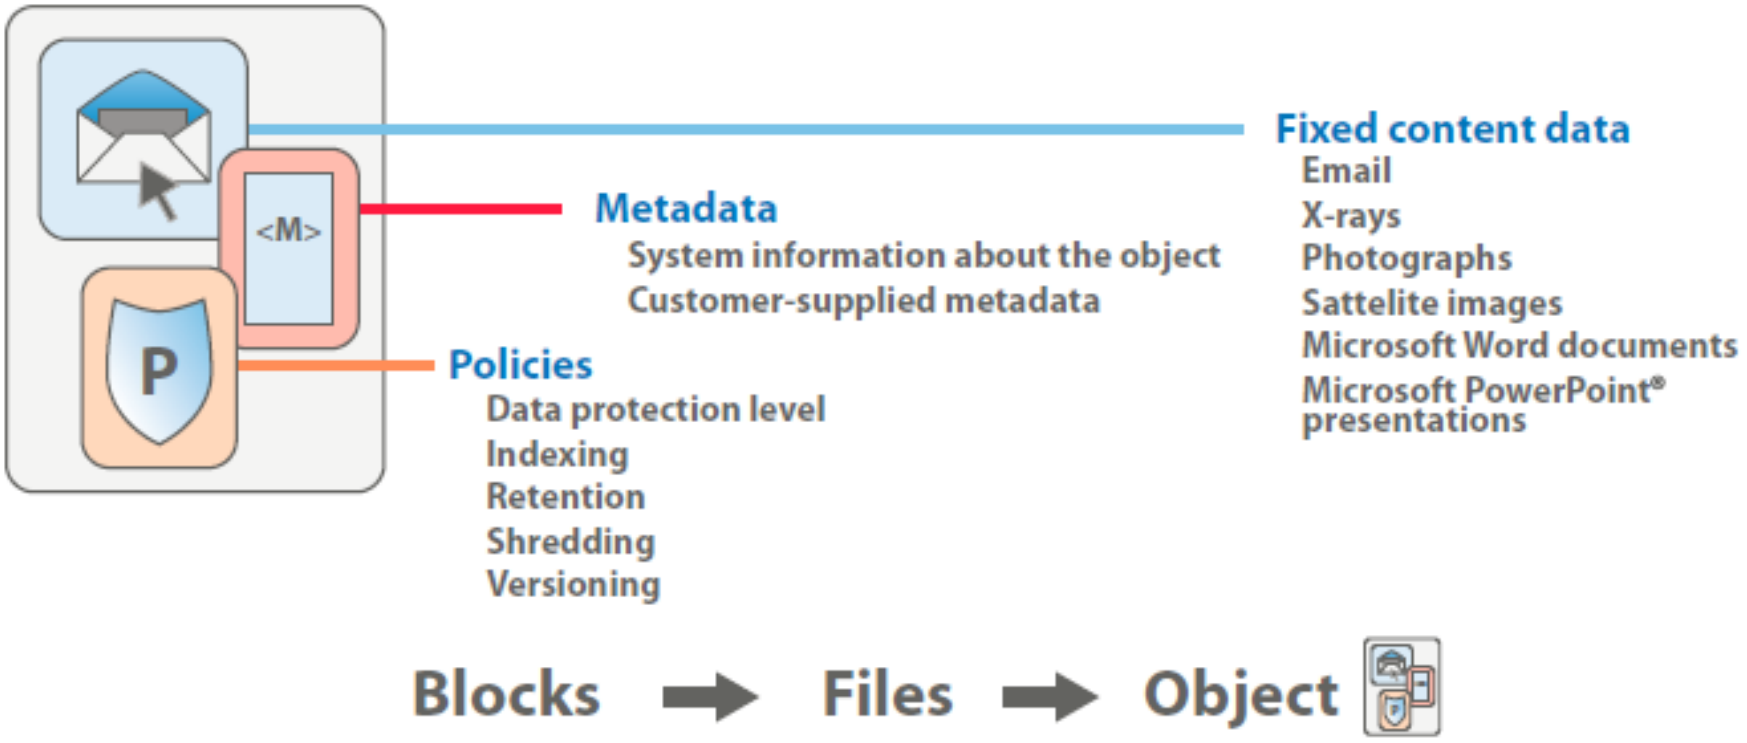

Block Level Storage Compared to Object Level Storage

A traditional block level storage system compared to an object level storage system. The object level storage system consists of powerful proprietary servers and management software. These servers are connected to a RAID array.

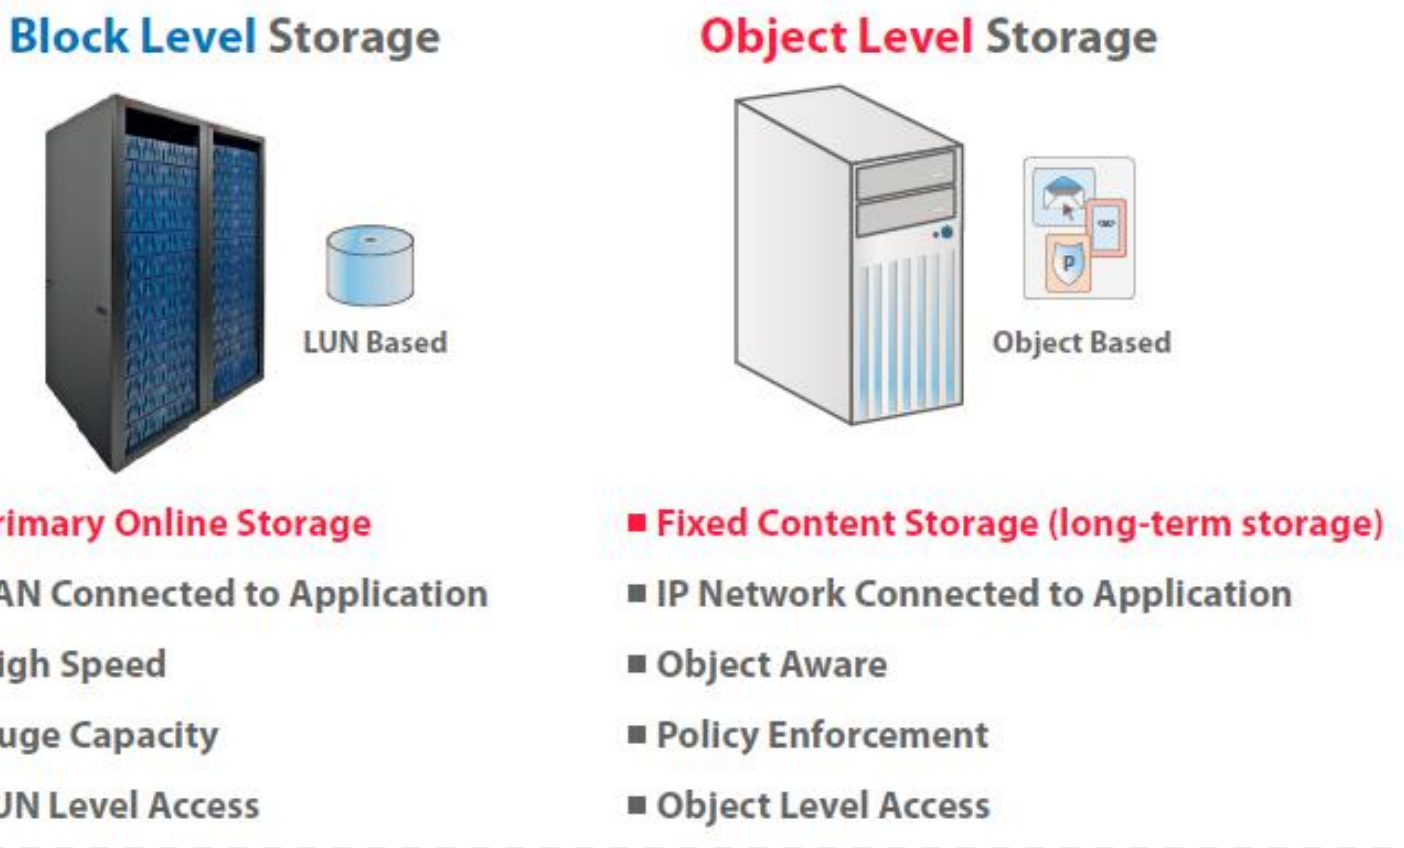

Internal Object Representation

A data object and its components in detail. This example illustrates how objects are handled by a object storage's digital archive.

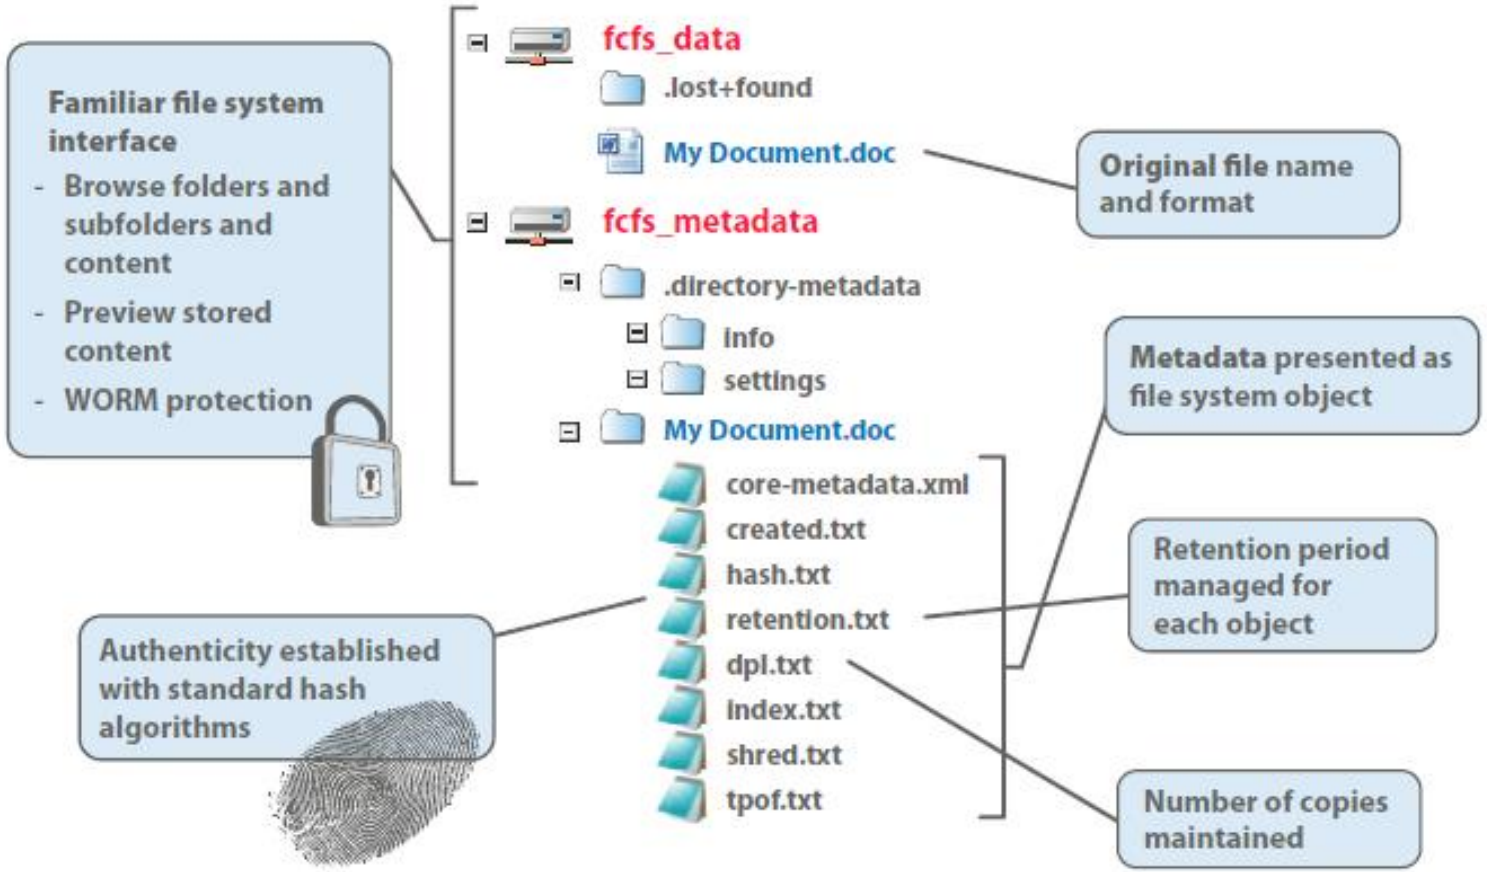

Digital Archive Features

Active functions of a digital archive are:

- Content verification – Ensures authenticity and integrity of each data object
- Protection service – Ensures stability of the digital archive
- Compression service – Achieves better utilization of storage space assigned to the digital archive
- Deduplication service – Detects and removes duplicities
- Replication service – Ensures redundancy of archived data
- Search capabilities – Allows users to search documents

Digital Archive Accessibility

In this example, a digital archive can be accessed using multiple independent standard protocols. WebDAV is an extension to HTTP protocol that allows remote management of files stored on web server

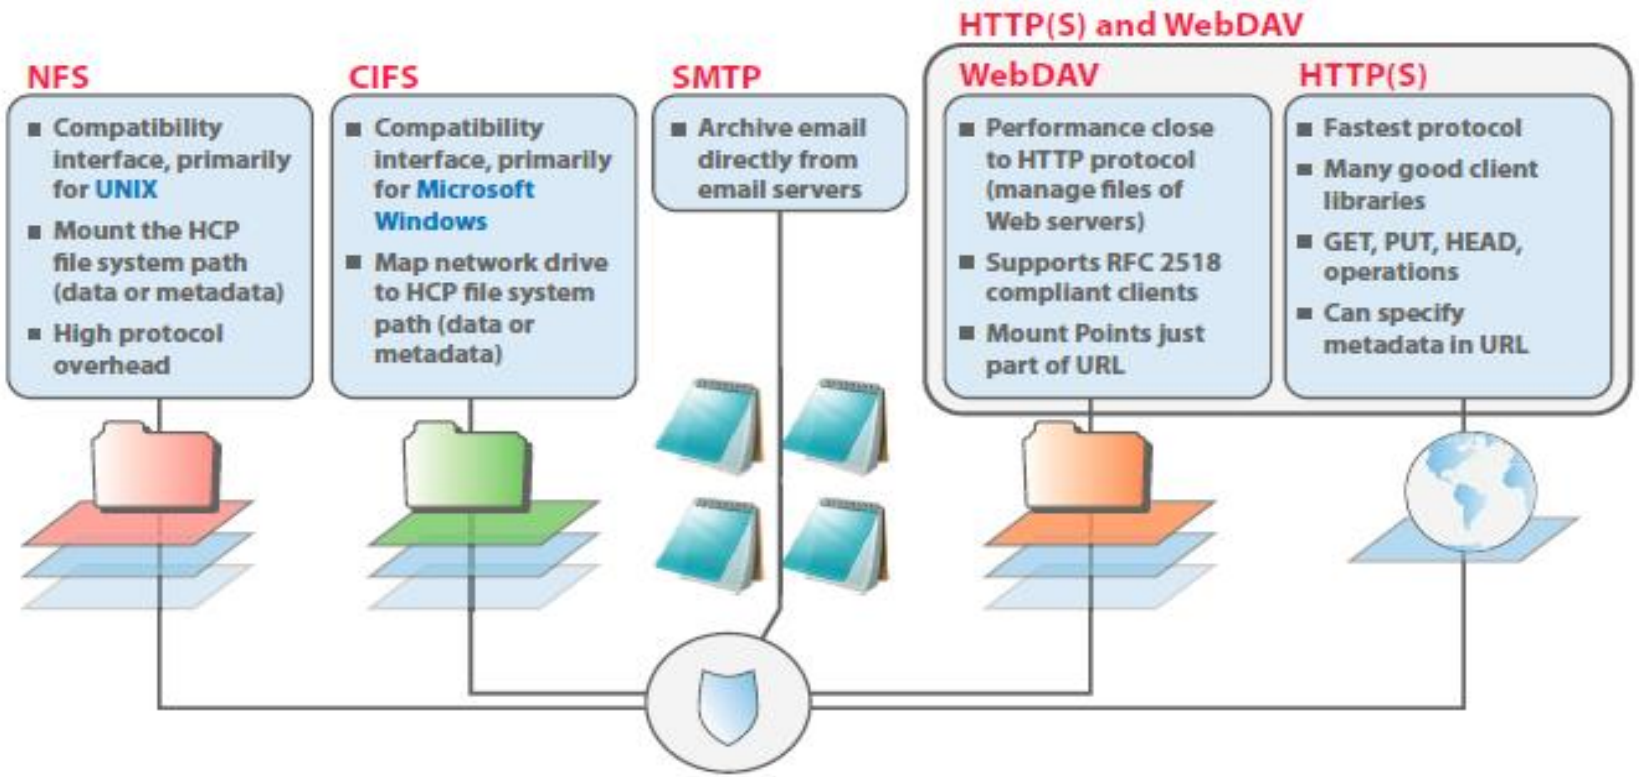

Digital Archive Compliance Features

The most important compliance features of a digital archive are:

- **Write once read many (WORM)**
- **Retention period definition**
- **Data shredding**
- **Data encryption**

Exercise: Fixed Content

Which 2 statements are true about the definition of ***fixed content***? (Choose 2)

Fixed content is …

a. Content that cannot be archived and restored

b. Content that can only be changed by the system administrator

c. Static data that is in a final state

d. Content that will not / cannot change

Answers= ____ and ____

If vILT class, write your answers on blank lines,

Exercise: Backup versus Archive

Explain the differences between what is meant by Backup and Archive.

