Unit – 03

Introduction to Networked Storage

Unit - 03/Lecture - 01

Das(direct attached storage) [Rgpv/dec2014(2),Rgpv/dec2013 (7),Rgpv/dec2012(10)]

Das is an architecture where storage connects directly to servers. Applications access data from Das using block-level access protocols. The internal HDD of a host, tape libraries, and directly connected external HDD packs are some examples of das.

Types of Das

Das is classified as internal or external, based on the location of the storage device with respect to the host.

Internal das

In internal das architectures, the storage device is internally connected to the host by a serial or parallel bus. The physical bus has distance limitations and can only be sustained over a shorter distance for high-speed connectivity. In addition, most internal buses can support only a limited number of devices, and they occupy a large amount of space inside the host, making maintenance of other components difficult.

External das

In external das architectures, the server connects directly to the external storage device . In most cases, communication between the host and the storage device takes place over scsi or FC protocol. Compared to internal das, an external das overcomes the distance and device count limitations and provides centralized management of storage devices.

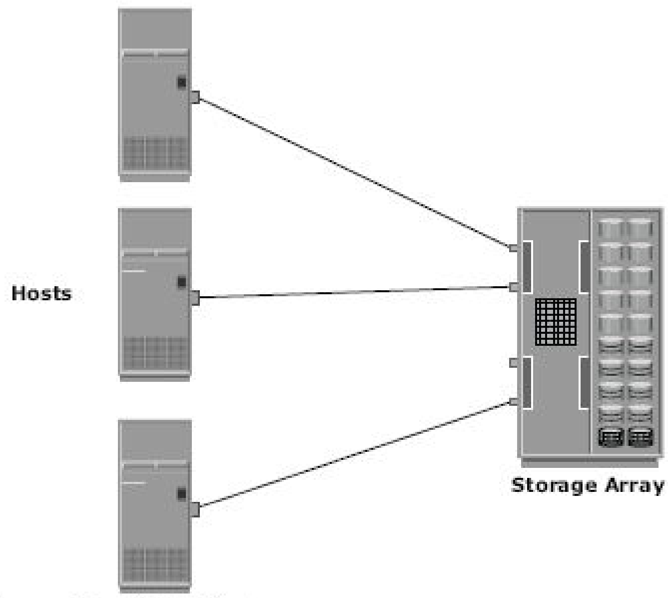

Das benefits and limitations

Das requires a relatively lower initial investment than storage networking. Storage networking architectures are discussed later in this book. Das con- figuration is simple and can be deployed easily and rapidly. Setup is managed using host-based tools, such as the host os, which makes storage management tasks easy for small and medium enterprises. Das is the simplest solution when compared to other storage networking models and requires fewer management tasks, and less hardware and software elements to set up and operate.

However, das does not scale well. A storage device has a limited number of ports, which restricts the number of hosts that can directly connect to the storage. A limited bandwidth in das restricts the available i/o processing capability. When capacities are being reached, the service availability may be com- promised, and this has a ripple effect on the performance of all hosts attached to that specific device or array. The distance limitations associated with implementing das because of direct connectivity requirements can be addressed by using fibre channel connectivity. Das does not make optimal use of resources due to its limited ability to share front end ports. In das environments, unused resources cannot be easily re-allocated, resulting in islands of over-utilized and under-utilized storage pools.

Disk utilization, throughput, and cache memory of a storage device, along with virtual memory of a host govern the performance of das. Raid-level configurations, storage controller protocols, and the efficiency of the bus are additional factors that affect the performance of das. The absence of storage interconnects and network latency provide das with the potential to outperform other storage networking configurations.

Das disk drive interfaces

The host and the storage device in das communicate with each other by using predefined protocols such as IDE/ATA, SATA, sas, scsi, and FC. These proto- cols are implemented on the HDD controller. Therefore, a storage device is also known by the name of the protocol it supports. This section describes each of these storage devices in detail.

IDE/ATA

An integrated device electronics/advanced technology attachment (IDE/ATA) disk supports the ide protocol. The term IDE/ATA conveys the dual-naming conventions for various generations and variants of this interface. The ide component in IDE/ATA provides the specification for the controllers connected.

To the computer's motherboard for communicating with the device attached. The ATA component is the interface for connecting storage devices, such as cd-roms, floppy disk drives, and HDDs, to the motherboard.

IDE/ATA has a variety of standards and names, such as ATA, ATA/atapi, eide, ATA-2, fast ATA, ATA-3, ultra ATA, and ultra dma. The latest version of ATA—ultra dma/133—supports a throughput of 133 mb per second.

In a master-slave configuration, an ATA interface supports two storage devices per connector. However, if the performance of the drive is important, sharing a port between two devices is not recommended.

SATA

A SATA (serial ATA) is a serial version of the IDE/ATA specification. SATA is a disk-interface technology that was developed by a group of the industry's leading vendors with the aim of replacing parallel ATA.

