

Storage

EMC,EVA,XP,SUN Storage,Storage,Brocade Switch,Switch



SAN Interview questions

Posted by Elumalai M

1.WHAT ARE THE BENEFITS OF FIBRE CHANNEL SANS?

Fibre Channel SANs are the de facto standard for storage networking in the corporate data center because they provide exceptional reliability, scalability, consolidation, and performance. Fibre Channel SANs provide significant advantages over direct-attached storage through improved storage utilization, higher data availability, reduced management costs, and highly scalable capacity and performance.

2.WHAT ENVIRONMENT IS MOST SUITABLE FOR FIBRE CHANNEL SANS?

AdChoices

- ▶ [Cloud Storage Drive](#)
- ▶ [EMC San Storage](#)
- ▶ [Data Storage](#)

About this blog

This blog useful for learning Storage technology.This site will be helpful for storage administrator in their day to day activities.Your comments on posts are welcome.

Typically, Fibre Channel SANs are most suitable for large data centers running business-critical data, as well as applications that require high-bandwidth performance such as medical imaging, streaming media, and large databases. Fibre Channel SAN solutions can easily scale to meet the most demanding performance and availability requirements.

3.WHAT CUSTOMER PROBLEMS DO FIBRE CHANNEL SANS SOLVE?

The increased performance of Fibre Channel enables a highly effective backup and recovery approach, including LAN-free and server-free backup models. The result is a faster, more scalable, and more reliable backup and recovery solution. By providing flexible connectivity options and resource sharing, Fibre Channel SANs also greatly reduce the number of physical devices and disparate systems that must be purchased and managed, which can dramatically lower capital expenditures. Heterogeneous SAN management provides a single point of control for all devices on the SAN, lowering costs and freeing personnel to do other tasks.

4.HOW LONG HAS FIBRE CHANNEL BEEN AROUND?

Development started in 1988, ANSI standard approval occurred in 1994, and large deployments began in 1998. Fibre Channel is a mature, safe, and widely deployed solution for high-speed (1 GB, 2 GB, 4 GB) communications and is the foundation for the majority of SAN installations throughout the world.

5.WHAT IS THE FUTURE OF FIBRE CHANNEL SANS?

Fibre Channel is a well-established, widely deployed technology with a proven track record and a very large installed base, particularly in high-performance, business-critical data center environments. Fibre Channel SANs continue to grow and will be enhanced for a long time to come. The reduced costs of Fibre Channel components, the availability of SAN kits, and the next generation of Fibre Channel (4 GB) are helping to fuel that growth. In addition, the Fibre Channel roadmap includes plans to double performance every three years

6.WHAT ARE THE BENEFITS OF 4GB FIBRE CHANNEL?

Storage

[EMC](#) (2)
[EVA](#) (1)
[NetAPP](#) (1)
[SAN](#) (1)
[Switch](#) (1)

AdChoices 

- ▶ [Raid 5 Storage](#)
- ▶ [NAS San](#)
- ▶ [Network Storage NAS](#)

SAN

[NetAPP](#)
[SAN Reference site](#)
[Raid Manager](#)
[EMC Client installation and checking](#)
[Configure lun for solaris](#)

Storage Docs Center

- ▼ [2011](#) (1)
 - ▼ [November](#) (1)
 - [SAN Interview questions](#)
- ▶ [2009](#) (5)

Benefits include twice the performance with little or no price increase, investment protection with backward compatibility to 2 GB, higher reliability due to fewer SAN components (switch and HBA ports) required, and the ability to replicate, back up, and restore data more quickly. 4 GB Fibre Channel systems are ideally suited for applications that need to quickly transfer large amounts of data such as remote replication across a SAN, streaming video on demand, modeling and rendering, and large databases. 4 GB technology is shipping today.

7.HOW IS FIBRE CHANNEL DIFFERENT FROM ISCSI?

Fibre Channel and iSCSI each have a distinct place in the IT infrastructure as SAN alternatives to DAS. Fibre Channel generally provides high performance and high availability for business-critical applications, usually in the corporate data center. In contrast, iSCSI is generally used to provide SANs for business applications in smaller regional or departmental data centers.

8.WHEN SHOULD I DEPLOY FIBRE CHANNEL INSTEAD OF ISCSI?

For environments consisting of high-end servers that require high bandwidth or data center environments with business-critical data, Fibre Channel is a better fit than iSCSI. For environments consisting of many midrange or low-end servers, an IP SAN solution often delivers the most appropriate price/performance.