*If vILT class, Send your answers to the instructor via the **WebEx CHAT tool** or **phone**.*

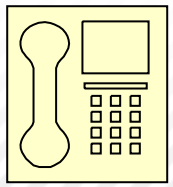

Exercise: Object Representation

Fixed-content data (Data)

System metadata

Custom metadata

1. Describe what each is.
2. What does an object contain?

Send your answers using CHAT

If vILT class, write your answers on blank lines

Module Summary

- Upon completion of this module, you should have learned to:
  - Explain the basic concepts and features of archiving
  - Describe what is meant by Fixed Content
  - Understand the differences between archiving and backup

Storage Concepts

Storage System Administration

Module Objectives

- Upon completion of this module, you should be able to:
  - Describe everyday storage administrator tasks
  - Explain how to configure and monitor storage systems
  - Describe tools used by the storage administrator in managing storage

Module Topics

- Storage system administrator tasks and common functions
- Storage system management software
- Storage system implementer tasks


Storage
Administrator

Who is a Storage System Administrator?

- In charge of maintaining a storage system infrastructure
  - Tasks based on Service-level Agreements (SLAs)

Storage Administrator Tasks

- **Capacity Management**

  Amount of data to be stored; size and performance of LUNs; hard drive performance; I/O performance and R/W operations

- **Availability Management**

  Replication, backup, and archive strategies; protection against component failures

- **Continuity Management**

  Part of business continuity planning and disaster recovery procedures

- **Financial Management**

  Budget preparation; cost calculation and invoicing and TCO

Storage Administrator Tasks – Other Common Operations

- Configuration of RAID groups and volumes
- Implementation of changes in volume configuration
- Data replication optimization
- Configuration of cache
- Cache partitioning
- Backup of storage system configuration

Integration of a New Storage System

When purchasing a new storage system, the storage administrator must think through the whole implementation process, including the following items:

- Storage system model
- Switch model
- Cabling
- Rack usage and floor space
- Power requirements
- Air conditioning
- LAN infrastructure

Tasks of a Storage System Implementer

- Installation and initial configuration of the storage system
- Basic training to familiarize the customer with the new device.
- Conduct all hardware replacement and upgrade procedures
- Monitor the storage system remotely
- Help with performance tuning
- Microcode updates

Module Summary

- Upon completion of this module, you should have learned to:
  - Describe everyday storage administrator tasks
  - Explain how to configure and monitor storage systems
  - Describe tools used by the storage administrator in managing storage

Storage Concepts

Business Challenges

Module Objectives

- Upon completion of this module, you should be able to:
  - Identify business challenges driving the need for storage
  - Understand why storage systems are important for business
  - Explain the energy and green issues faced by today's businesses

Module Topics

- Business challenges companies face
- Advanced data classification and tiered storage
- Data center operations environmental concerns

Business Challenges

Business Challenges

Business challenges can be classified as follows:

- **Accelerating storage growth** — need more and more capacity
- **Increasing requirements on high availability** — cannot afford any disruptions because our data has become too important for our business
- **Fast and effective response to business growth** — need to be able to react quickly to new conditions
- **Heterogeneous infrastructure** — result of fast infrastructure growth, which was not properly planned, causing TCO to increase rapidly
- **Compliance and security challenges** — need to process, protect and retain data according to legal requirements and regulations
- **Power and cooling challenges** — pay too much for electricity and cooling and may be running out of space in the server room
- **Data center challenges** — specific needs for those whose business is focused primarily on cloud type provision of storage and computing capacity

Data Growth Forecast

An overview of the storage requirements in the past years provide you with the necessary information to forecast data growth.