A SATA provides point-to-point connectivity up to a distance of one meter and enables data transfer at a speed of 150 mb/s. Enhancements to the SATA have increased the data transfer speed up to 600 mb/s.

A SATA bus directly connects each storage device to the host through a dedicated link, making use of low-voltage differential signaling. Lvds is an electrical signaling system that can provide high-speed connectivity over. Lowcost, twisted-pair copper cables. For data transfer, a SATA bus uses lvds With a voltage of 250 mv.A SATA bus uses a small 7-pin connector and a thin cable for connectivity. A SATA port uses 4 signal pins, which improves its pin efficiency compared to The parallel ATA that uses 26 signal pins, for connecting an 80-conductor ribbon Cable to a 40-pin header connector. SATA devices are hot-pluggable, which means that they can be connected or Removed while the host is up and running. A SATA port permits single-device Connectivity. Connecting multiple SATA drives to a host requires multiple ports .To be present on the host. Single-device connectivity enforced in SATA, eliminates.The performance problems caused by cable or port sharing in IDE/ATA.

Parallel scsi [Rgpv Dec2015(3)]

Scsi is available in a variety of interfaces. Parallel scsi (referred to as scsi) is one of the oldest and most popular forms of storage interface used in hosts. Scsi is a set of standards used for connecting a peripheral device to a computer and transferring data between them. Often, scsi is used to connect HDDs and tapes to a host. Scsi can also connect a wide variety of other devices such as scanners and printers. Communication between the hosts and the storage devices uses the scsi command set.Since its inception, scsi has undergone rapid revisions, resulting in continu- ous performance improvements. The oldest scsi variant, called scsi-1 provided data transfer rate of 5 mb/s; scsi ultra 320 provides data transfer speeds of 320Mb/s.

S.no	Rgpv question	Year	Marks
Q.1	Write down advantage nd disadvantage of DAS.	Dec2014 Dec2013	2 7
Q,2	What are the Limitations of DAS?	Dec2012	10

Unit-03/Lecture-02

NAS(network attached storage) – [Rgpv/dec 2015(2), Rgpv/dec2012(10), Rgpv/dec2011(10)]

Network attached storage (NAS) is an ip-based file-sharing device attached to a local area network. NAS provides the .Advantages of server consolidation by eliminating the need for multiple file servers. It provides storage consolidation through file-level data access and sharing. NAS is a preferred storage solution that enables clients to share files quickly and directly with minimum storage management overhead. NAS uses network and file-sharing protocols to perform filing and storage.Functions. These (ftp) and other Protocols for both environments. Recent advancements in networking technology have enabled NAS to scale up to enterprise requirements for improved Performance and reliability in accessing data.

A NAS device is a dedicated, high-performance, high-speed, single-purpose File serving and storage system. NAS serves a mix of clients and servers over an Ip network. Most NAS devices support multiple interfaces and networks. A NAS device uses its own operating system and integrated hardware, soft-Ware components to meet specific file service needs. Its operating system is Optimized for file i/o and, therefore, performs file i/o better than a general- Purpose server. As a result, a NAS device can serve more clients than traditional File servers, providing the benefit of server consolidation.

Benefits of NAS

NAS offers the following benefits:

- Supports comprehensive access to information: enables efficient file sharing and supports many-to-one and one-to-many configurations. The many-to-one configuration enables a NAS device to serve many clients simultaneously. The one-to-many configuration enables one client to connect with many NAS devices simultaneously.

- Improved efficiency: eliminates bottlenecks that occur during file access from a general-purpose file server because NAS uses an operating system specialized for file serving. It improves the utilization of general-purpose servers by relieving them of file-server operations.
- Improved flexibility: compatible for clients on both unix and windows platforms using industry-standard protocols. NAS is flexible and can serve requests from different types of clients from the same source.
- Centralized storage: centralizes data storage to minimize data dupli- cation on client workstations, simplify data management, and ensures greater data protection.
- Simplified management: provides a centralized console that makes it possible to manage file systems efficiently.
- scalability: scales well in accordance with different utilization profiles and types of business applications because of the high performance and low-latency design.
- High availability: offers efficient replication and recovery options, enabling high data availability. NAS uses redundant networking components that provide maximum connectivity options. A NAS device can use clustering technology for failover.
- Security: ensures security, user authentication, and file locking in conjunc- tion with industry-standard security schemas.