9.Name some of the SAN topologies

Point-to-point, arbitrated loop, and switched fabric topologies

10.What's the need for separate network for storage why LAN cannot be used?

LAN hardware and operating systems are geared to user traffic, and LANs are tuned for a fast user response to messaging requests.

With a SAN, the storage units can be secured separately from the servers and totally apart from the user network enhancing storage access in data blocks (bulk data transfers), advantageous for server-less backups.

11.What are the advantages of RAID?

“Redundant Array of Inexpensive Disks”

Depending on how we configure the array, we can have the

- data mirrored [RAID 1] (duplicate copies on separate drives)
- striped [RAID 0] (interleaved across several drives), or
- parity protected [RAID 5](extra data written to identify errors).

These can be used in combination to deliver the balance of performance and reliability that the user requires.

12. Define RAID? Which one you feel is good choice?

RAID (Redundant array of Independent Disks) is a technology to achieve redundancy with faster I/O. There are Many Levels of RAID to meet different needs of the customer which are: R0, R1, R3, R4, R5, R10, R6.

Generally customer chooses R5 to achieve better redundancy and speed and it is cost effective.

R0 – Striped set without parity/[Non-Redundant Array].

Provides improved performance and additional storage but no fault tolerance. Any disk failure destroys the array, which becomes more likely with more disks in the array. A single disk failure destroys the entire array because when data is written to a RAID 0 drive, the data is broken into fragments. The number of fragments is dictated by the number of disks in the drive. The fragments are written to their respective disks simultaneously on the same sector. This allows smaller sections of the entire chunk of data to be read off the drive in parallel, giving this type of arrangement huge bandwidth. RAID 0 does not implement error checking so any error is unrecoverable. More disks in the array means higher bandwidth, but greater risk of data loss

R1 - Mirrored set without parity.

Provides fault tolerance from disk errors and failure of all but one of the drives. Increased read performance occurs when using a multi-threaded operating system that supports split seeks, very small performance reduction when writing. Array continues to operate so long as at least one drive is functioning. Using RAID 1 with a separate controller for each disk is sometimes called duplexing.

R3 - Striped set with dedicated parity/Bit interleaved parity.

This mechanism provides an improved performance and fault tolerance similar to RAID 5, but with a dedicated parity disk rather than rotated parity stripes. The single parity disk is a bottle-neck for writing since every write requires updating the parity data. One minor benefit is the dedicated parity disk allows the parity drive to fail and operation will continue without parity or performance penalty.

R4 - Block level parity.

Identical to RAID 3, but does block-level striping instead of byte-level striping. In this setup, files can be distributed between multiple disks. Each disk operates independently which allows I/O requests to be performed in parallel, though data transfer speeds can suffer due to the type of parity. The error detection is achieved through dedicated parity and is stored in a separate, single disk unit.

R5 - Striped set with distributed parity.

Distributed parity requires all drives but one to be present to operate; drive failure requires replacement, but the array is not destroyed by a single drive failure. Upon drive failure, any subsequent reads can be calculated from the distributed parity such that the drive failure is masked from the end user. The array will have data loss in the event of a second drive failure and is vulnerable until the data that was on the failed drive is rebuilt onto a replacement drive.

R6 - Striped set with dual distributed Parity.

Provides fault tolerance from two drive failures; array continues to operate with up to two failed drives. This makes larger RAID groups more practical, especially for high availability systems. This becomes increasingly important because large-capacity drives lengthen the time needed to recover from the failure of a single drive. Single parity RAID levels are vulnerable to data loss until the failed drive is rebuilt: the larger the drive, the longer the rebuild will take. Dual parity gives time to rebuild the array without the data being at risk if one drive, but no more, fails before the rebuild is complete.

13.What is the difference between RAID 0+1 and RAID 1+0

RAID 0+1 (Mirrored Stripped)

In this RAID level all the data is saved on striped volumes which are in turn mirrored, so any disk failure saves the data loss but it makes whole stripe unavailable. The key difference from RAID 1+0 is that RAID 0+1 creates a second striped set to mirror a primary striped set. The array continues to operate with one or more drives failed in the same mirror set, but if drives fail on both sides of the mirror the data on the RAID system is lost. In this RAID level if one disk is failed full mirror is marked as inactive and data is saved only one striped volume.