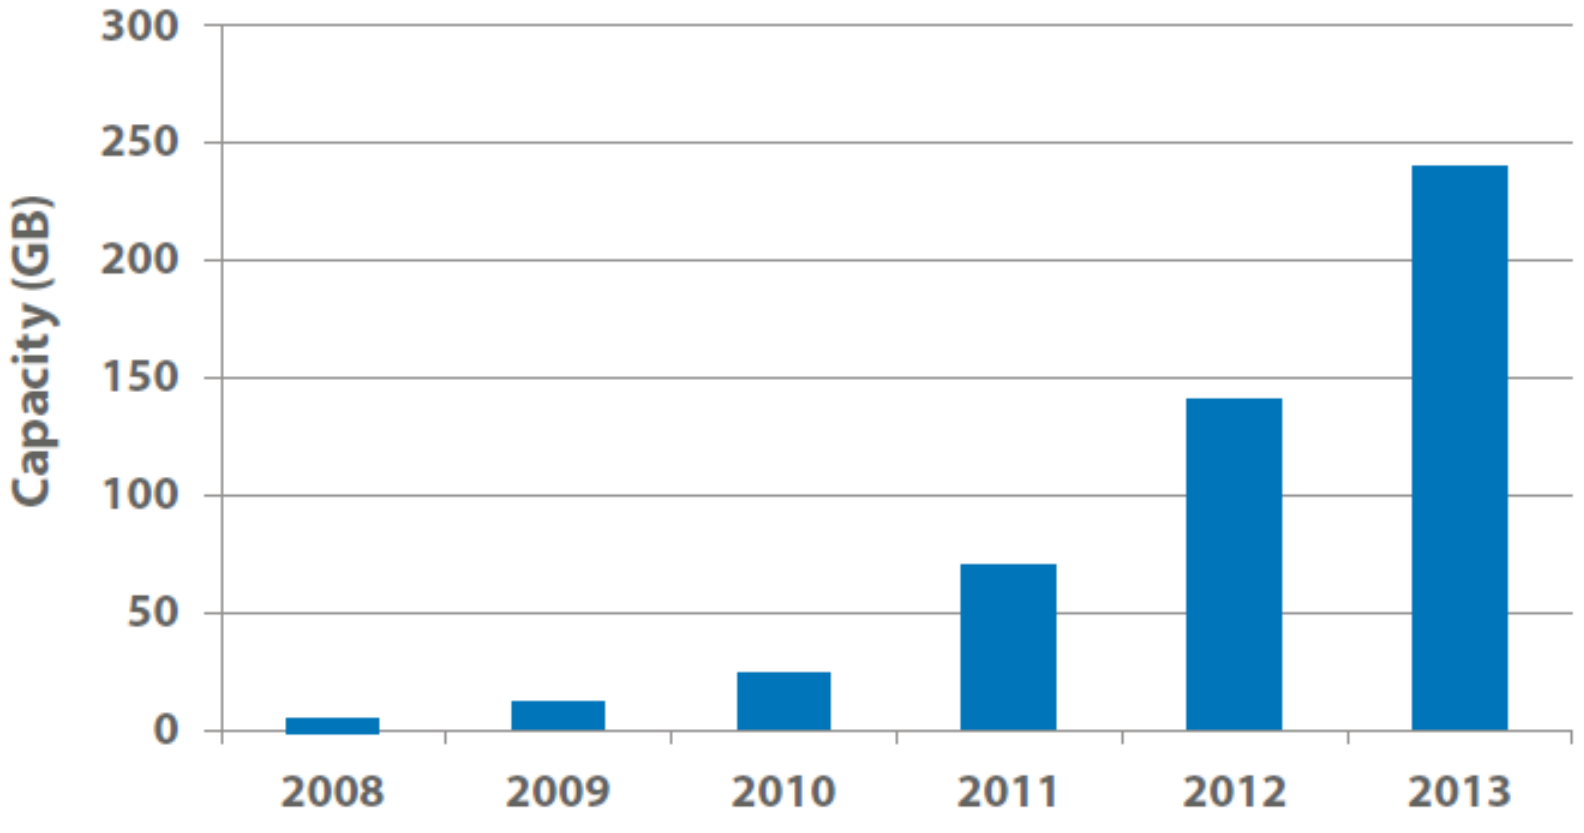

Data Growth Forecast – Tier View

Data growth forecast in relation to performance tiers

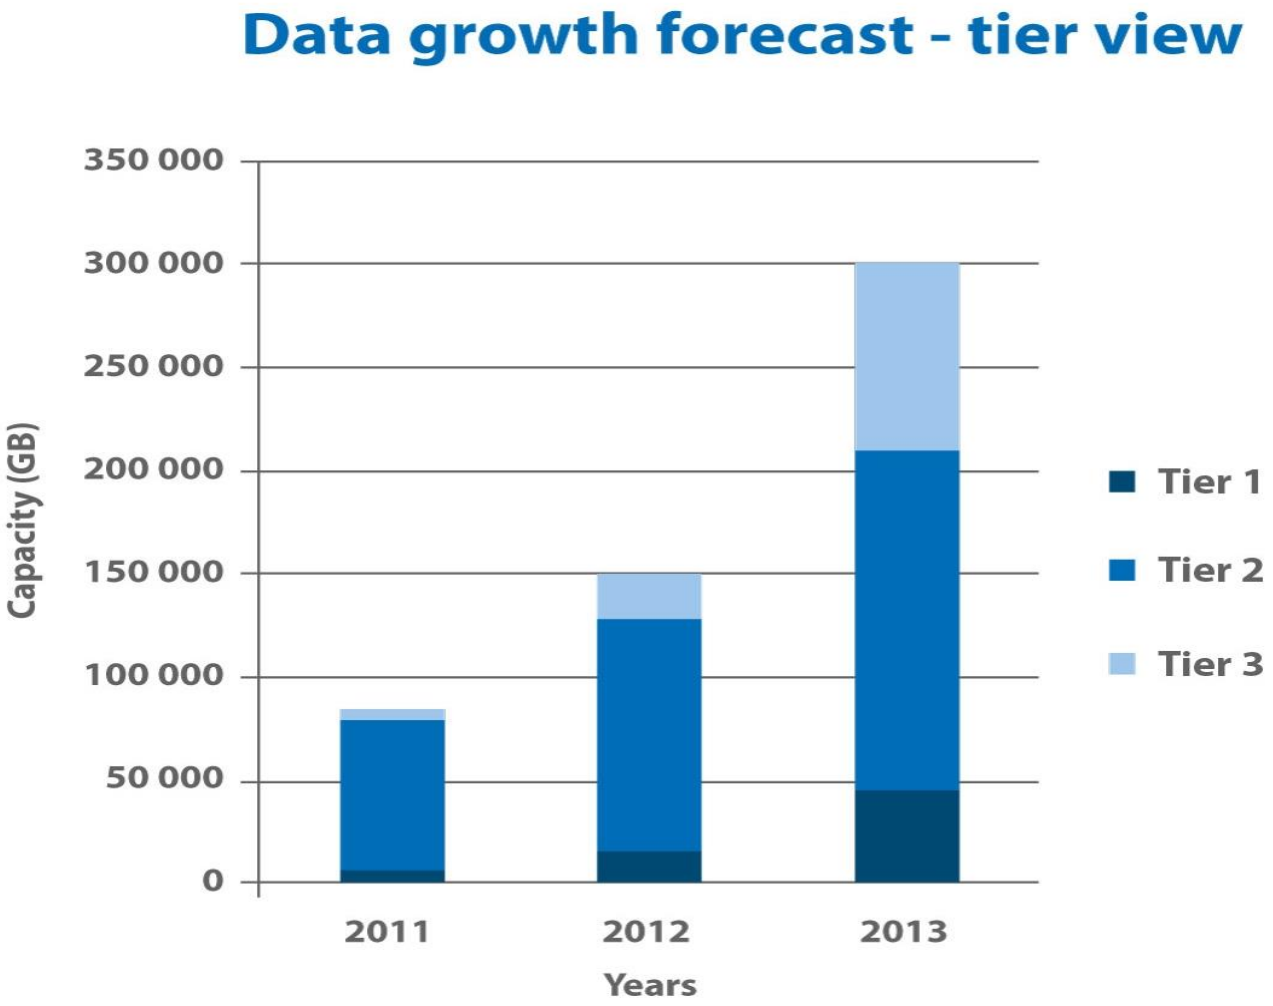

Structured, Unstructured and Replicated Data Growth

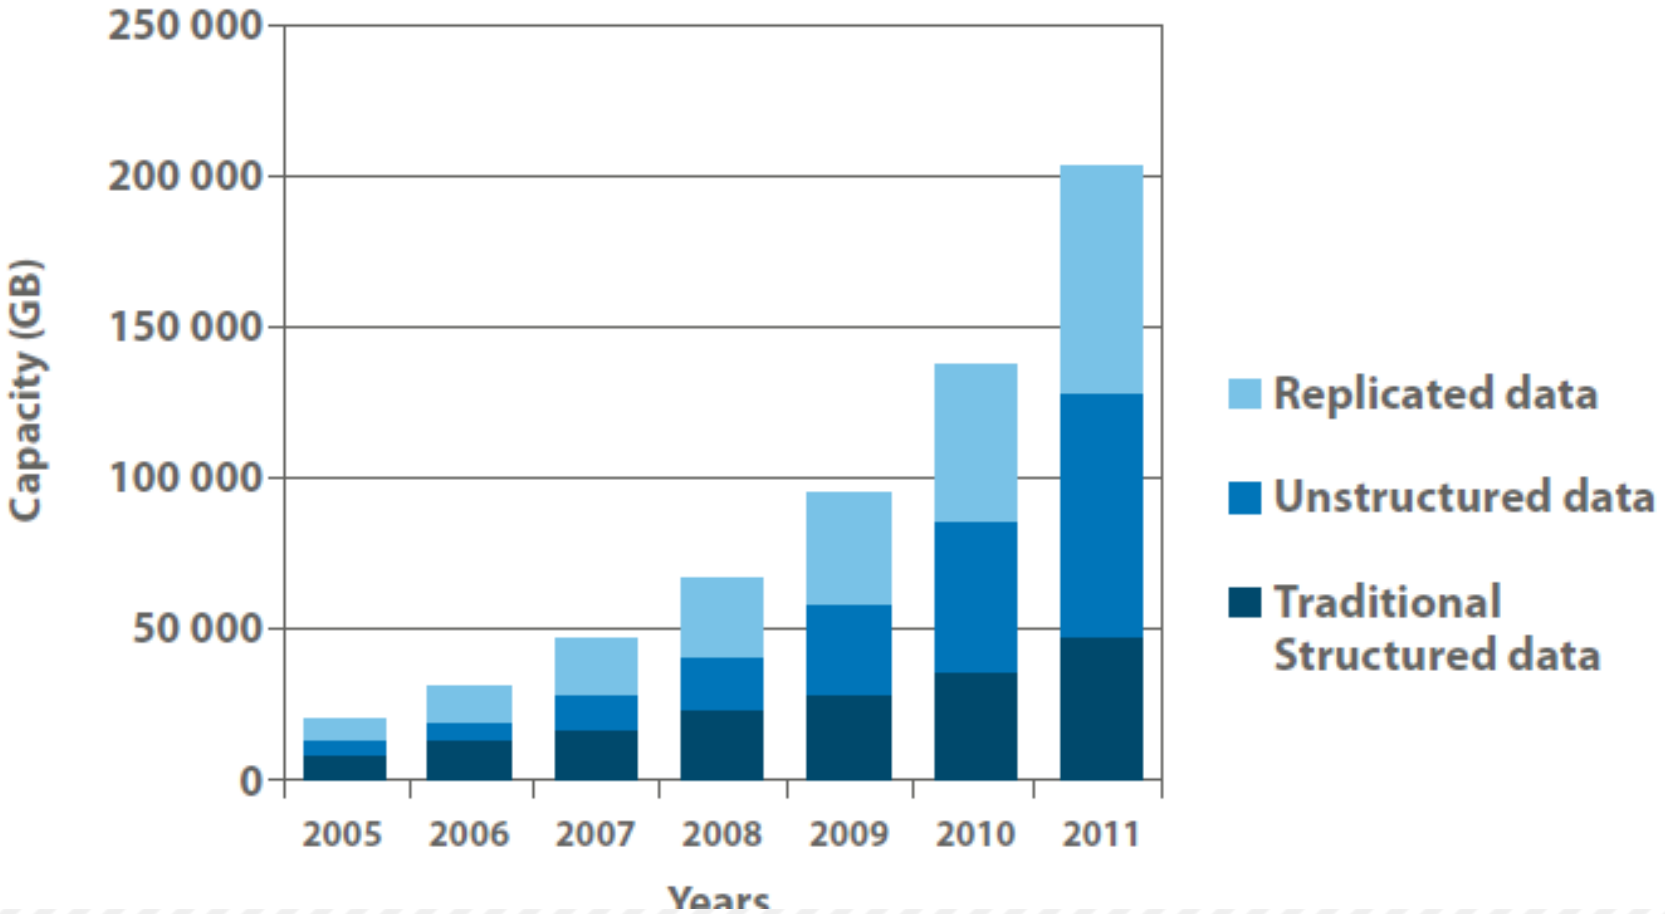

Tiered storage infrastructure

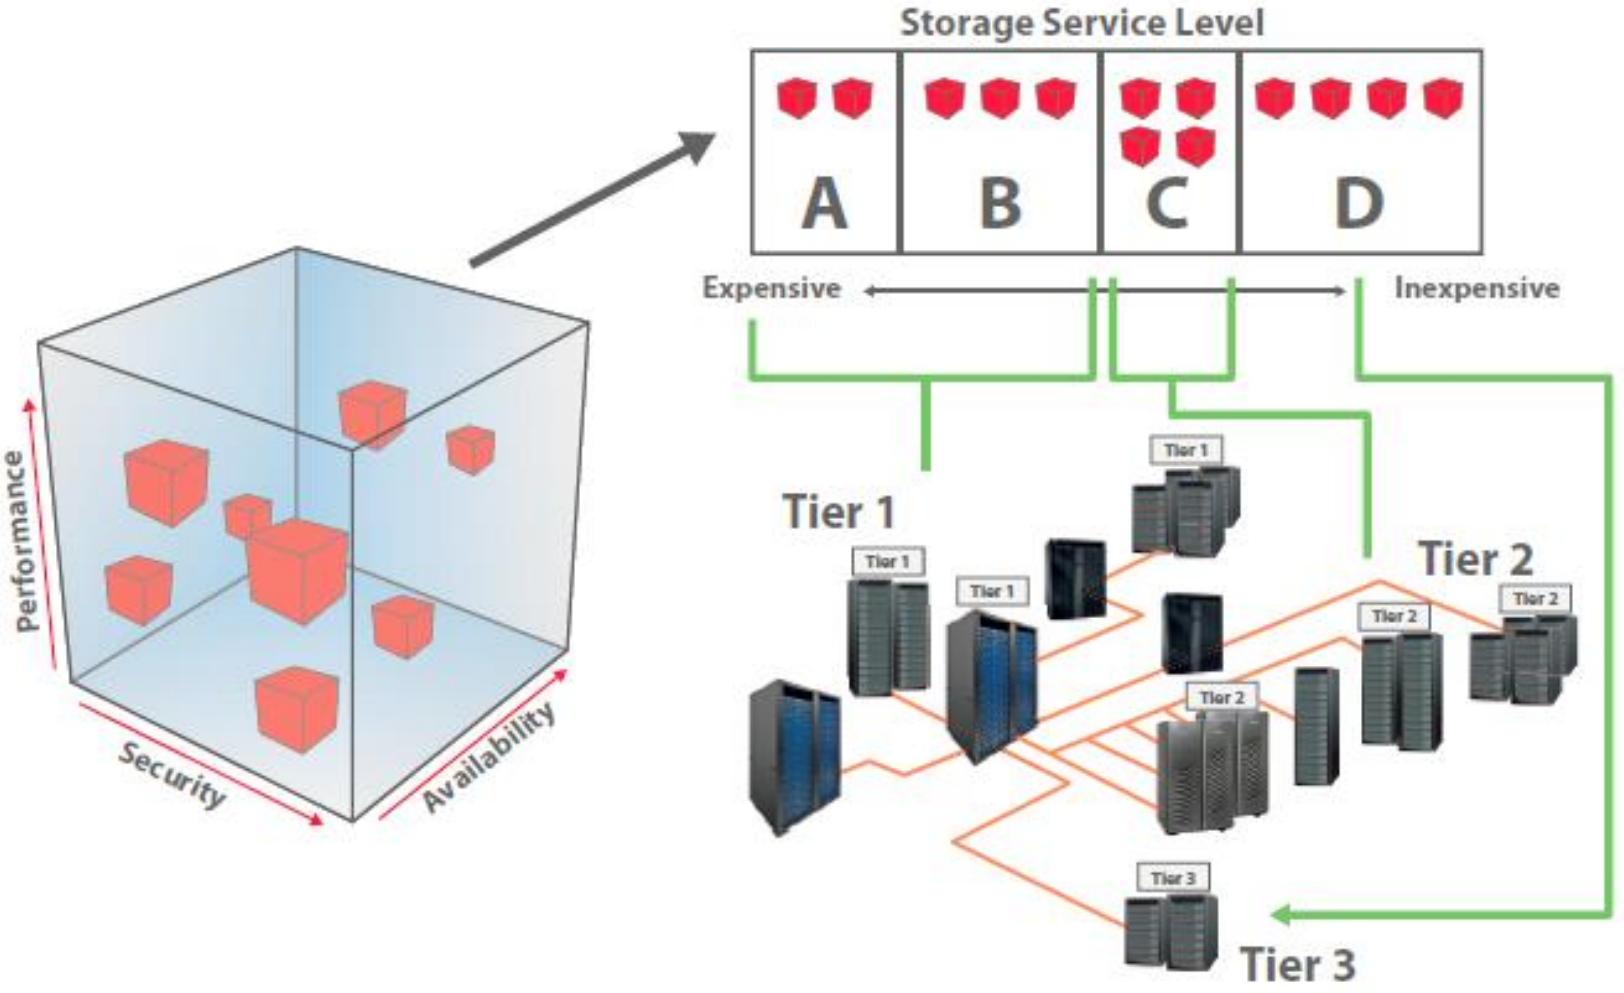

Power Requirements and Cooling

Greatest issues organizations and companies face

What is the greatest facility problem with your primary data center?

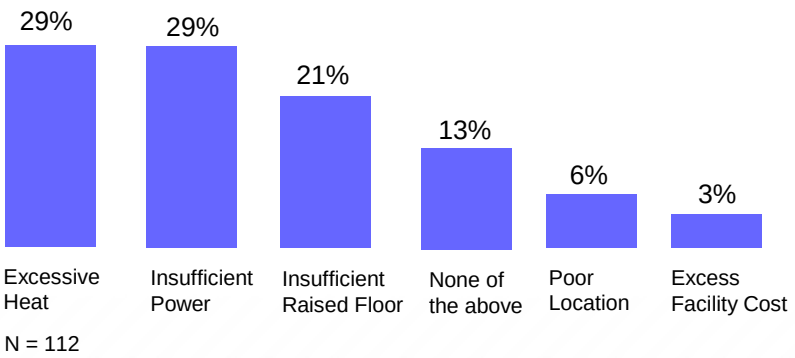

Power and cooling exceeds server Spending

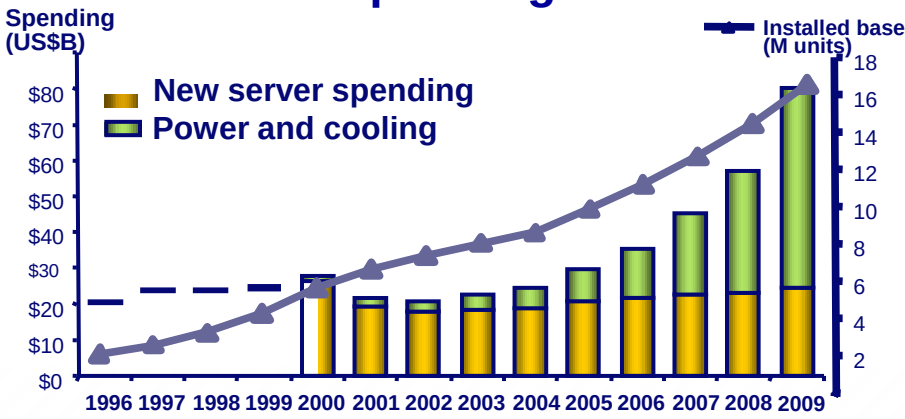

Gartner, Best Practices in Data Center Facilities

Power Challenges Facing Data Centers

Some of the challenges facing data centers with respect to electricity, cooling and environmental requirements include:

- Running out of power, cooling and space
- Growing energy costs
- Increasing regulatory compliance issues
- Data center expansion without consideration for future power and cooling requirements
- Data storage configured without adequate consideration to heat distribution (equipment racks should be installed with cold rows and hot rows
- Difficulty relieving data center hot spots without disrupting applications

Other Challenges Facing Data Centers

In addition to power consumption metrics in kW, there are other metrics that should be considered:

- Total five-year power and space costs
- Heat loading (kW/sq ft)
- Space requirements (sq ft)
- Floor loading (lbs/sq ft)

Controller-based virtualization and thin provisioning can also yield substantial environmental advantages because they reduce the need for storage capacity.

HDD and Fan Power Savings

Features that work on the HDD and fan level and lead to significant power savings:

- Spin down drives in selected RAID groups
- SATA drives will park heads when idle for more than 2 hours
- Adjust fan speeds to maintain correct temperatures
- Keep data in cache as long as possible

Green Data Center

Applies principles from Eco friendy Data Center Project and IT Power Saving Plan

- Location: Yokohama
- Operation: 2009
- Dimension: 10,000m²
- Power Usage effectiveness rating: 1.6

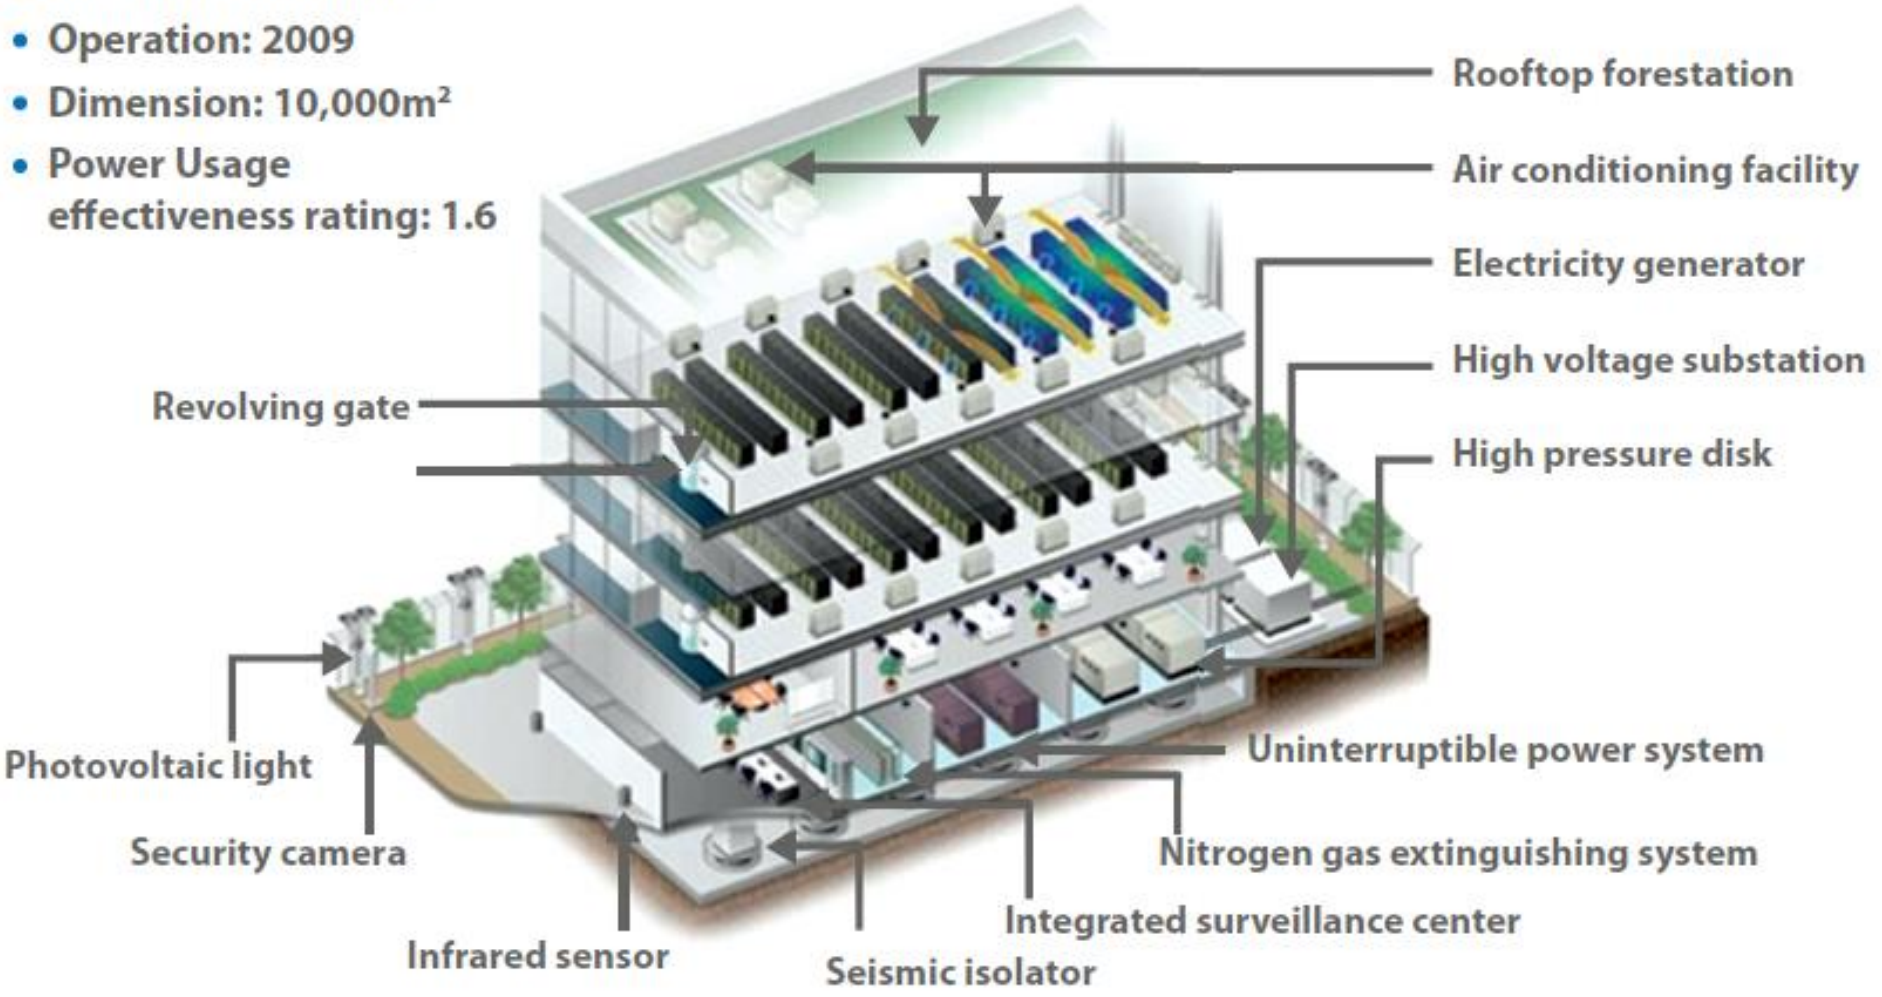

Hot and Cold Rows

Arrange racks in alternating rows with cold air intakes facing one way and hot air exhausts facing the other

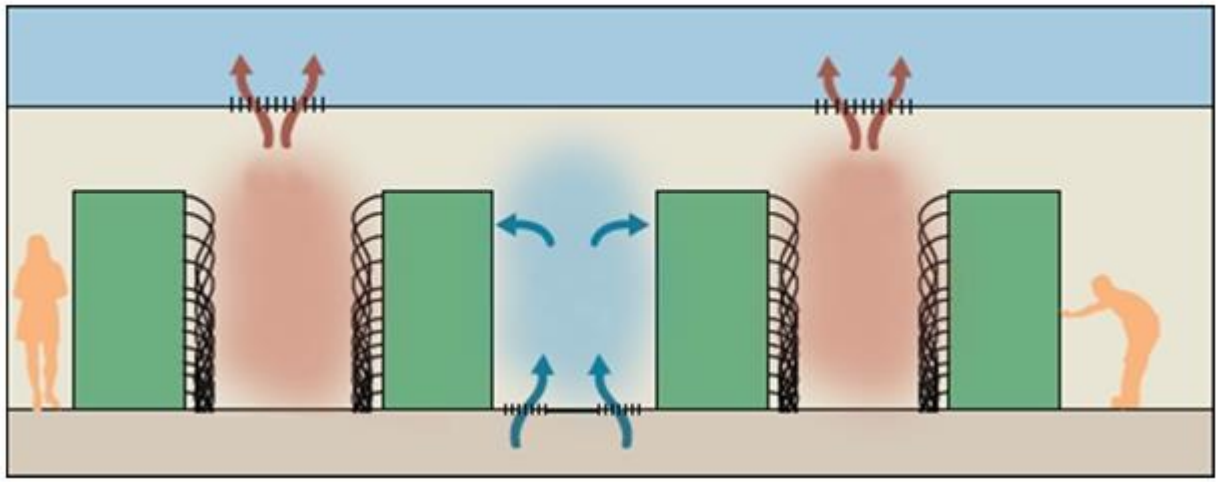

Module Summary

- Upon completion of this module, you should have learned to:
  - Identify business challenges driving the need for storage
  - Understand why storage systems are important for business
  - Explain the energy and green issues faced by today's businesses

Storage Concepts

Storage Networking and Security

Module Objectives

- Upon completion of this module, you should be able to:
  - Describe basic networking concepts
  - Explain how common network devices operate
  - Explain how devices communicate in a network
  - Explain storage area network security

Module Topics

- Basic networking concepts
- Operations of common network devices
- Possibilities we have in storage system networking
- How devices communicate with each other through the network
- How to secure Storage Area Networks

Introduction to Networks – Components

Twisted Pair Cable structure, RJ45 connectors

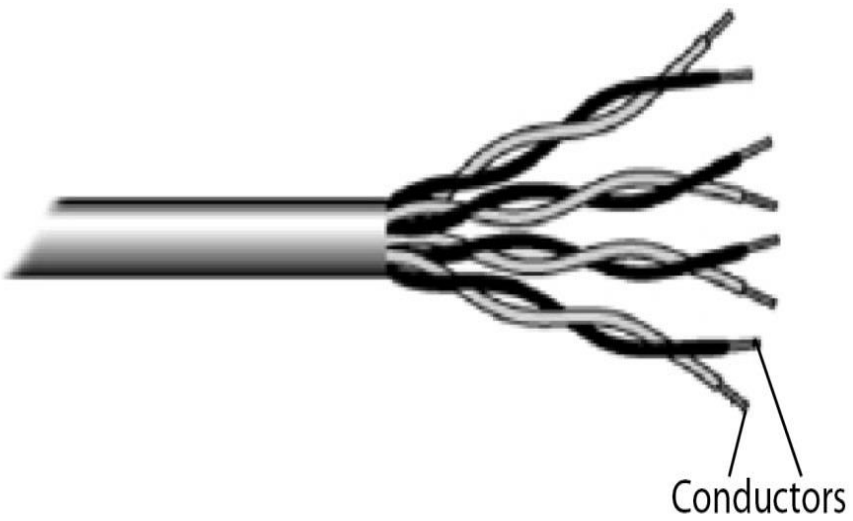

Fiber Optic Cable

Fiber Optic Cables – Fiber optics is a technology that uses glass or plastic fibers to transmit data as light impulses. A fiber optic cable consists of a bundle of fibers and each fiber can transmit millions of messages modulated onto light waves.