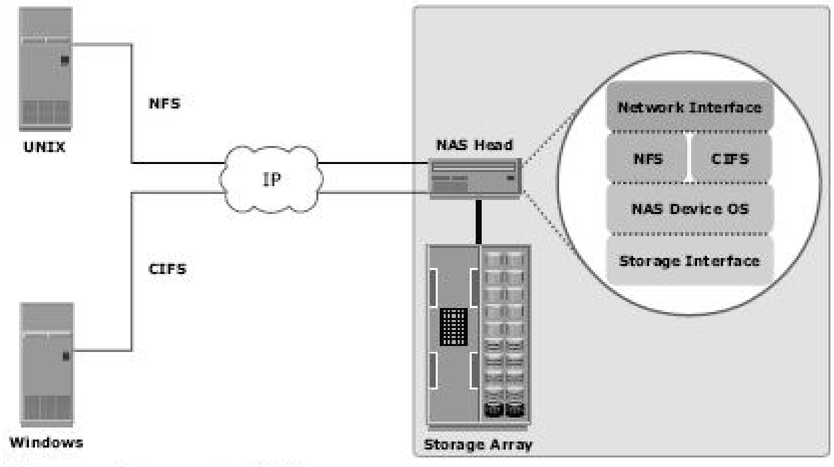

NAS(network attached storage) implementation

There are two types of NAS implementations: integrated and gateway. The integrated NAS device has all of its components and storage system in a single enclosure. In gateway implementation, NAS head shares its storage with SAN environment.

integrated NAS

An integrated NAS device has all the components of NAS, such as the NAS head and storage, in a single enclosure, or frame. This makes the integrated NAS a self-contained environment. The NAS head connects to the ip network to provide connectivity to the clients and service the file i/o requests. The storage consists of a number of disks that can range from low-cost ATA to high- throughput FC disk drives. Management software manages the NAS head and storage configurations.

An integrated NAS solution ranges from a low-end device, which is a single enclosure, to a high-end solution that can have an externally connected storage array.

A low-end appliance-type NAS solution is suitable for applications that a small department may use, where the primary need is consolidation of storage, rather than high performance or advanced features such as disaster recovery and business continuity. This solution is fixed in capacity and might not be upgradable beyond its original configuration. To expand the capacity, the solution must be scaled by deploying additional units, a task that increases management overhead because multiple devices have to be administered.

In a high-end NAS solution, external and dedicated storage can be used. This enables independent scaling of the capacity in terms of NAS heads or storage. However, there is a limit to scalability of this solution.

Gateway NAS

A gateway NAS device consists of an independent NAS head and one or more storage arrays. The NAS head performs the same functions that it does in the integrated solution; while the storage is shared with other applications that require block-level i/o. Management functions in this type of solution are more complex than those in an integrated environment because there are separate administrative tasks for the NAS head and the storage. In addition to the components that are explicitly tied to the NAS solution, a gateway solution can also utilize the FC infrastructure, such as switches, directors, or direct-attached storage arrays.

The gateway NAS is the most scalable because NAS heads and storage arrays can be independently scaled up when required. Adding processing capacity to the NAS gateway is an example of scaling. When the storage limit is reached, it can scale up, adding capacity on the SAN independently of the NAS head. Administrators can increase performance and i/o processing capabilities for their environments without purchasing additional interconnect devices and storage. Gateway NAS enables high utilization of storage capacity by sharing it with SAN environment.

Integrated NAS connectivity

An integrated solution is self-contained and can connect into a standard ip network. Although the specifics of how devices are connected within a NAS implementation vary by vendor and model. In some cases, storage is embedded within a NAS device and is connected to the NAS head through internal connections, such as ATA or scsi controllers. In others, the storage may be external but connected by using scsi controllers. In a high-end integrated NAS model, external storage can be directly connected by FC or by dedicated FC switches. In the case of a low-end integrated NAS model, backup traffic is shared on the same public ip network along with the regular client access traffic. In the case of a high-end integrated NAS model, an isolated backup network can be used to segment the traffic from impeding client access. More complex solutions may include an intelligent storage subsystem, enabling faster backup and larger capacities while simultaneously enhancing performance. Figure 7-4 illustrates an example of integrated NAS connectivity.

Gateway NAS connectivity

In a gateway solution, front-end connectivity is similar to that in an integrated solution. An integrated environment has a fixed number of NAS heads, making it relatively easy to determine ip networking requirements. In contrast, networking requirements in a gateway environment are complex to determine due to scalability options. Adding more NAS heads may require additional networking connectivity and bandwidth.

Communication between the NAS gateway and the storage system in a gate- way solution is achieved through a traditional FC SAN. To deploy a stable NAS solution, factors such as multiple paths for data, redundant fabrics, and load distribution must be considered.

Factors affecting NAS performance

As NAS uses IP network, bandwidth and latency issues associated with IP affect NAS performance. Network congestion is one of the most significant sources of latency in a NAS environment. Other factors that affect NAS performance at different levels are:

1. **Number of hops:** A large number of hops can increase latency because IP processing is required at each hop, adding to the delay caused at the router.

2. **Authentication with a directory service such as LDAP, Active Directory, or NIS:** The authentication service must be available on the network, with adequate bandwidth, and must have enough resources to accommodate the authentication load. Otherwise, a large number of authentication requests are presented to the servers, increasing latency. Authentication adds to latency only when authentication occurs.