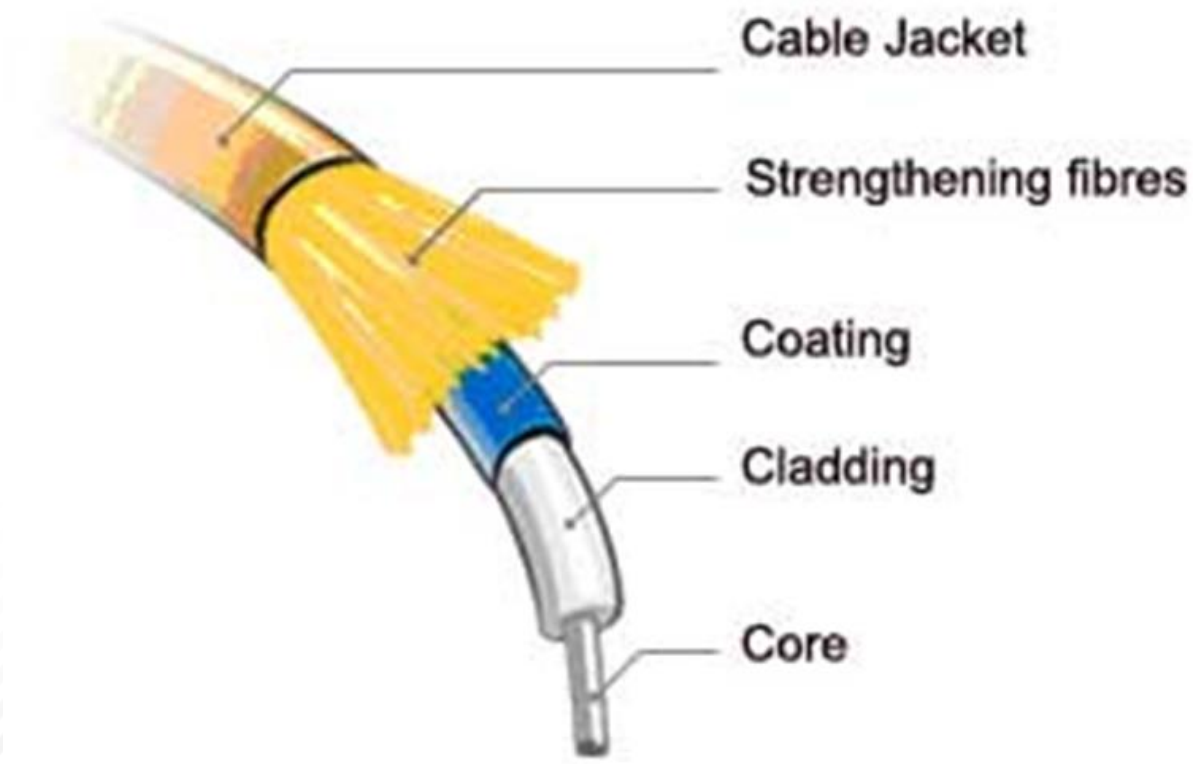

Fiber Optic Connectors

Fiber Optic Cable LC connector and SFP transceiver

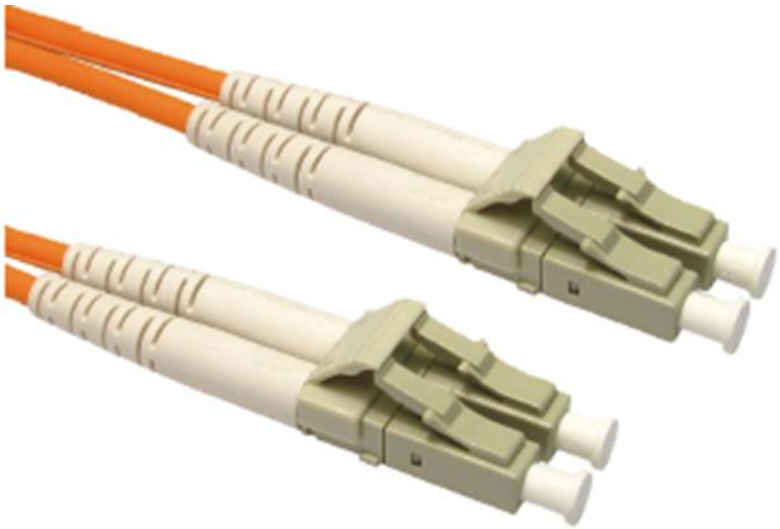

Storage Network Components – Nodes

Network nodes are all the devices connected in the network. We distinguish between endpoint communication nodes and data redistribution nodes.

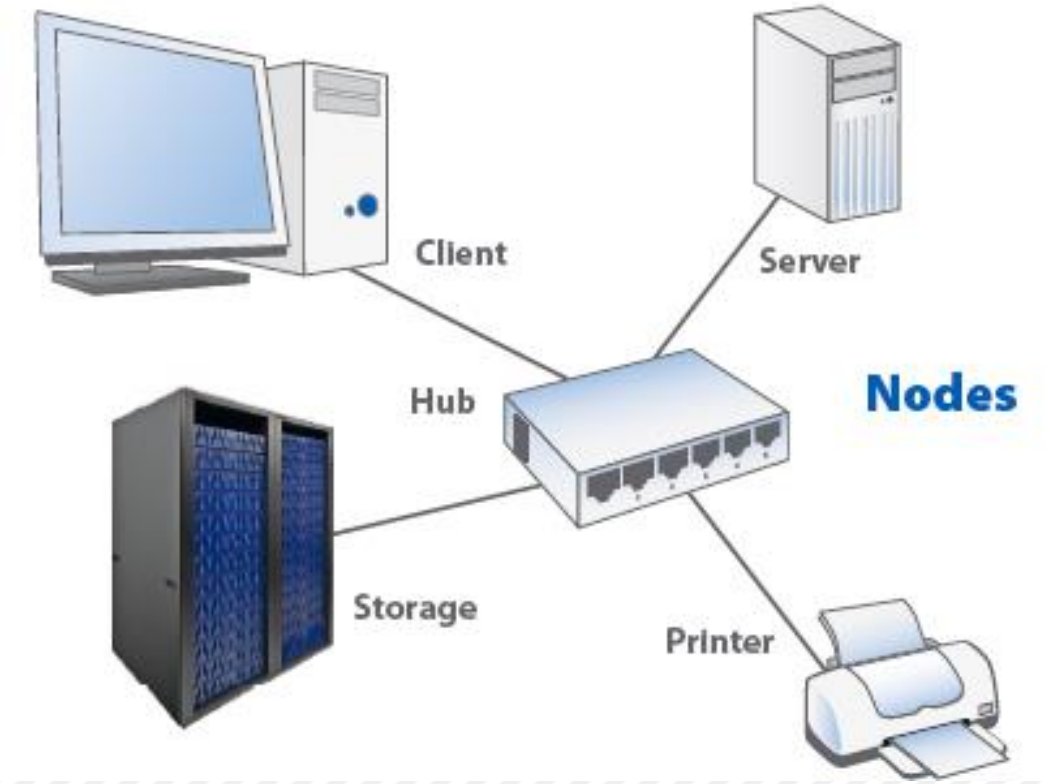

Storage Network Components – Ports

- **Ports**

On a storage network, a port enables a node to communicate with another node over a Fibre Channel connection.

- A node can contain multiple ports.

On a storage network, a port enables the following connections:

- Server to switch
- Switch to switch
- Switch to storage

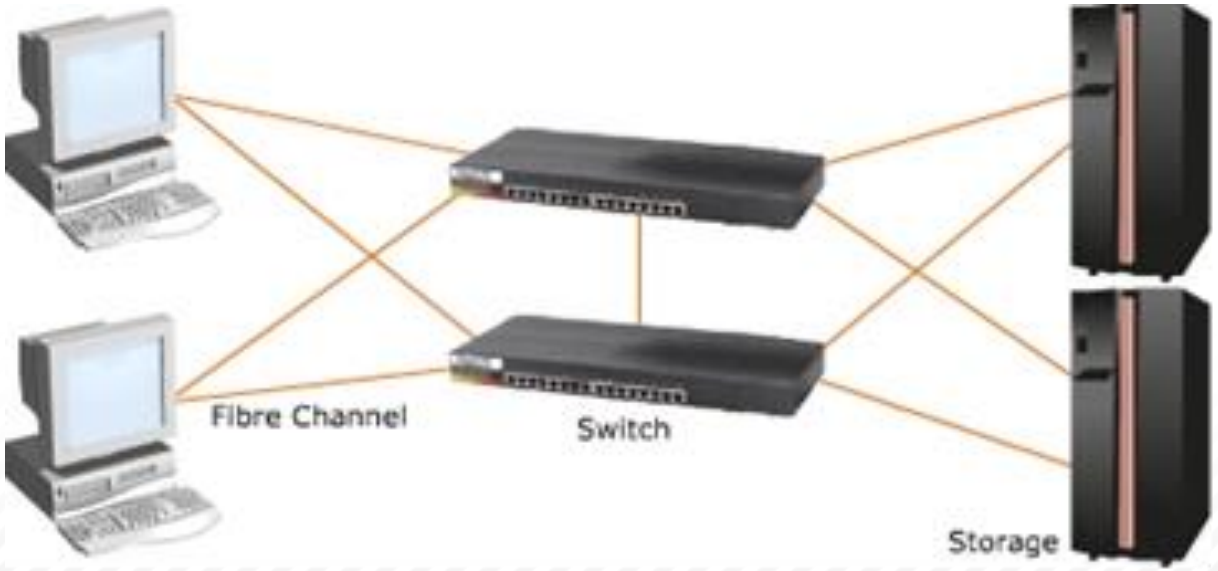

Storage Network Components – HBAs

- **Host Bus Adapter (HBA)** – In a storage system, an HBA is a Fibre Channel interface card installed in a server. It connects a computer and storage devices on a network.

  Each HBA has a unique WWN. The two types of WWNs on an HBA are these:

  - Node WWN: Shared by all ports on an HBA
  - Port WWN: Unique to each port on the HBA

Storage Network Components – WWNs

WWN address, World Wide Name, is a unique label, which identifies a particular device in Fiber Channel network

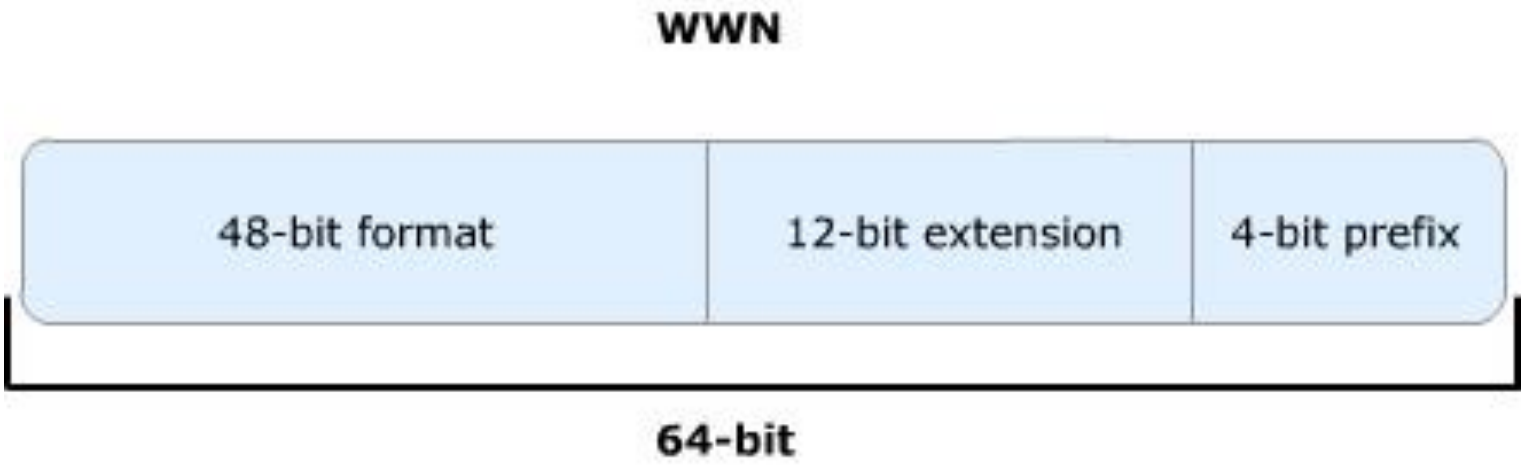

WWN example - 5 0 0 6 0 E 8 0 1 0 4 5 3 0 3 0 1 6

The Open Systems Interconnection (OSI) Model

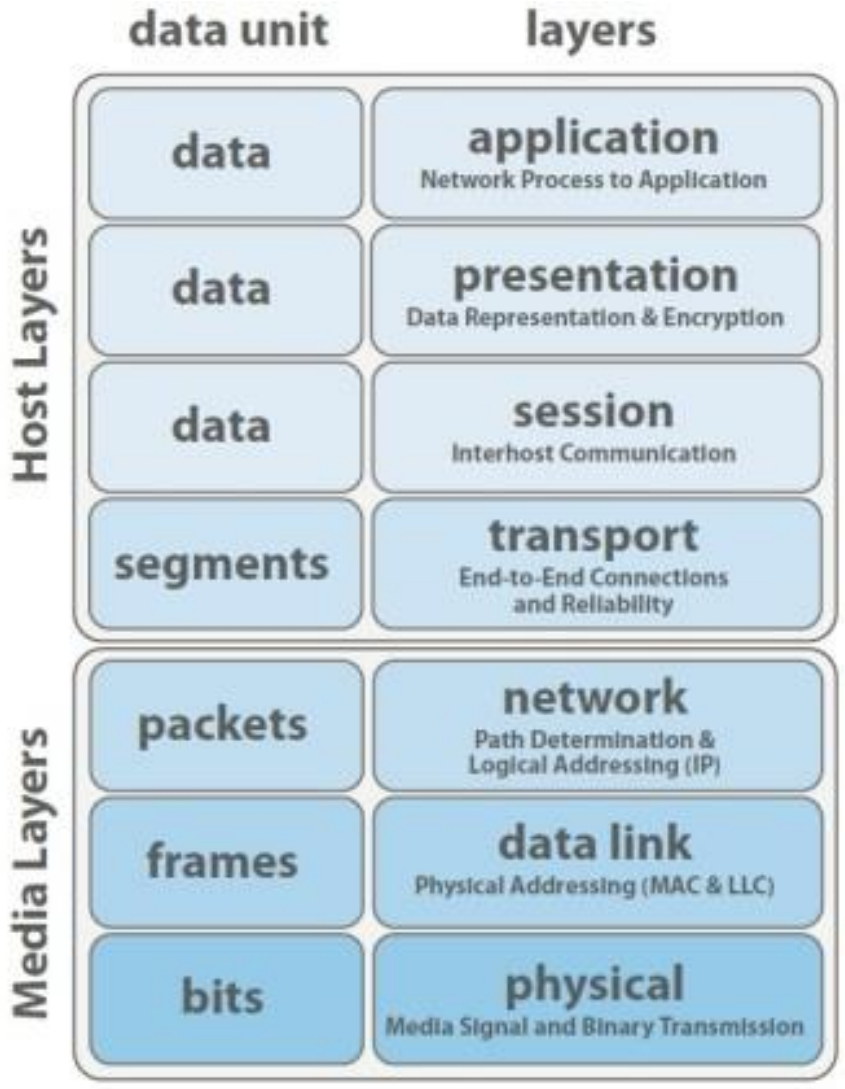

The Open Systems Interconnection (OSI) model is a conceptual model that characterizes and standardizes the internal functions of a communications system by partitioning it into abstraction layers.

Storage Network Components – HUB

- **Hub** – a simple device that allows interconnection and communication among nodes.

Storage Network Components – Fibre Channel Switches

- **Switch –** also allows interconnection and communication among nodes within one network

Storage Network Components – Routers

- **Router** – provides an interface between two different networks

Storage Network Components – Directors

- **Directors**

  A director is a large and complex switch. It is:

  - Highly available, reliable, scalable, and manageable
  - Fault tolerant with the ability to recover from a non-fatal error
  - Designed with redundant hardware components
  - Capable of supporting Fibre Channel and fiber connectivity (**FICON)**
  - Potentially expensive and complex
  - Designed for enterprises with large data centers

- Large networks often use Fibre Channel switches and directors in the same implementation.

Exercise: Storage Network Components

Match the storage network component with the appropriate description:

a. Node

b. Port

c. WWN

d. HBA

e. Cable

_____ 1. Connects and transmits signals between nodes

_____ 2. Transmits or receives data over a network

_____ 3. Fibre Channel interface card

_____ 4. Enables a node to communicate with another node

_____ 5. Unique number used to identify elements on a FC storage network

If vILT class, write your answers on blank lines

Storage Networking Topologies

SAN Topologies

- **Point-to-Point (FC-P2P) –** A point-to-point (P-P) topology is considered the simplest topology, in which two devices are directly connected using Fibre Channel.
  - It has fixed bandwidth; data is transmitted serially over a single cable.
  - It can be used with DAS.

SAN Topologies

- **Arbitrated Loop (AL)** – An FC topology where all devices are part of a loop and only one device can communicate with another device at a time

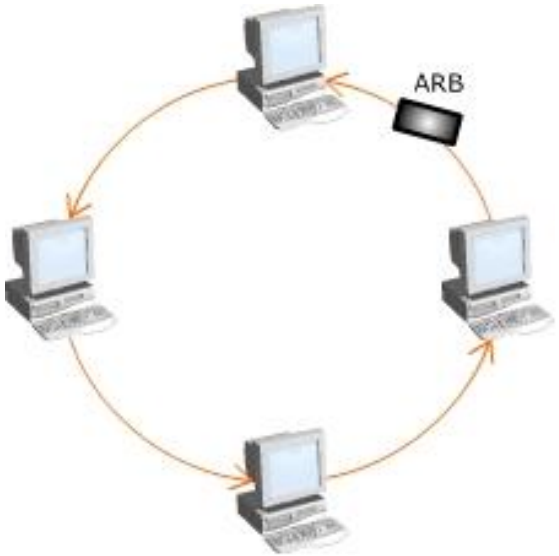

- In AL, devices use an access request mechanism called *arbitrate* (ARB), which circles the loop.
  - A device can use ARB depending on its priority and access rights.
  - The device with the highest priority gets first access.

SAN Topologies

- **Switched Fabric (FC-SW)** – A Fibre Channel topology that connects multiple devices by using Fibre Channel switches

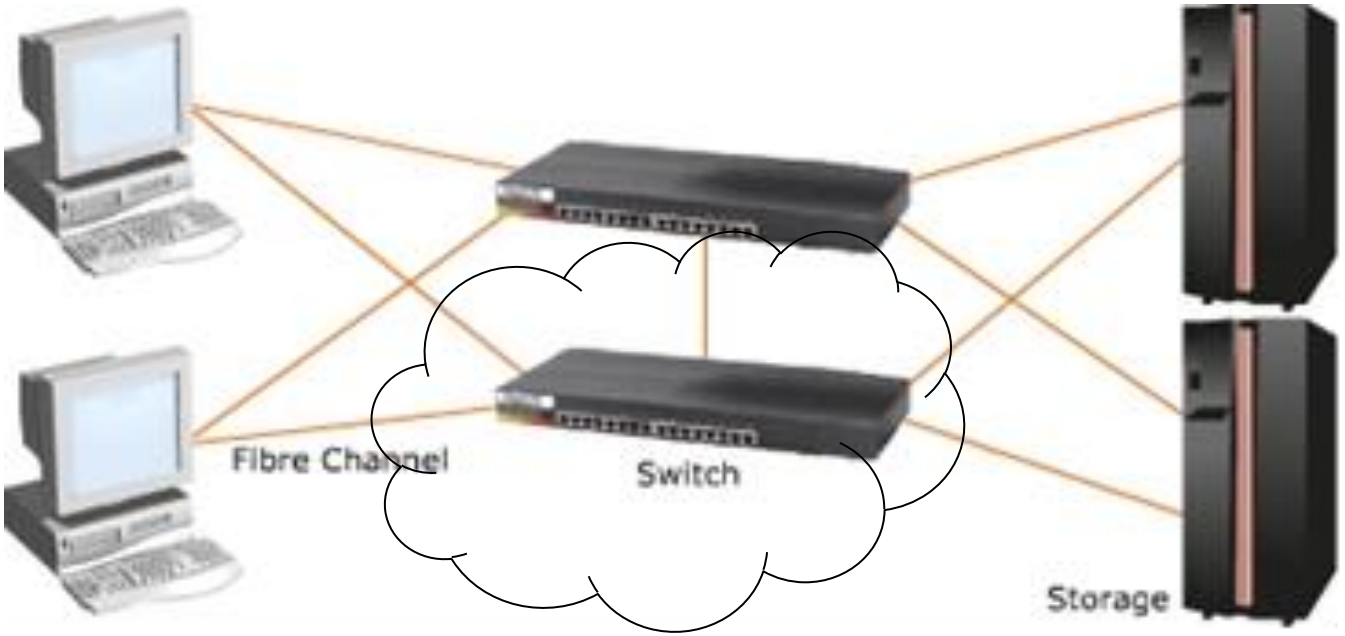

- In a switched fabric topology, bandwidth is not shared between devices, enabling devices to transmit and receive data at full speed at all times.

Direct Attached Storage

- Direct attached storage infrastructure. Server is directly connected to a storage system. Storage system can be accessed only through the server. Server can be accessed from Local Area Network (LAN).

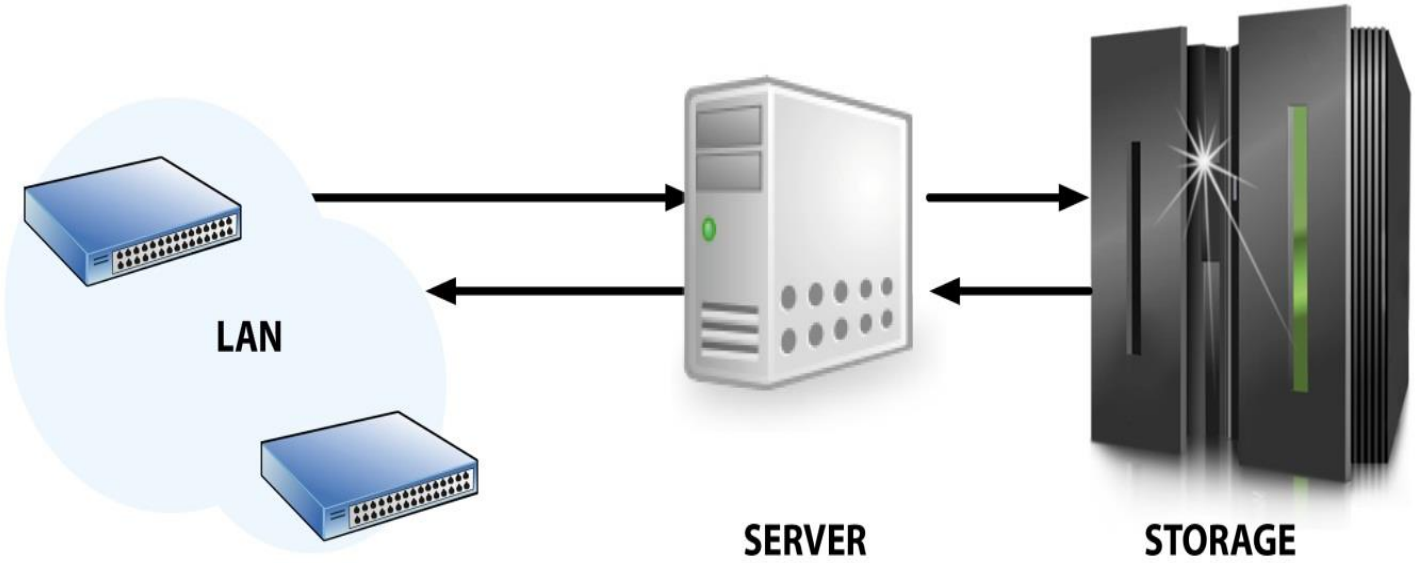

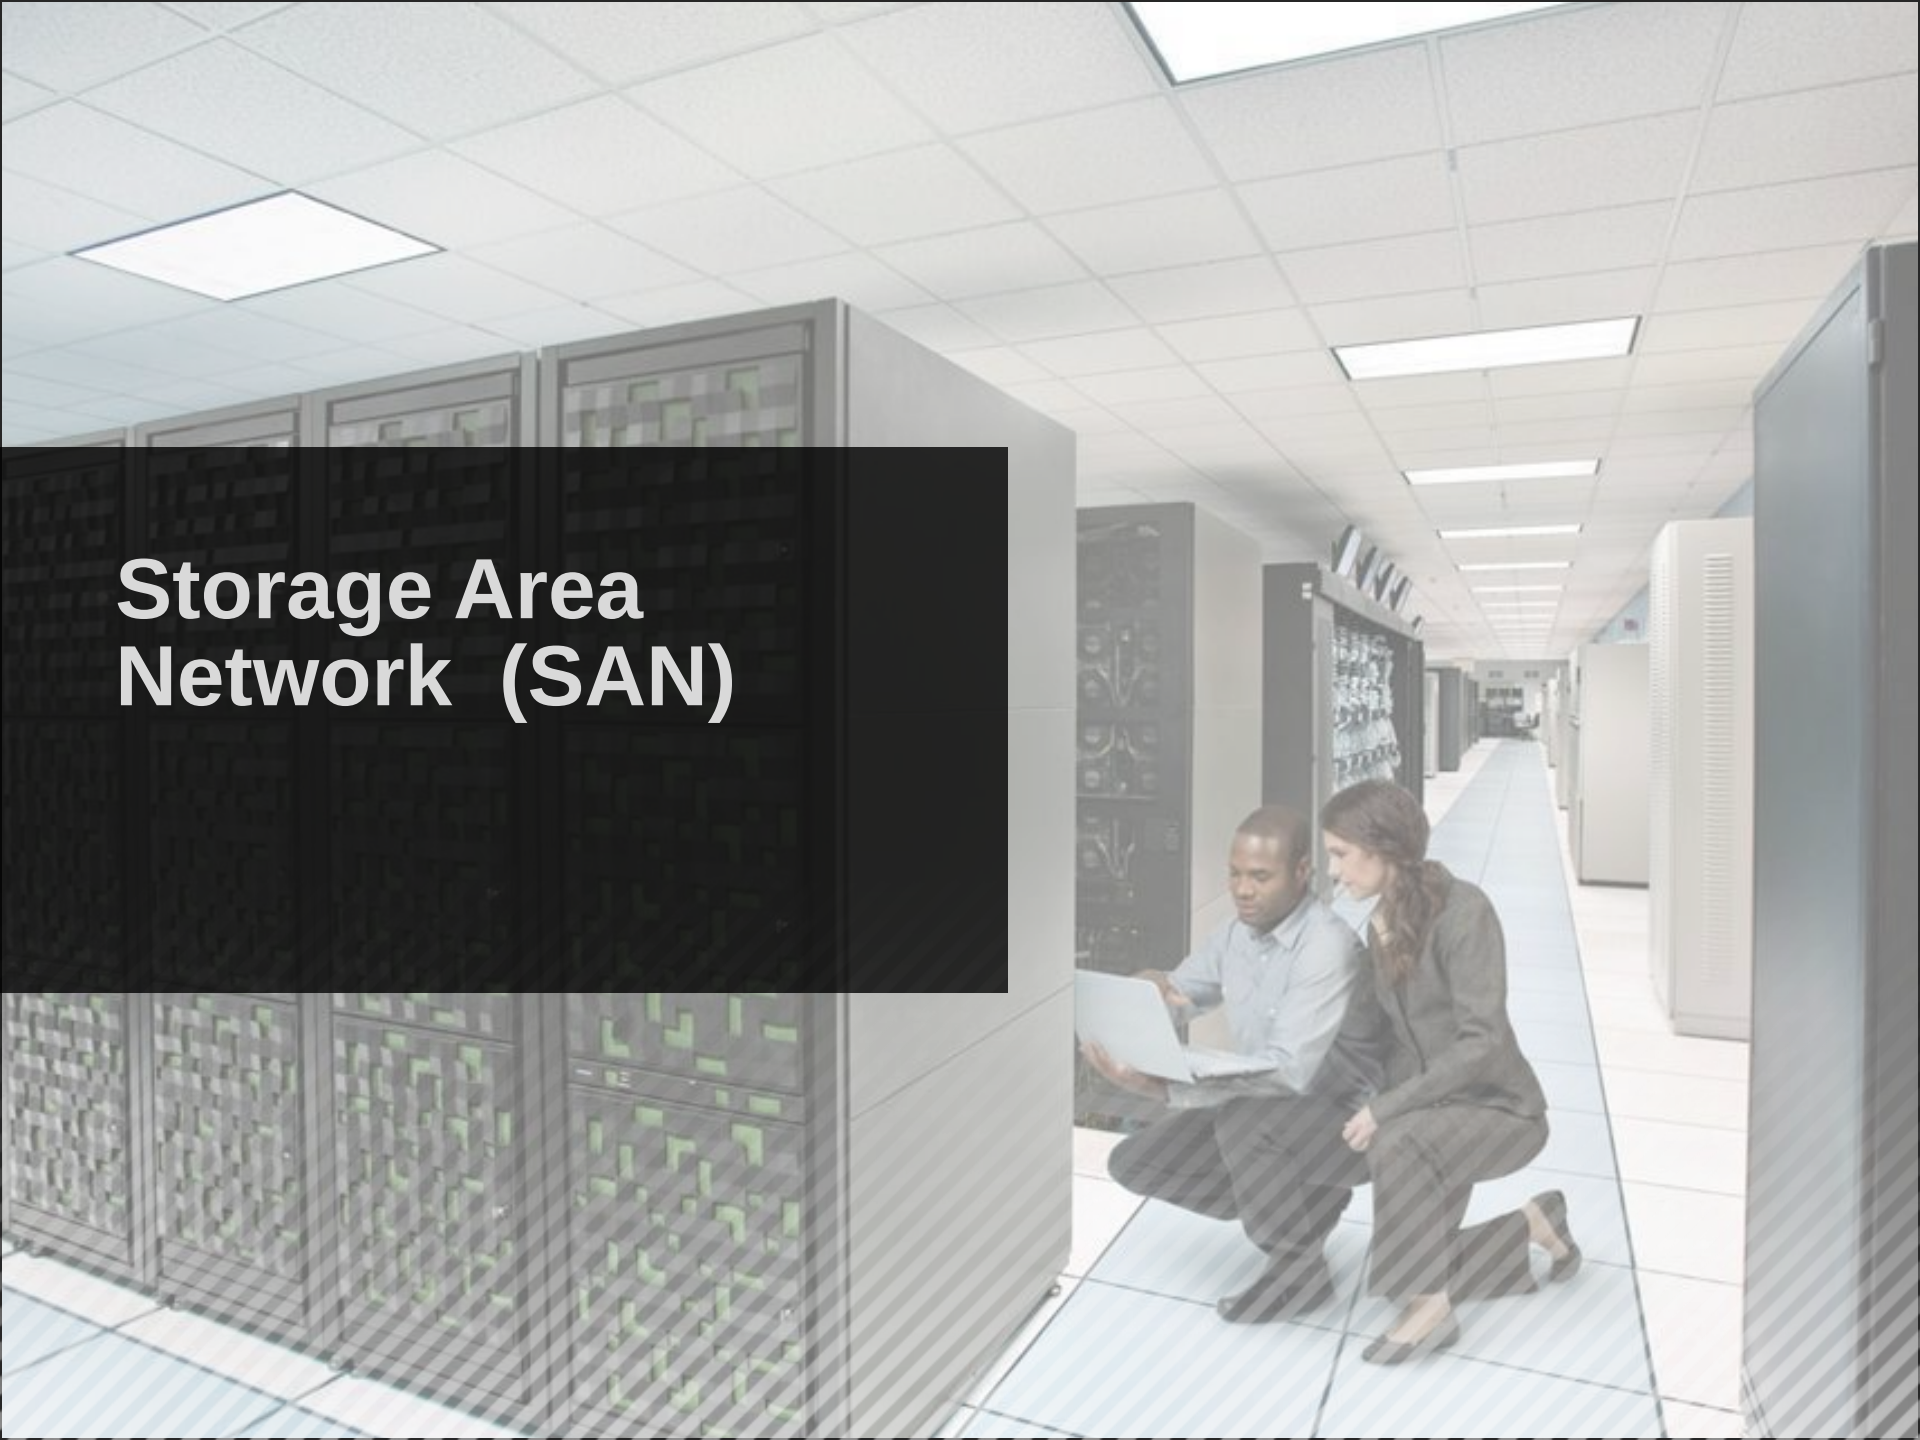

Storage Area Network (SAN)

Storage Area Network

- **Storage Area Network** (SAN) is a high-speed network of shared storage devices. Servers attached to a SAN can access any SAN attached storage devices.