3. **Retransmission:** Link errors, buffer overflows, and flow control mechanisms can result in retransmission. This causes packets that have not reached the specified destination to be resent. Care must be taken when configuring parameters for speed and duplex settings on the network devices and the NAS heads so that they match. Improper configuration may result in errors and retransmission, adding to latency.

4. **Overutilized routers and switches:** The amount of time that an overutilized device in a network takes to respond is always more than the response time of an optimally utilized or underutilized device. Network administrators can view vendor-specific statistics to determine the utilization of switches and routers in a network. Additional devices should be added if the current devices are over utilized.

5. **File/directory lookup and metadata requests:** NAS clients access files on NAS devices. The processing required before reaching the appropriate file or directory can cause delays. Sometimes a delay is caused by deep directory structures and can be resolved by flattening the directory structure. Poor file system layout and an over utilized disk system can also degrade performance.

6. **Overutilized NAS devices:** Clients accessing multiple files can cause high utilization levels on a NAS device which can be determined by viewing utilization statistics. High utilization levels can be caused by a poor file system structure or insufficient resources

in a storage subsystem.

7. **Overutilized clients:** The client accessing CIFS or NFS data may also be over utilized. An overutilized client requires longer time to process the responses received from the server, increasing latency. Specific performance-monitoring tools are available for various operating systems to help determine the utilization of client resources.

S.no	Rgpv question	Year	Marks
Q.1	What is NAS? Explain how the performance of NAS can be affected if sender and receiver window is not synchronized?	Dec2015 Dec2012	2 10
Q.2	Discuss various factors that affect NAS performance?	Dec2012	10

Unit-03/Lecture -03

SAN(storage area network) - [Rgpv/dec2013(7)]

Direct-attached storage (das) is often referred to as a stovepiped storage environment. Hosts “own” the storage and it is difficult to manage and share resources on these isolated storage devices. Efforts to organize this dispersed data led to the emergence of the storage area network (SAN). SAN is a high- speed, dedicated network of servers and shared storage devices. Traditionally connected over fibre channel (FC) networks, a SAN forms a single-storage pool and facilitates data centralization and consolidation. SAN meets the stor- age demands efficiently with better economies of scale. A SAN also provides effective maintenance and protection of data.

Components of SAN

A SAN consists of three basic components: servers, network infrastructure, and storage. These components can be further broken down into the following key elements: node ports, cabling, interconnecting devices (such as FC switches or hubs), storage arrays, and SAN management software.

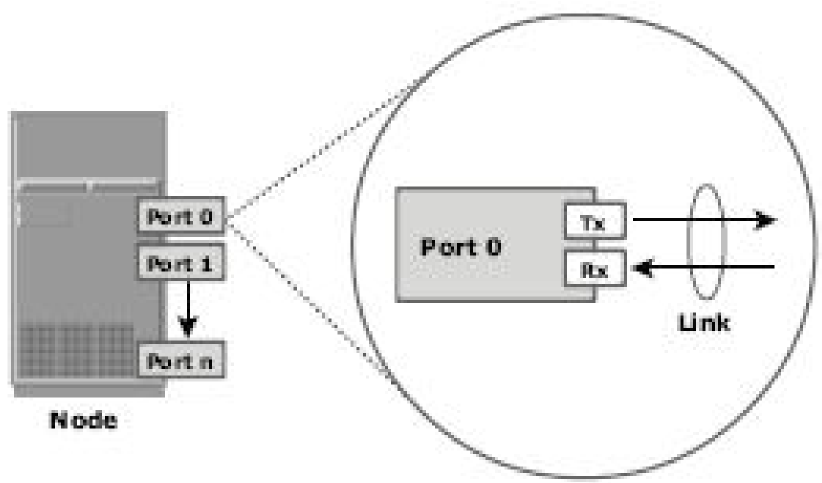

Nodes links and lines

Cabling

SAN implementations use optical fiber cabling. Copper can be used for shorter distances for back-end connectivity, as it provides a better signal-to-noise ratio for distances up to 30 meters. Optical fiber cables carry data in the form of light. There are two types of optical cables, multi-mode and single-mode.

Multi-mode fiber (mmf) cable carries multiple beams of light projected at different angles simultaneously onto the core of the cable. Based on the bandwidth, multi-mode fibers are classified as om1 (62.5µm), om2 (50µm) and laser optimized om3 (50µm). In an mmf transmission, multiple light beams traveling inside the cable tend to disperse and collide. This collision weakens the signal strength after it travels a certain distance — a process known as modal dispersion. An mmf cable is usually used for distances of up to 500 meters because of signal degradation (attenuation) due to modal dispersion.