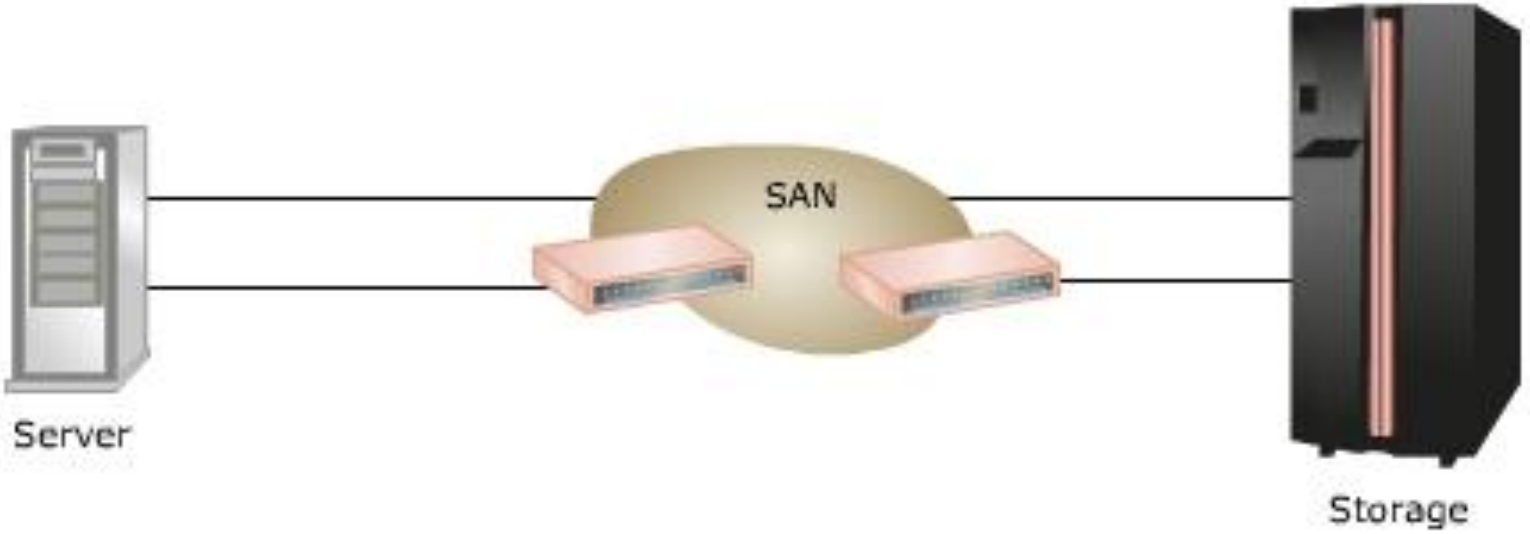

- The only components in a SAN are storage devices and switches. SAN is designed to connect computer storage devices, such as disk array controllers and tape libraries, to multiple servers, or hosts.

SAN Components

Storage Area Networks are using Fiber Channel infrastructure, which includes Host Bus Adaptors installed in servers, Fiber Channel cables and switches, Fiber Channel ports installed in the storage system front end and proprietary network protocols.

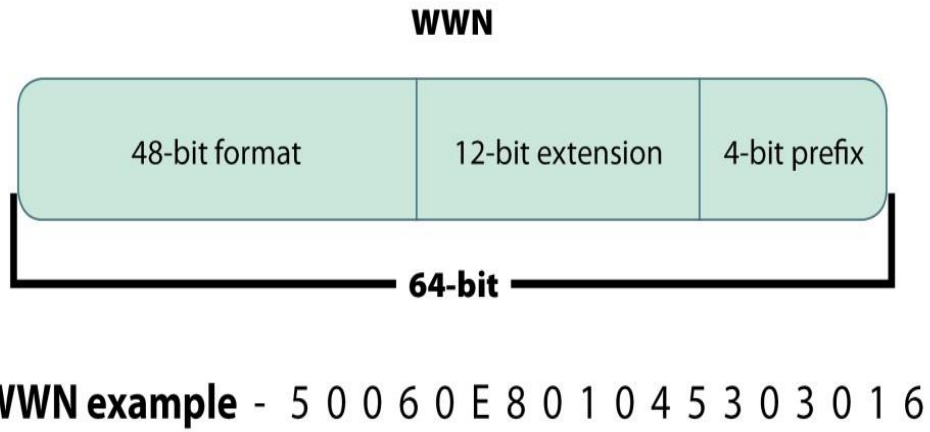

Host Bus Adaptor (HBA) and an example of WWN number.

SAN over iSCSI Interface

- **Internet Small Computer Systems Interface (iSCSI) –** SAN can be implemented by using a network protocol standard called iSCSI which uses the SCSI protocol to transmit data over TCP/IP networks.
- iSCSI allows organizations to use their existing TCP/IP network infrastructure without investing in expensive Fibre Channel switches.

iSCSI HBA

Network Attached Storage

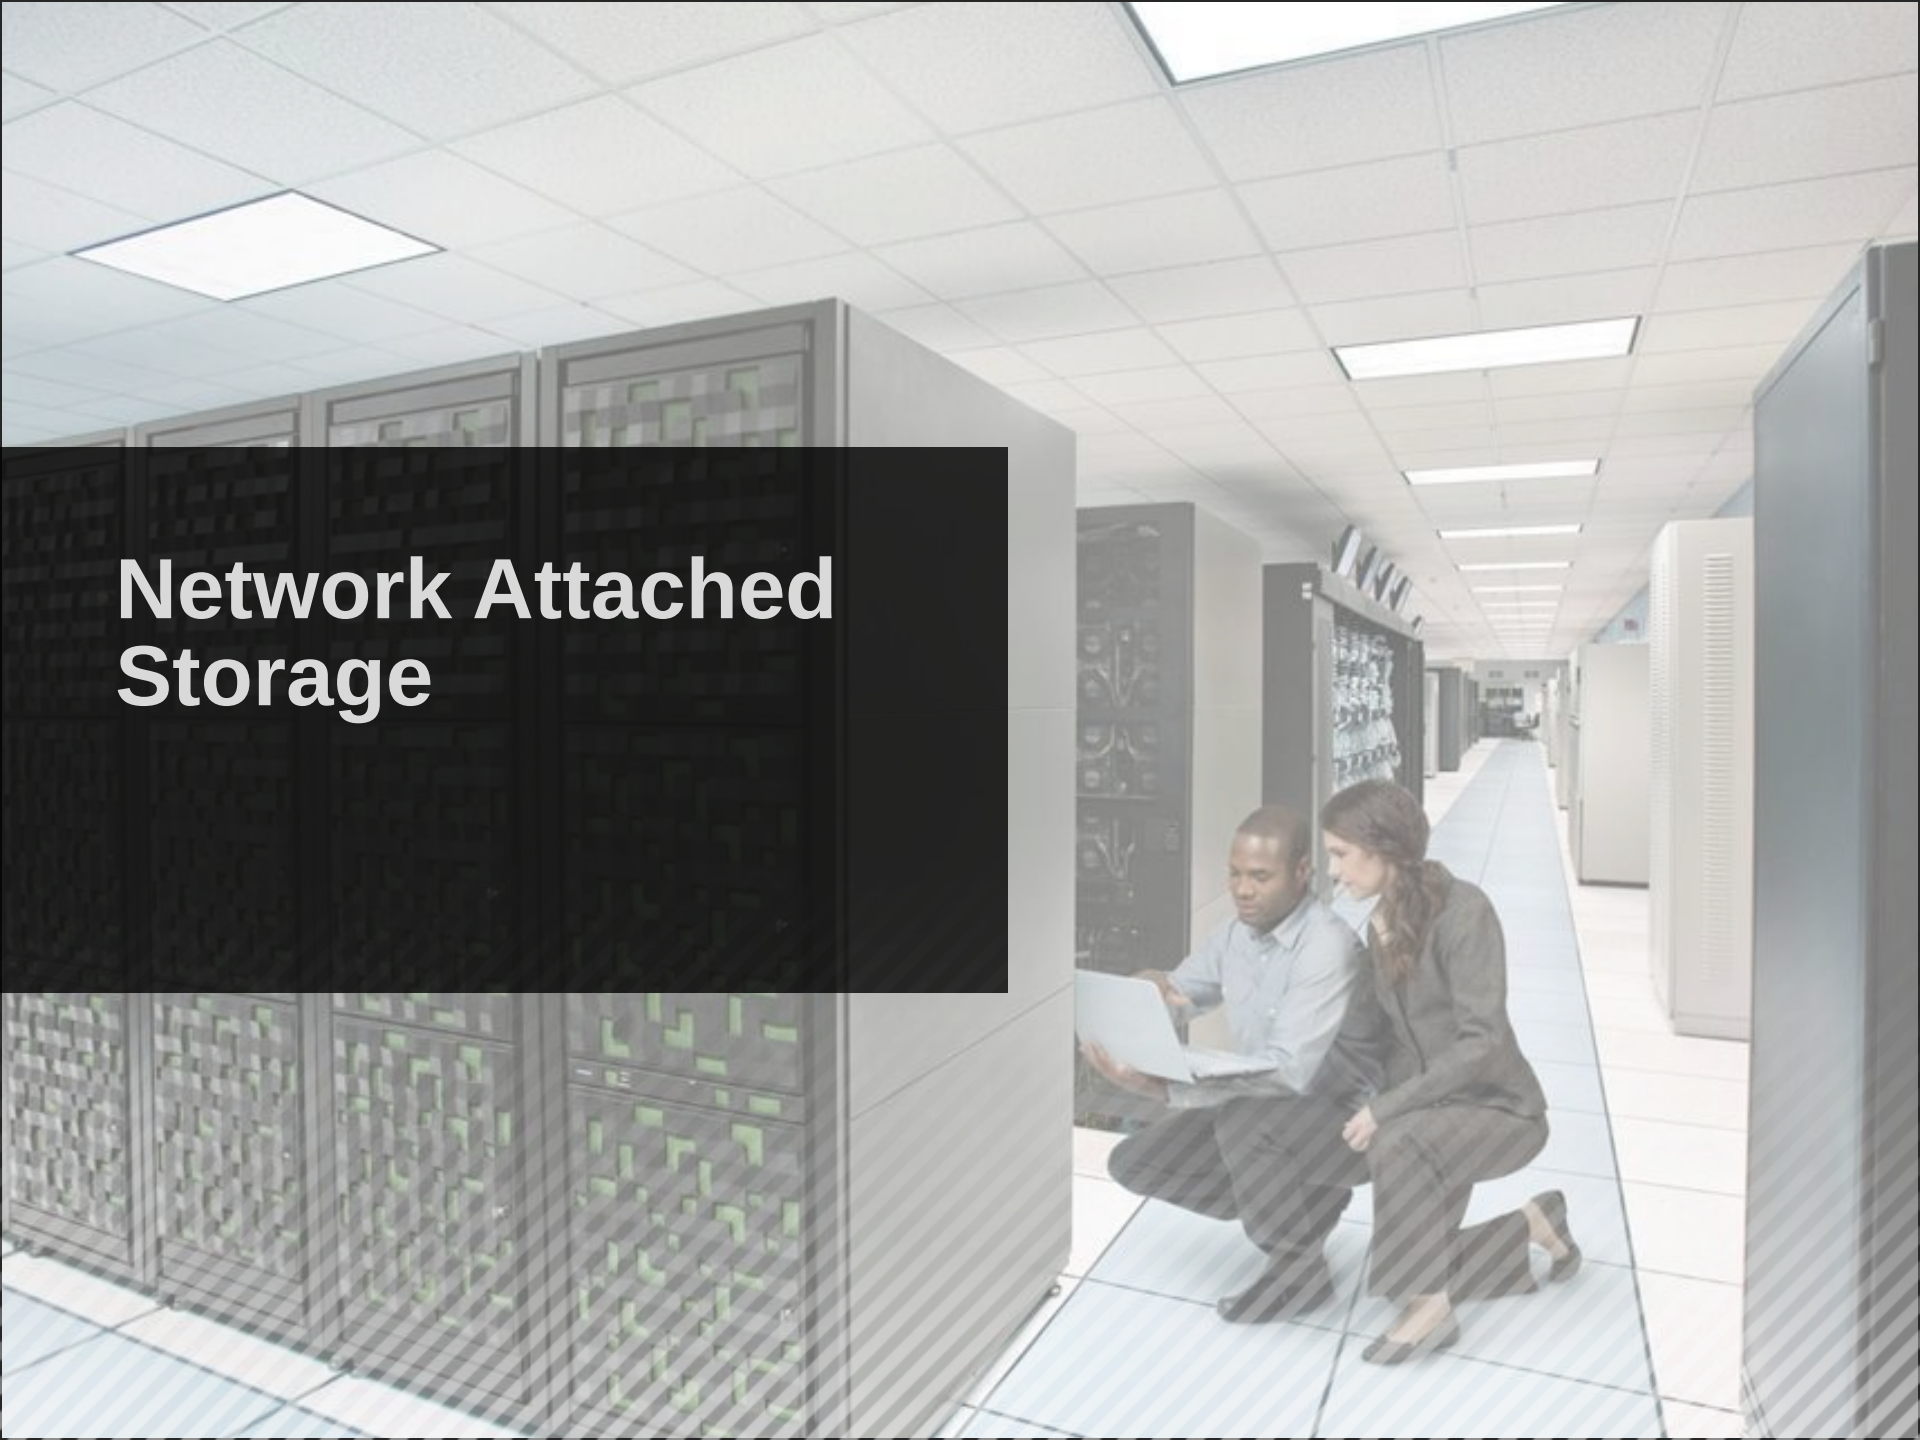

Network Attached Storage

Network Attached Storage (NAS) is represented by the server that functions as the NAS Head and common storage system. There are solutions that integrate both these functionalities in one package (NAS Appliances). NAS devices work relatively independently; they do not require servers with applications. All clients, application and other servers can access **files** stored in a NAS device.

File Access Protocols supported include:

- CIFS
- NFS,
- HTTP
- HTTPS
- FTP

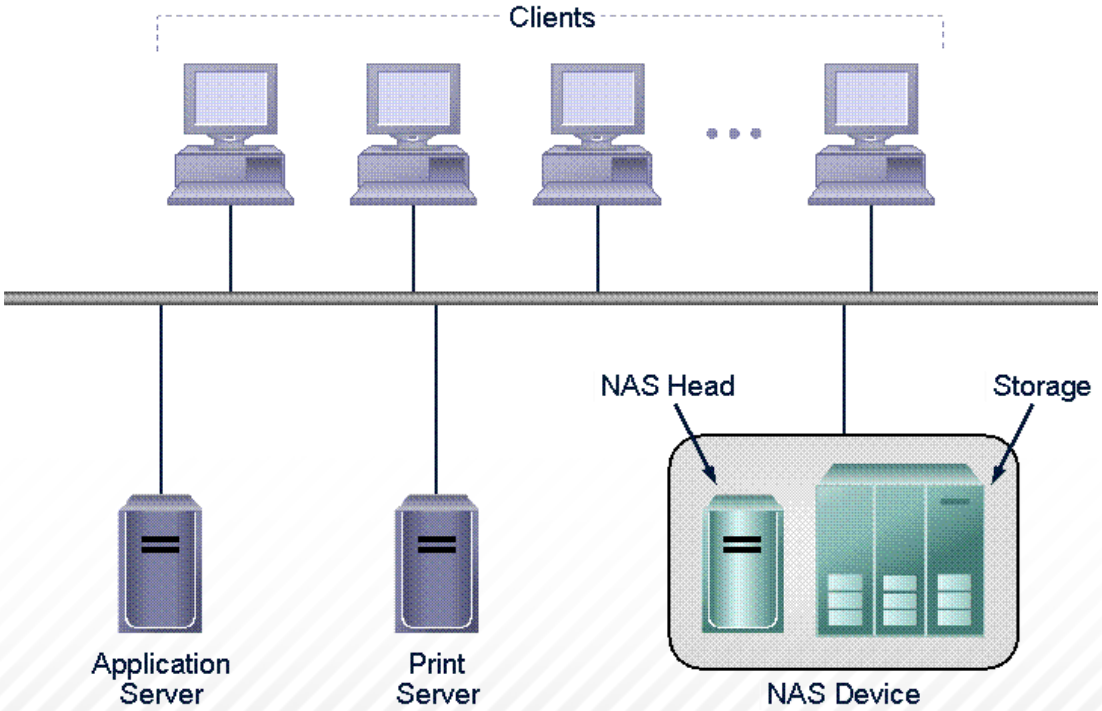

Network Attached Storage

- In **NAS**, the storage device is directly connected to the LAN and there is no server between the data and other network devices.
- Data is presented to the server at file level.

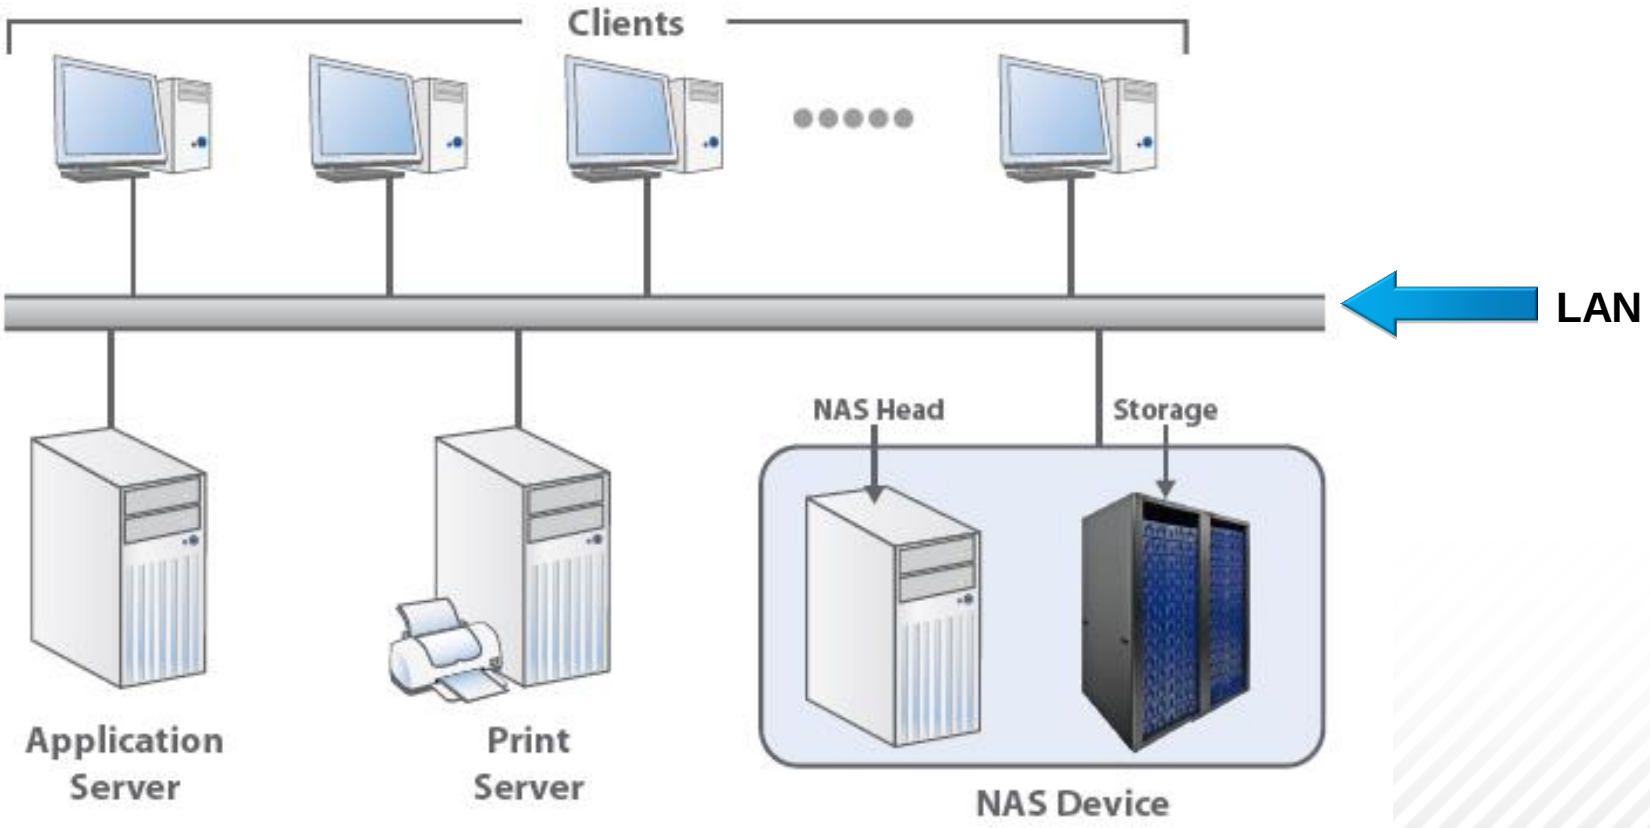

Network Attached Storage

- **Advantages of NAS**
  - Offers storage to different open-systems operating systems over LAN
  - Data is presented to servers at file level, reducing server overhead
  - Dedicated file server, optimized for sharing files between many users
  - Minimizes overhead by centrally managing storage
  - Facilitates easy and inexpensive implementation

- **Disadvantages of NAS**
  - Relies on the client-server model for communication and data transport which creates network overhead
  - Lower performance than a SAN

NAS Implementations

- **Methods of NAS Implementation –** An organization can implement NAS architecture by following methods:
  - NAS appliance or filer
  - NAS blade
  - NAS gateway

NAS Appliance

- **NAS Appliance –** Combines a front-end file server and back-end storage system in a single unit. This approach is called a closed-box approach.

- NAS appliance has the following advantages:
  - Combines a file server with the storage array
  - Provides efficient performance
  - Has high reliability
  - Enables easy installation, management, and use
  - Provides the least expensive NAS implementation

- NAS appliance has the following disadvantages:
  - Is not scalable
  - No pool storage, which makes it hard to achieve high utilization.

NAS Blade

- NAS Blade – Allows multi-protocol data storage in a large disk array

- NAS blade has following advantages in addition to a NAS appliance:
  - Is scalable
  - Provides backup of storage data
  - Supports multiple NAS blades

NAS Gateway

- **NAS Gateway** – All devices communicate directly with the file system.

  A NAS gateway overcomes the limitations of a NAS appliance.

- It has the following advantages:
  - Separates file server from storage device
  - Is less expensive than a NAS appliance
  - Supports multiple NAS gateways
  - Has better utilization rates
  - Combines NAS with SAN capacity to meet growing storage requirements
  - Provides NAS functionalities to SAN storage

NAS Gateway

- NAS gateway controller uses FC protocol to connect to external storage.

Converged Solution – SAN and NAS Together

- Converged solution – SAN and NAS together
  - NAS head with storage over the SAN
  - NAS scales to the limits of the SAN
    - Limited by NAS file system's capacity
- Co-exists with application servers
- Centrally managed

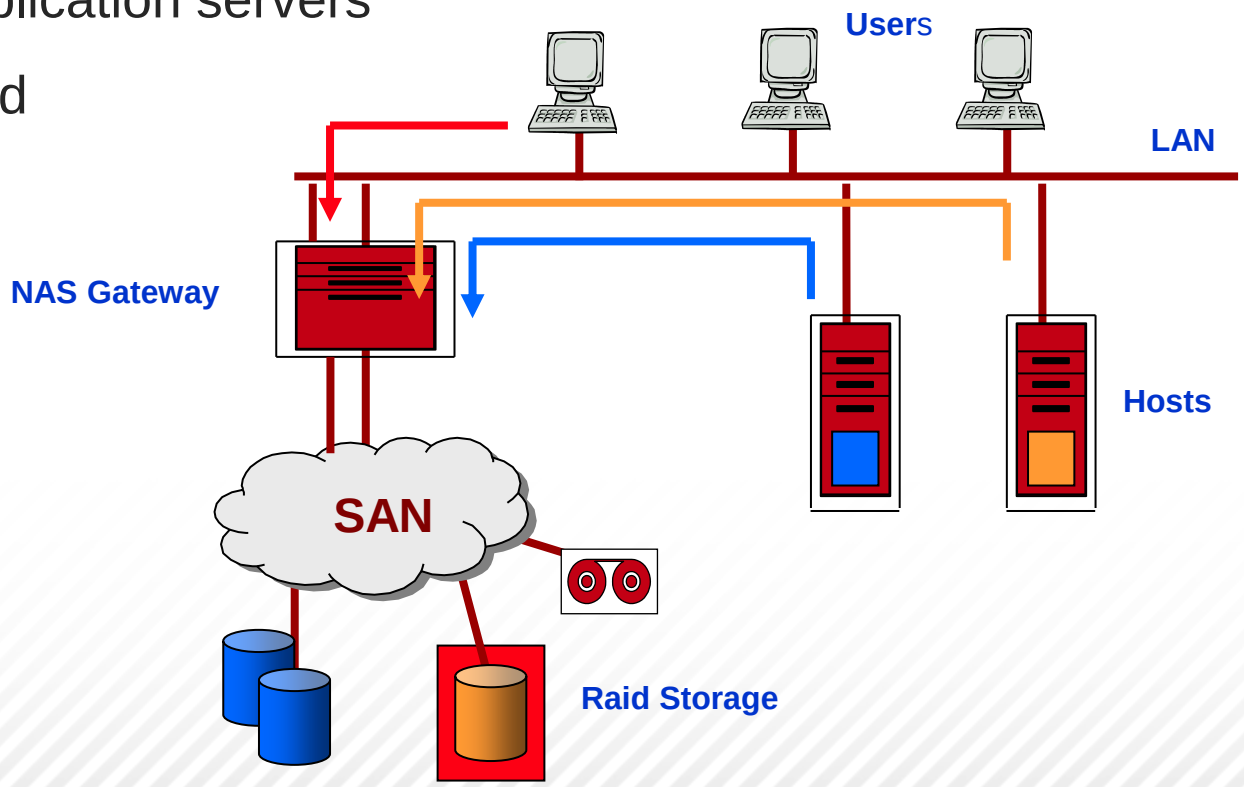

Storage Networking Architectures…Side by Side

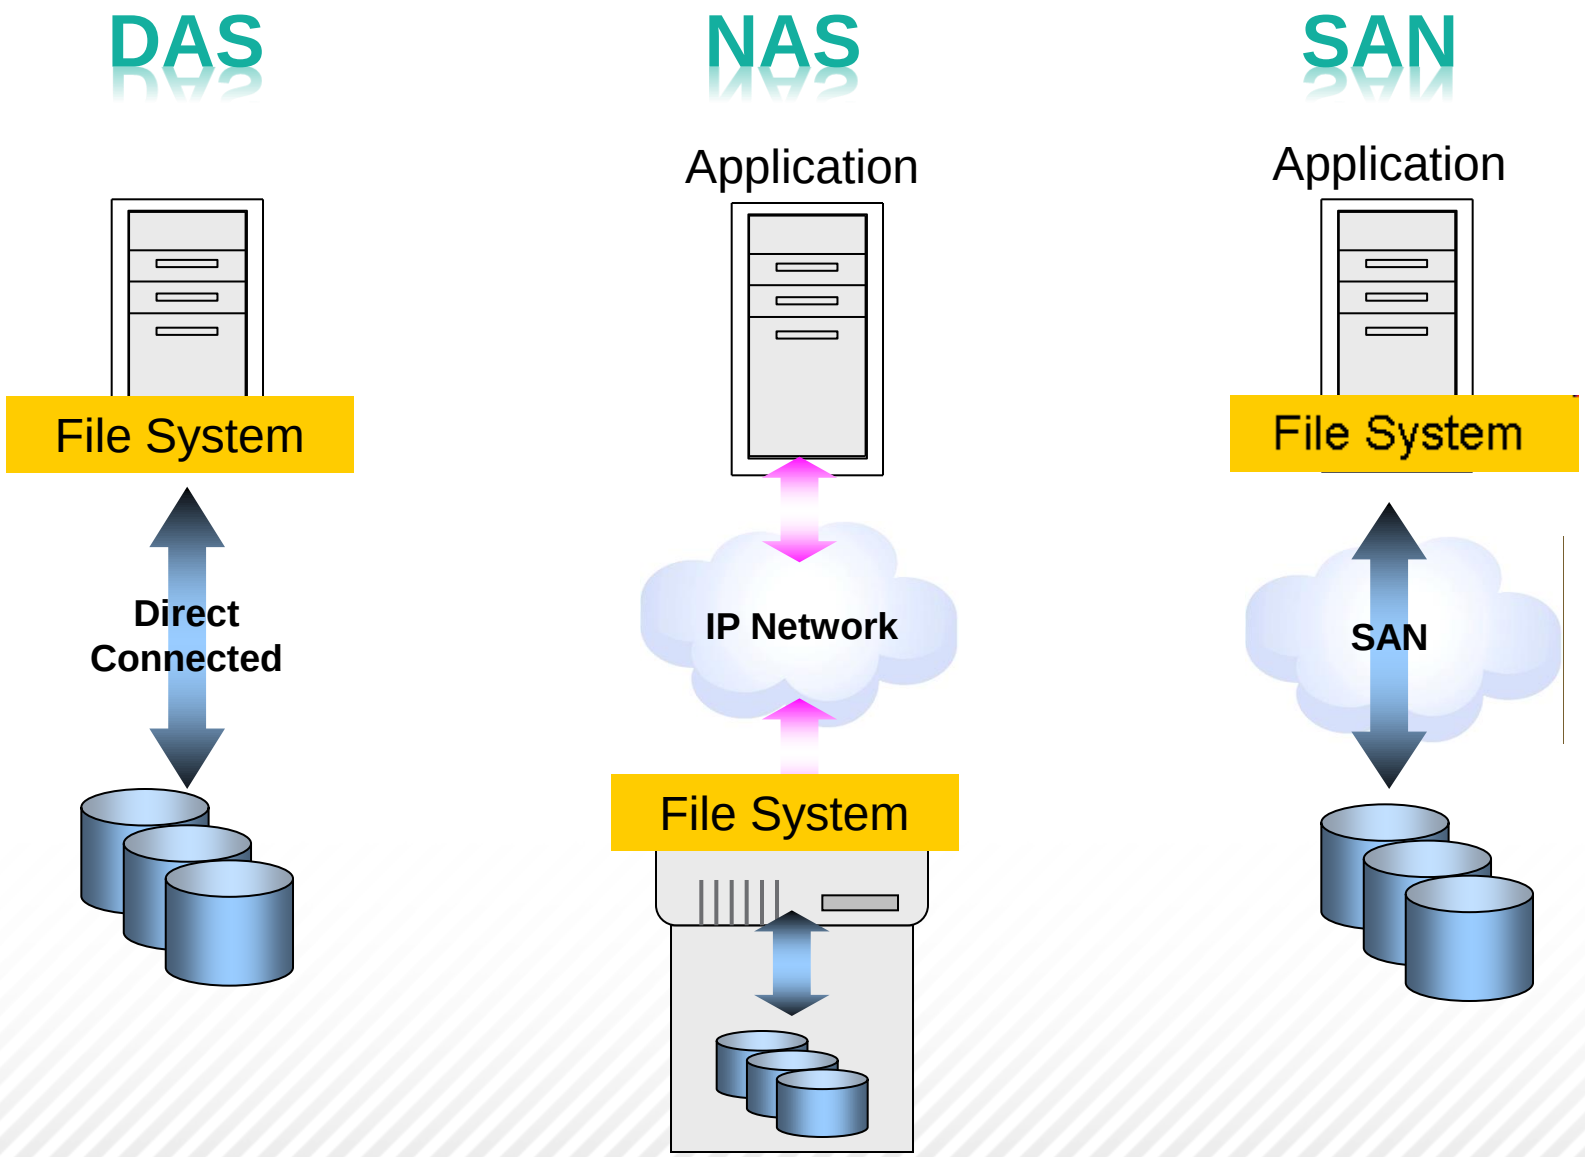

Exercise: Storage Networking Concepts

1. A DAS device is not shared, so no other network device can access the data without first accessing the server. True or False?

2. Select the best description for the Fibre Channel topology known as FC-AL.

   a) Two devices, data transmitted serially over a single cable

   b) Multiple devices connected in a loop, highest priority device gets first access

   c) Multiple devices connected using Fibre Channel switch, devices transmit and receive data at full speed at all times