Single-mode fiber (smf) carries a single ray of light projected at the center of the core. These cables are available in diameters of 7–11 microns; the most common size is 9 microns. In an smf transmission, a single light beam travels in a straight line through the core of the fiber. The small core and the single light wave limits modal dispersion. Among all types of fibre cables, single-mode provides minimum signal attenuation over maximum distance (up to 10 km). A single-mode cable is used for long-distance cable runs, limited only by the power of the laser at the transmitter and sensitivity of the receiver.

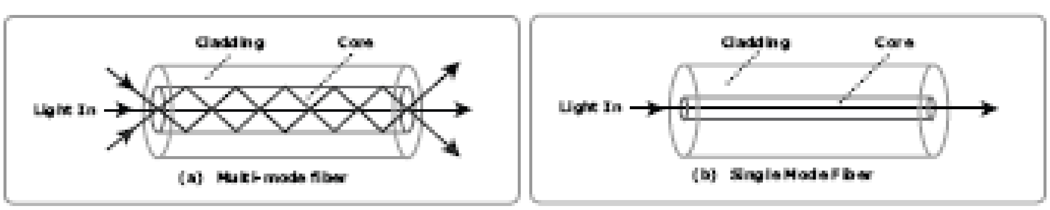

Multi mode and single mode fiber

Interconnect devices

Hubs, switches, and directors are the interconnect devices commonly used in SAN.
Hubs are used as communication devices in FC-al implementations. Hubs physically connect nodes in a logical loop or a physical star topology. All the nodes must share the bandwidth because data travels through all the connection points. Because of availability of low cost and high performance switches, hubs are no longer used in sans. Switches are more intelligent than hubs and directly route data from one physical port to another. Therefore, nodes do not share the bandwidth. Instead, each node has a dedicated communication path, resulting in bandwidth aggregation.

SAN management software

SAN management software manages the interfaces between hosts, interconnect devices, and storage arrays. The software provides a view of the SAN environment and enables management of various resources from one central console. It provides key management functions, including mapping of storage devices, switches, and servers, monitoring and generating alerts for discovered devices, and logical partitioning of the SAN, called zoning. In addition, the software provides management of typical SAN components such as storage components, and interconnecting devices .

Connectivity of SAN(storage area network)

FC connectivity
The FC architecture supports three basic interconnectivity options: point-to- point, arbitrated loop (FC-al), and fabric connect.
Point-to-point
Point-to-point is the simplest FC configuration — two devices are connected directly to each other, as shown in figure 6-6. This configuration provides a dedicated connection for

data transmission between nodes. However, the point-to-point configuration offers limited connectivity, as only two devices can communicate with each other at a given time. Moreover, it cannot be scaled to accommodate a large number of network devices. Standard das usess point- to-point connectivity.

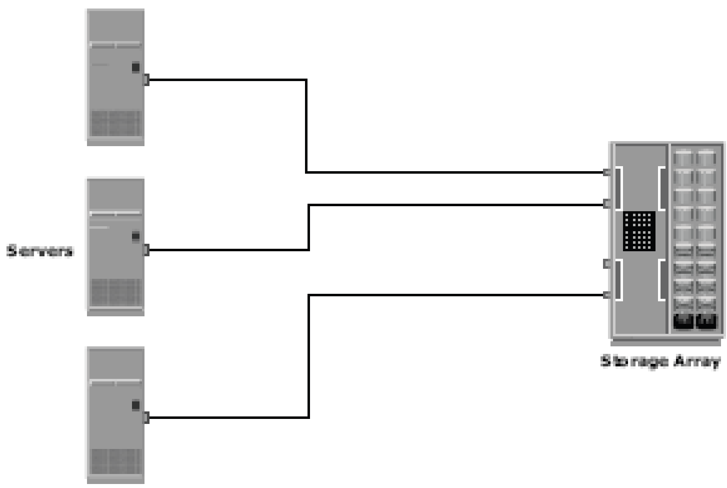

Point to point topology

Fibre channel arbitrated loop

In the FC-al configuration, devices are attached to a shared loop, FC-al has the characteristics of a token ring topology and a physical star topology. In FC-al, each device contends with other devices to perform i/o operations. Devices on the loop must "arbitrate" to gain control of the loop. At any given time, only one device can perform i/o operations on the loop.

Fiber channel ports – [Rgpv/dec2013(7)]

Ports are the basic building blocks of an FC network. Ports on the switch can be one of the following types:

- n_port: an end point in the fabric. This port is also known as the node port. Typically, it is a host port (hba) or a storage array port that is connected to a switch in a switched fabric.

■ nl_port: a node port that supports the arbitrated loop topology. This port is also known as the node loop port.

■ e_port: an FC port that forms the connection between two FC switches. This port is also known as the expansion port. The e_port on an FC switch Connects to the e_port of another FC switch in the fabric through a Link, which is called an inter-switch link (isl). Isls are used to transfer Host-to-storage data as well as the fabric management traffic from one switch to another. Isl is also one of the scaling mechanisms in SAN Connectivity.