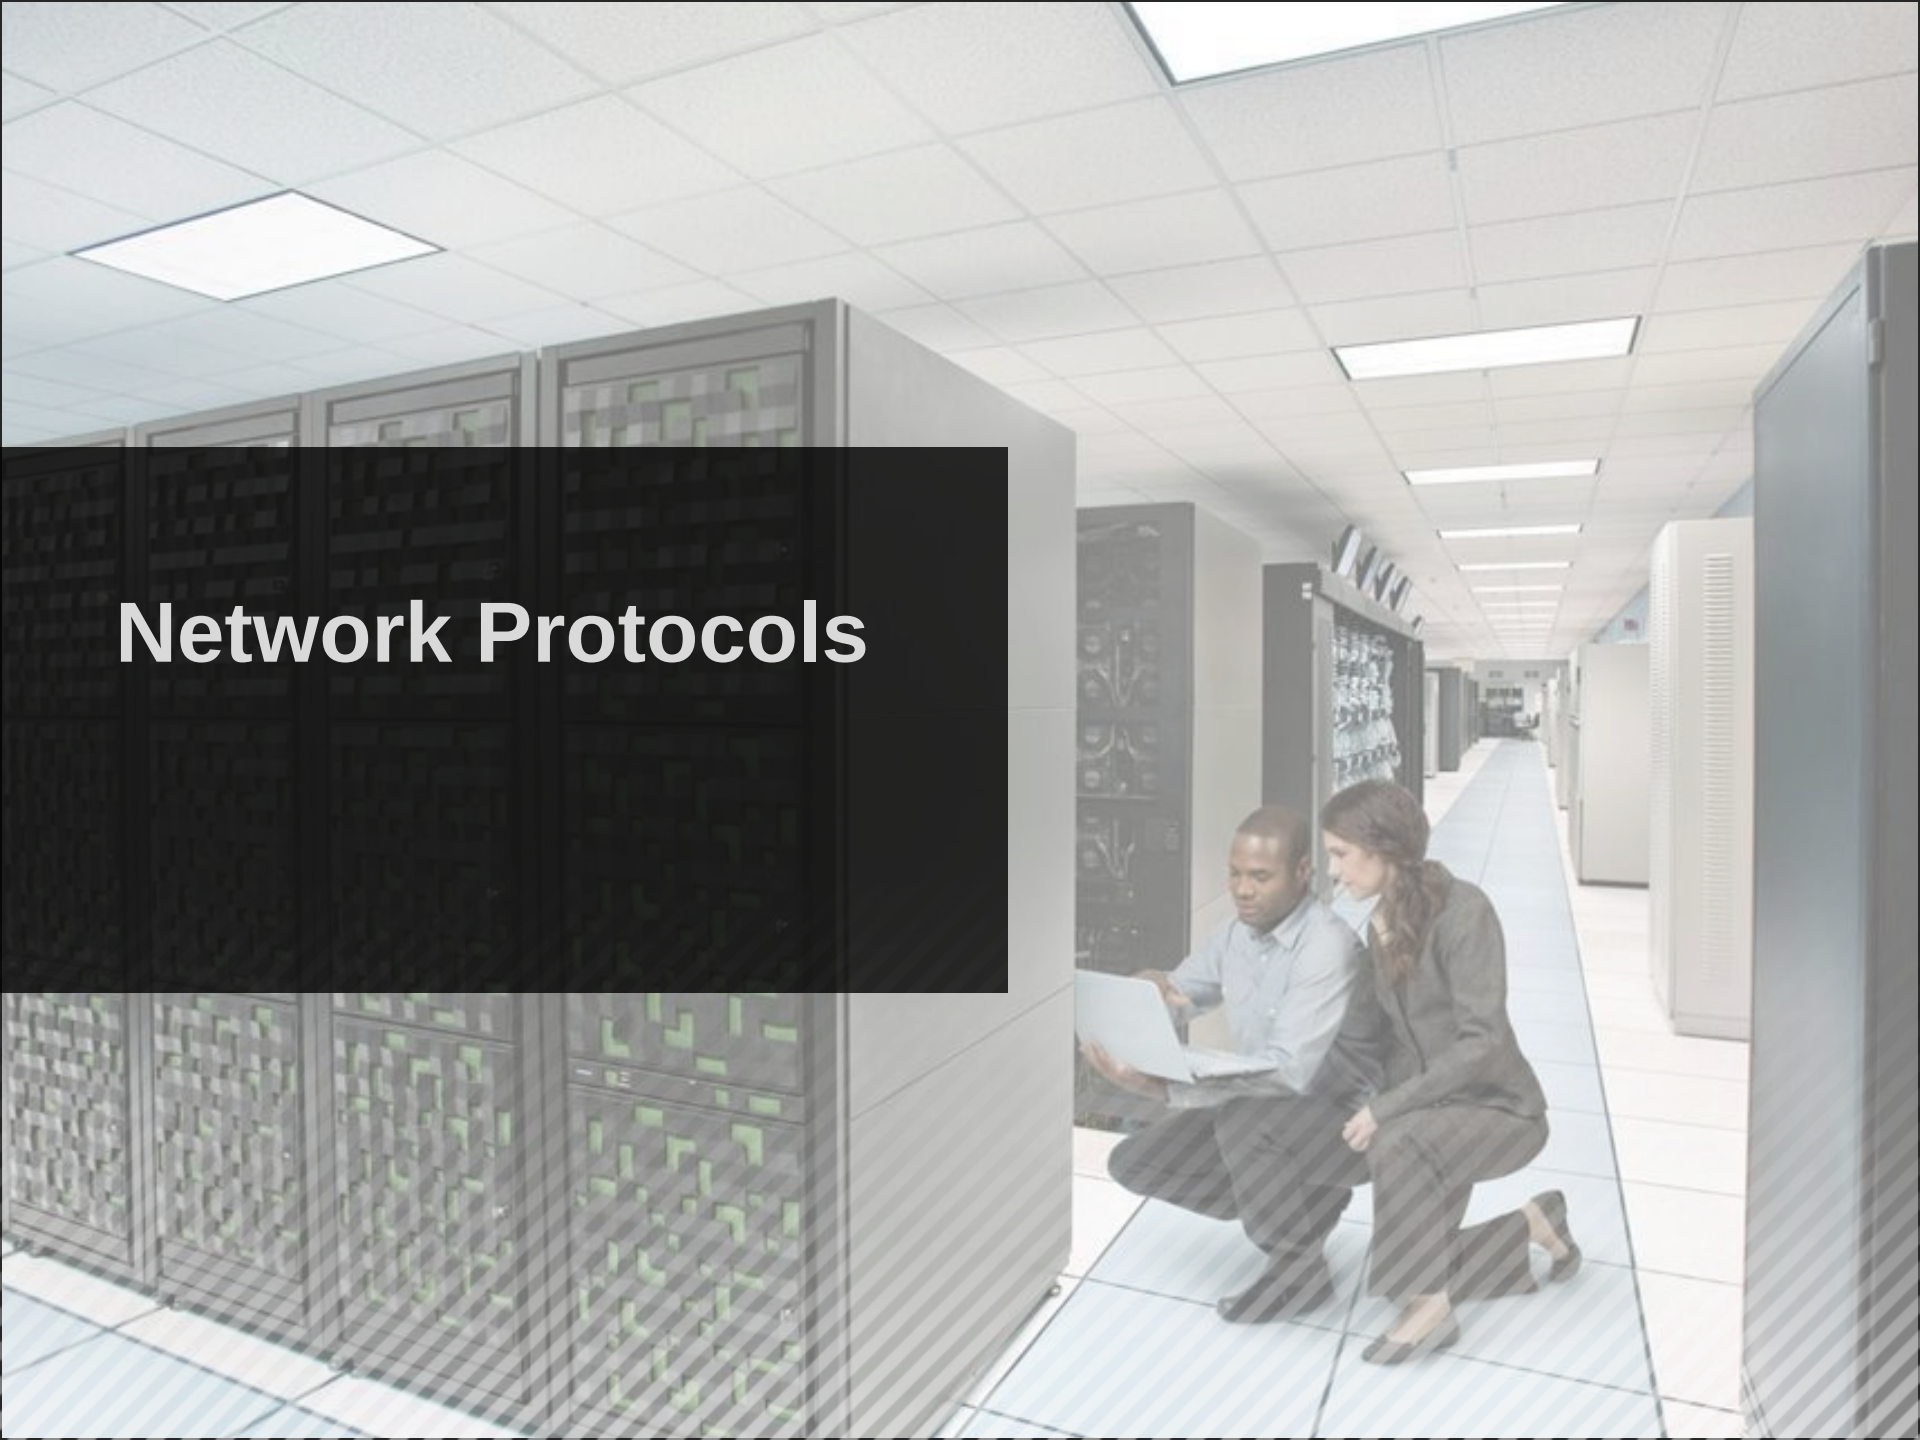
Network Protocols

Protocol

- **Protocol** – A set of rules that govern communication between computers on a network. It regulates the following characteristics of a network:

  - Access method
  - Physical topologies allowed in the network
  - Types of cable that can be used in the network
  - Speed of data transfer

Protocols

The different types of protocols that can be used in a network are:

- Ethernet
- Fibre Channel protocol (FCP)
- Fiber connectivity (FICON)
- Internet protocol (IP)
- Internet small computer system interface (iSCSI)
- Fibre Channel over IP (FCIP)
- Internet Fibre Channel protocol (iFCP)
- Fibre Channel over Ethernet (FCoE)

Protocols

Ethernet

Uses an access method called carrier sense multiple access/collision detection (CSM/CD). Before transmitting, a node checks whether any other node is using the network. If clear, the node begins to transmit. Ethernet allows data transmission over twisted pair or fiber optic cables and is mainly used in LANs. There are various versions of Ethernet with various speed specifications.

FCP

Defines a multi-layered architecture for moving data. FCP packages SCSI commands into Fibre Channel frames ready for transmission. FCP also allows data transmission over twisted pair and over fiber optic cables. It is mainly used in large data centers for applications requiring high availability, such as transaction processing and databases.

FICON

Connects a mainframe to its peripheral devices and disk array. Ficon is based on FCP and has evolved from the older ESCON protocol.

Protocols

IP/TCP

IP is used to transfer data across a network. Each device on the network has a unique IP address that identifies it. IP works in conjunction with the TCP, iSCSI and FCIP protocols. When you transfer messages over a network by using IP, IP breaks the message into smaller units called packets (third layer in OSI model). Each packet is treated as an individual unit. IP delivers the packets to the destination. **TCP** Is the protocol that combines the packets into the correct order to reform the message that was sent from the source.

iSCSI

Establishes and manages connection between IP-based storage devices, and it hosts and enables deployment of IP-based storage area networks. It facilitates data transfers over intranets, manages storage over long distances and is cost-effective, robust and reliable. iSCSI is best-suited for web server, email and departmental business applications in small to medium sized businesses.

Protocols

FCIP

Fibre Channel over IP is a TCP/IP based tunnelling protocol that connects geographically distributed Fibre Channel SANs. FCIP encapsulates Fibre Channel frames into frames that comply with TCP/IP standards. It can be useful when connecting two SAN networks over the Internet tunnel, in a similar fashion to virtual private networks (VPNs) allowing connection to a distant LAN over the Internet.

iFCP

iFCP is again TCP/IP based. It is basically an adaptation of FCIP using routing instead of tunneling. It interconnects Fibre Channel storage devices or SANs by using an IP infrastructure. iFCP moves Fibre Channel data over IP networks by using iSCSI protocols.

Both FCIP and iFCP provide means to extend Fibre Channel networks over distance. Both these protocols are highly reliable and scalable. They are best suited for connecting two data centers for centralized data management or disaster recovery.

FCoE

Fibre Channel over Ethernet is an encapsulation of Fibre Channel frames over Ethernet networks. This allows Fibre Channel to use 10Gb Ethernet networks while preserving the Fibre Channel protocol. FCoE provides these advantages:

- Network (IP) and storage (SAN) data traffic can be consolidated using a single network switch.
- It reduces the number of network interface cards required to connect disparate storage and IP networks.
- Reduces the number of cables and switches.
- Reduces power and cooling costs.

Thus, you can build your SAN using Ethernet cables (mostly twisted pair). You can use one switch for your IP-based network traffic (LAN) and for creating SAN infrastructure. Even though the switch and cabling are the same, LAN will run on TCP/IP while SAN runs on FCP.

Exercise: Storage Networking

Match the following list of components with the appropriate definition:

a. Client-server

b. Protocol

c. LAN

d. WAN

_____ 1. Network connecting devices in a small geographic area

_____ 2. Relationship between two computers – one sends requests; one responds with data

_____ 3. Set of rules governing communication among computers on a network

_____ 4. Network connecting devices across larger geographical areas

If vILT class, write your answers on blank lines

Storage Area Network Security

LUN Mapping

- A LUN is a logical device mapped to a storage port.

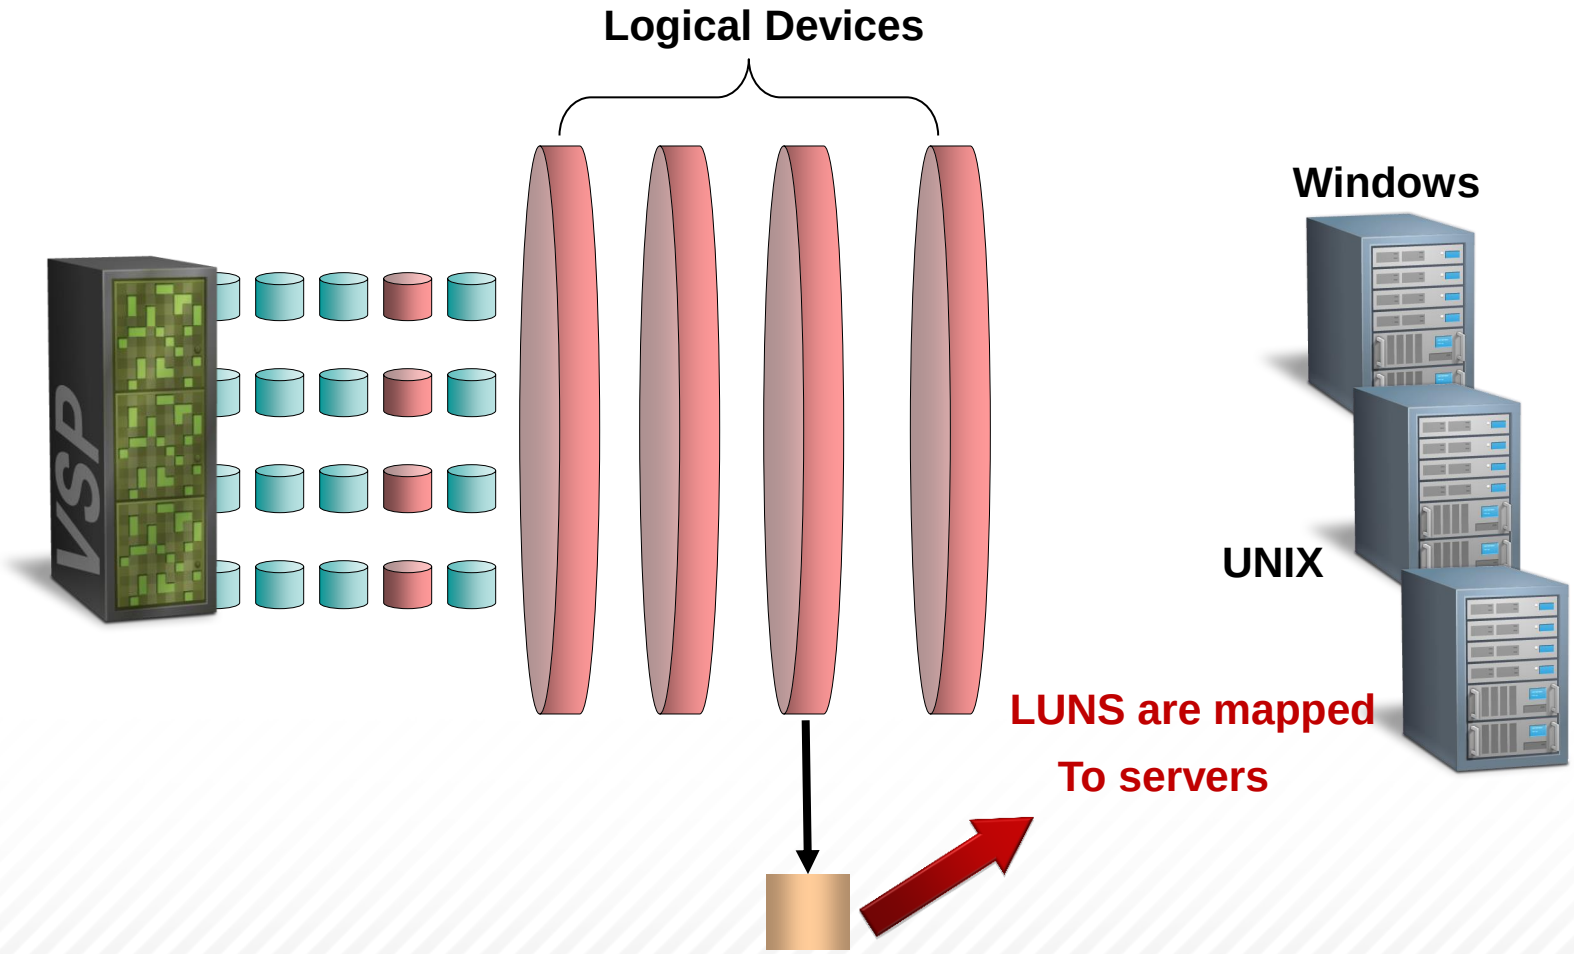

Zoning

- Zones – Defined to establish rules governing communication of network devices
  - WWN Zoning (Soft Zoning)
  - Port Based Zoning (Hard Zoning)
  - Mixed Zoning

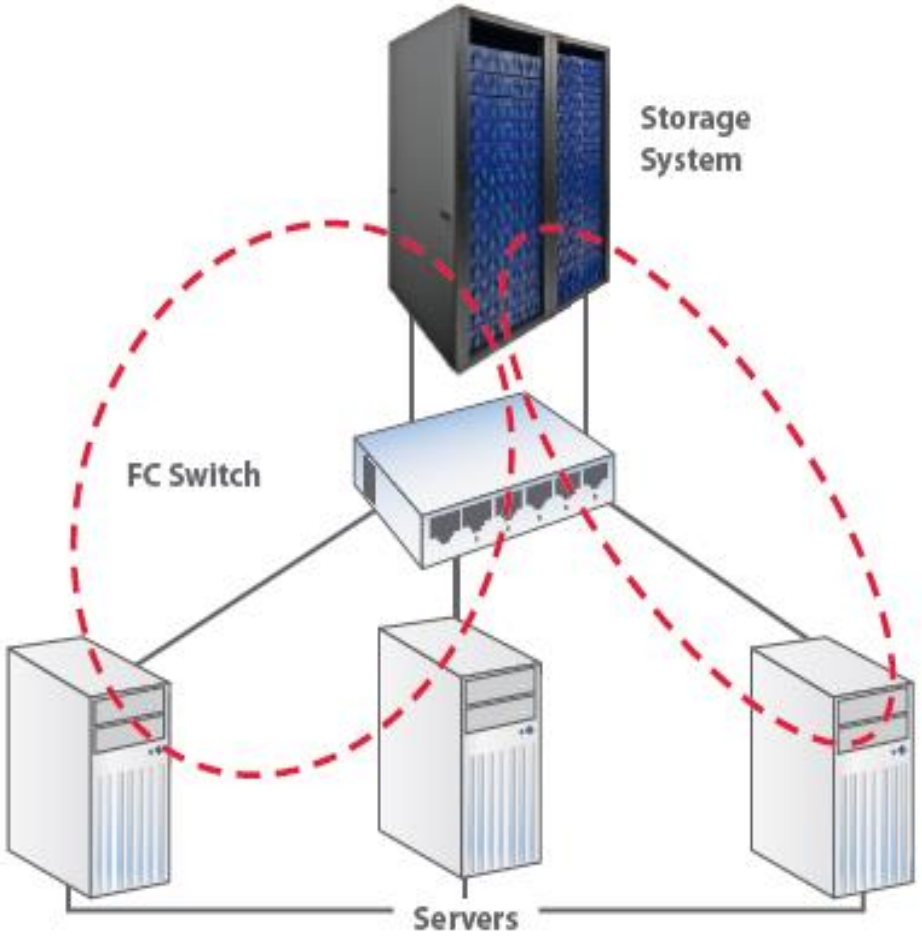

An example of WWN based zoning

Module Summary

- Upon completion of this module, you should have learned to:
  - Describe basic networking concepts
  - Explain how common network devices operate
  - Explain how devices communicate in a network
  - Explain storage area network security