■ f_port: a port on a switch that connects an n_port. It is also known as a Fabric port and cannot participate in FC-al.

■ fl_port: a fabric port that participates in FC-al. This port is connected to the nl_ports on an FC-al loop. A fl_port also connects a loop to a switch in a switched fabric. As a result, all nl_ports in the loop can participate in FC-sw. This configuration is referred to as a public loop. In contrast, an arbitrated loop without any switches is referred to as a private loop. A private loop contains nodes with nl_ports, and does not contain fl_port.

■ g_port: a generic port that can operate as an e_port or an f_port and determines its functionality automatically during initialization.

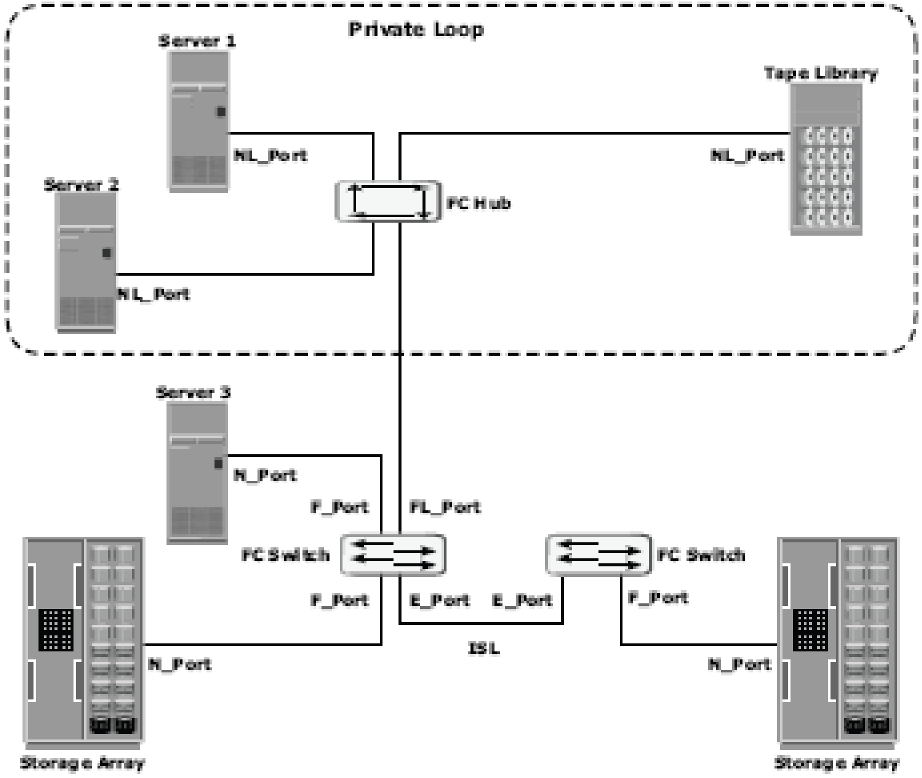

Fibre channel port

Fibre channel topologies[Rgpv/dec2012(10)]

There are three major fibre channel topologies, describing how a number of ports are connected together. A port in fibre channel terminology is any entity that actively communicates over the network, not necessarily a hardware port. This port is usually implemented in a device such as disk storage, an HBA on a server or a fibre channel switch.[1]

- **Point-to-point** (FC-P2P). Two devices are connected directly to each other. This is the simplest topology, with limited connectivity.

- **Arbitrated loop** (FC-AL). In this design, all devices are in a loop or ring, similar to token ring networking. Adding or removing a device from the loop causes all activity on the loop to be interrupted. The failure of one device causes a break in the ring. Fibre Channel hubs exist to connect multiple devices together and may bypass failed ports. A loop may also be made by cabling each port to the next in a ring.

- A minimal loop containing only two ports, while appearing to be similar to FC-P2P, differs considerably in terms of the protocol.
- Only one pair of ports can communicate concurrently on a loop.
- Maximum speed of 8GFC.

- **switched fabric** (FC-SW). All devices or loops of devices are connected to fibre channel switches, similar conceptually to modern ethernet implementations. Advantages of this topology over FC-P2P or FC-AL include:
  - The switches manage the state of the fabric, providing optimized interconnections.
  - The traffic between two ports flows through the switches only, it is not transmitted to any other port.
  - Failure of a port is isolated and should not affect operation of other ports.
  - Multiple pairs of ports may communicate simultaneously in a fabric.

S.no	Rgpv question	Year	Marks
Q.1	Explain different types of FC ports?	Rgpv Dec2013	7
Q.2	Explain different interfaces in FC?	Rgpv Dec2012	10
Q.3	Discuss the advantages FC-SW and FC-AL?		

Unit-03/Lecture -04

Content-addressed storage(CAS) – [Rgpv/dec 2014(3),Rgpv/dec2013(10)]

CAS is an object-based system that has been purposely built for storing fixed content data. It is designed for secure online storage and retrieval of fixed content. Unlike file-level and block-level data access that use file names and the physical location of data for storage and retrieval, CAS stores user data and its attributes as separate objects. The stored object is assigned a globally unique address known as a content address (ca). This address is derived from the object's binary representation. CAS provides an optimized and centrally managed storage solution that can support single-instance storage (sis) to eliminate multiple copies of the same data.

Features and benefits of CAS

CAS has emerged as an alternative to tape and optical solutions because it over- comes many of their obvious deficiencies. CAS also meets the demand to improve data accessibility and to properly protect, dispose of, and ensure service level agreements for archived data. The features and benefits of CAS include the following:

Content authenticity: it assures the genuineness of stored content. This is achieved by generating a unique content address and automating the process of continuously checking and recalculating the content address for stored objects. Content authenticity is assured because the address assigned to each piece of fixed content is as unique as a fingerprint. Every time an object is read, CAS uses a hashing algorithm to recalculate the object's content address as a validation step and compares the result to its original content address. If the object fails validation, it is rebuilt from its mirrored copy.

Content integrity: refers to the assurance that the stored content has not been altered.

Use of hashing algorithm for content authenticity also ensures content integrity in CAS. If the fixed content is altered, CAS assigns a new address to the altered content, rather than overwrite the original fixed content, providing an audit trail and maintaining the fixed content in its original state. As an integral part of maintaining data integrity and audit trail capabilities, CAS supports parity raid protection in addition to mirroring. Every object in a CAS system is systematically checked in the background. Over time, every object is tested, guarantee‐ ing content integrity even in the case of hardware failure, random error, or attempts to alter the content with malicious intent.

Llocation independence: CAS uses a unique identifier that applications can leverage to retrieve data rather than a centralized directory, path

Names, or urls. Using a content address to access fixed content makes the physical location of the data irrelevant to the application request‐ ing the data. Therefore the location from which the data is accessed is transparent to the application. This yields complete content mobility to applications across locations.

Single-instance storage (sis): the unique signature is used to guarantee the storage of only a single instance of an object. This signature is derived from the binary representation of the object. At write time, the CAS system is polled to see if it already has an object with the same signature. If the object is already on the system, it is not stored, rather only a pointer to that object is created. Sis simplifies storage resource management tasks, especially when handling hundreds of terabytes of fixed content.

Retention enforcement: protecting and retaining data objects is a core requirement of an archive system. CAS creates two immutable components: a data object and a meta‐object for every object stored. The meta‐ object stores object's attributes and data handling policies. For systems that support object‐retention capabilities, the retention policies are enforced until the policies expire.

Record-level protection and disposition: all fixed content is stored in CAS once and is backed up with a protection scheme. The array is com‐ posed of one or more storage clusters. Some CAS architectures provide an extra level of protection by replicating the content onto arrays located at a different location. The disposition of records also follows the stringent guidelines established by regulators for shredding and disposing of data in electronic formats.

technology independence: the CAS system interface is impervious to technology changes. As long as the application server is able to map the original content address the data remains accessible. Although hardware changes are inevitable, the goal of CAS hardware vendors is to ensure compatibility across platforms.

Fast record retrieval: CAS maintains all content on disks that provide subsecond "time to first byte" (200 ms–400 ms) in a single cluster. Random disk access in CAS enables fast record retrieval.

S.no	Rgpv question	Year	Marks
Q.1	To access data in a SAN , a host uses a physical address known a logical address. A host using a CAS device does not use or need a physical address Why?	Dec2013	10

Unit -03 /Lecture-06

Hub, switches, storage array [Rgpv/dec2013(7)]

A hub is the most basic networking device that connects multiple computers or other network devices together. Unlike a network switch or router, a network hub has no routing tables or intelligence on where to send information and broadcasts all network data across each connection. Most hubs can detect basic network errors such as collisions, but having all information broadcast to multiple ports can be a security risk and cause bottlenecks. In the past network hubs were popular because they were much cheaper than a switch and router, but today most switches do not cost much more than a hub and are a much better solution for any network.

In general, a hub refers to a hardware device that enables multiple devices or connections to be connected to a computer. Another example besides the one given above is a usb hub, which allows dozens of usb devices to be connected to one computer, even though that computer may only have a few usb connections. The picture is an example of a usb hub.

Switches

A switch is a device used on a computer network to physically connect devices together. Multiple cables can be connected to a switch to enable networked devices to communicate with each other. Switches manage the flow of data across a network by only transmitting a received message to the device for which the message was intended. Each networked device connected to a switch can be identified using a mac address, allowing the switch to regulate the flow of traffic. This maximises security and efficiency of the network.

Because of these features, a switch is often considered more "intelligent" than a network hub. Hubs neither provide security, or identification of connected devices. This means that messages have to be transmitted out of every port of the hub, greatly degrading the efficiency of the network. Switches may operate at one or more layers of the osi model, including the data link and network layers. A device that operates simultaneously at more than one of these layers is known as a multilayer switch.

In switches intended for commercial use, built-in or modular interfaces make it possible to connect different types of networks, including ethernet, fibre channel, atm, itu-t g.hn and 802.11. This connectivity can be at any of the layers mentioned. While layer-2 functionality is adequate for bandwidth-shifting within one technology, interconnecting technologies such as ethernet and token ring is easier at layer 3.

Devices that interconnect at layer 3 are traditionally called routers, so layer-3 switches can also be regarded as (relatively primitive) routers.

Where there is a need for a great deal of analysis of network performance and security, switches may be connected between wan routers as places for analytic modules. Some vendors provide firewall,[3][4] network intrusion detection, and performance analysis modules that can plug into switch ports. Some of these functions may be on combined modules.

In other cases, the switch is used to create a mirror image of data that can go to an external device. Since most switch port mirroring provides only one mirrored stream, network hubs can be useful for fanning out data to several read-only analyzers, such as intrusion detection systems and packet sniffers.

Storage array

The fundamental purpose of a SAN is to provide host access to storage resources. The large storage capacities offered by modern storage arrays have been exploited in SAN environments for storage consolidation and centralization. SAN implementations complement the standard features of storage arrays by providing high availability and redundancy, improved performance, business continuity, and multiple host connectivity.

S.no	Rgpv question	Year	Marks
Q.1	Explain briefly following A)node port, B)storage array, C)SAN, D)hub, E)switches	Dec2013	7

Additional Topic Unit - 03/Lecture - 07

Jbod: just a bunch of disks [Rgpv/dec2015(2)]

If we compare disk subsystems with regard to their controllers we can differentiate between three levels of complexity: (1) no controller; (2) raid controller and (3) intelligent controller with additional services such as instant copy and remote mirroring.

If the disk subsystem has no internal controller, it is only an enclosure full of disks (jbods). In this instance, the hard disks are permanently fitted into the enclosure and the connections for i/o channels and power supply are taken outwards at a single point. Therefore, a jbod is simpler to manage than a few loose hard disks. Typical jbod disk subsystems have space for 8 or 16 hard disks. A connected server recognises all these hard disks as independent disks. Therefore, 16 device addresses are required for a jbod disk subsystem incorporating 16 hard disks. In some i/o techniques such as scsi and fibre channel arbitrated loop this can lead to a bottleneck at device addresses in contrast to intelligent disk subsystems, a jbod disk subsystem in particular is not capable of supporting raid or other forms of virtualisation. If required, however, these can be realised outside the jbod disk subsystem, for example, as software in the server or as an independent virtualisation entity in the storage network.

Fiber channel overview

The fc architecture forms the fundamental construct of the san infrastruc- ture. *Fibre channel* is a high-speed network technology that runs on high-speed optical fiber cables (preferred for front-end san connectivity) and serial cop- per cables (preferred for back-end disk connectivity). The fc technology was created to meet the demand for increased speeds of data transfer among com- puters, servers, and mass storage subsystems. Although fc networking was introduced in 1988, the fc standardization process began when the american national standards institute (ansi) chartered the fibre channel working group (fcwg). By 1994, the new high-speed computer interconnection standard was developed and the fibre channel association (fca) was founded with 70 charter member companies. Technical committee, which is the committee within incits (international committee for information technology standards), is responsible for fibre channel interfaces. T11 (previously known as x3t9.3) has been producing interface standards for high performance and mass storage applications.

Higher data transmission speeds are an important feature of the fc network- ing technology. The initial implementation offered throughput of 100 mb/s (equivalent to raw bit rate of 1gb/s i.e. 1062.5 mb/s in fibre channel), which was greater than the speeds of ultra scsi (20 mb/s) commonly used in das environments. Fc in full-duplex mode could sustain throughput of 200 mb/s. In comparison with ultra-scsi, fc is a significant leap in storage networking tech- nology. Latest fc implementations of 8 gfc (fibre channel) offers throughput of 1600 mb/s (raw bit rates of 8.5 gb/s), whereas ultra320 scsi is available with a throughput of 320 mb/s. The fc architecture is highly scalable and theoreti- cally a single fc network can accommodate approximately 15 million node.

S.no	Rgpv question	Year	Marks
Q.1	What is JBOD?	Dec 2015	2

References

Book	Author	Priority
Information storage management	G. Somasundaram Alok Shrivastava	1
Storage Network explained : Basic and application of fiber channels, SAN, NAS, iSESI	Ulf Troppens, Wolfgang Mueller-Friedt, Rainer Erkens, Rainer Wolafka, Nils Haustein	2
Nick Antonopoulos, Lee Gillam	Cloud Computing : Principles, System & Application	3