RAID 1+0 (Striped Mirrored)

In this RAID level all the data is saved on mirrored volumes which are in turn striped, so any disk failure saves data loss. The key difference from RAID 0+1 is that RAID 1+0 creates a striped set from a series of mirrored drives. In a failed disk situation RAID 1+0 performs better because all the remaining disks continue to be used. The array can sustain multiple drive losses so long as no mirror loses both its drives.

This RAID level is most preferred for high performance and high data protection because rebuilding of RAID 1+0 is less time consuming in comparison to RAID 0+1.

14. When JBOD's are used?

“Just a Bunch of Disks”

It is a collection of disks that share a common connection to the server, but don't include the mirroring, striping, or parity facilities that RAID systems do, but these capabilities are available with host-based software.

15. Differentiate RAID & JBOD?

RAID: “Redundant Array of Inexpensive Disks”

Fault-tolerant grouping of disks that server sees as a single disk volume

Combination of parity-checking, mirroring, striping

Self-contained, manageable unit of storage

JBOD: “Just a Bunch of Disks”

Drives independently attached to the I/O channel
Scalable, but requires server to manage multiple volumes
Do not provide protection in case of drive failure

16. What is a HBA?

Host bus adapters (HBAs) are needed to connect the server (host) to the storage.

17. What are the advantages of SAN?

Massively extended scalability

Greatly enhanced device connectivity

Storage consolidation

LAN-free backup

Server-less (active-fabric) backup

Server clustering

Heterogeneous data sharing

Disaster recovery - Remote mirroring

While answering people do NOT portray clearly what they mean & what advantages each of them have, which are cost effective & which are to be used for the client's requirements.

18. What is the difference b/w SAN and NAS?

The basic difference between SAN and NAS, SAN is Fabric based and NAS is Ethernet based.

SAN - Storage Area Network

It accesses data on block level and produces space to host in form of disk.

NAS - Network attached Storage

It accesses data on file level and produces space to host in form of shared network folder.

19. What is a typical storage area network consists of - if we consider it for implementation in a small business setup?

If we consider any small business following are essential components of SAN

- Fabric Switch
- FC Controllers
- JBOD's

20.Can you briefly explain each of these Storage area components?

Fabric Switch: It's a device which interconnects multiple network devices .There are switches starting from 16 port to 32 ports which connect 16 or 32 machine nodes etc. vendors who manufacture these kind of switches are Brocade, McData.

21.FC Controllers: These are Data transfer media they will sit on PCI slots of Server; you can configure Arrays and volumes on it.

JBOD: Just Bunch of Disks is Storage Box, it consists of Enclosure where set of hard-drives are hosted in many combinations such SCSI drives, SAS, FC, SATA.

22.What is the most critical component in SAN?

Each component has its own criticality with respect to business needs of a company.

23.How is a SAN managed?

There are many management software's used for managing SAN's to name a few

- Santricity
- IBM Tivoli Storage Manager.
- CA Unicenter.
- Veritas Volumemanger.

24.Which one is the Default ID for SCSI HBA?

Generally the default ID for SCSI HBA is 7.

SCSI- Small Computer System Interface

25.What is the highest and lowest priority of SCSI?

There are 16 different ID's which can be assigned to SCSI device 7, 6, 5, 4, 3, 2, 1, 0, 15, 14, 13, 12, 11, 10, 9, 8.

Highest priority of SCSI is ID 7 and lowest ID is 8.

26.How do you install device drivers for the HBA first time during OS installation?

In some scenarios you are supposed to install Operating System on the drives connected thru SCSI HBA or SCSI RAID Controllers, but most of the OS will not be updated with drivers for those controllers, that time you need to supply drivers externally, if you are installing windows, you need to press F6 during the installation of OS and provide the driver disk or CD which came along with HBA.

If you are installing Linux you need to type "linux dd" for installing any driver.

27.What is Array?

Array is a group of Independent physical disks to configure any Volumes or RAID volumes.

28.Can u describe at-least 3 troubleshooting scenarios which you have come across in detail ?

SCENARIO 1: How do you find/debug when there is error while working SCSI devices?

In our daily SAN troubleshooting there are many management and configuration tools we use them to see when there is a failure with target device or initiator device.

Some time it is even hard to troubleshoot some of the things such as media errors in the drives, or some of the drives taking long time to spin-up. In such cases these utilities will not come to help. To debug this kind of information most of the controller will be implemented with 3-pin serial debug port. With serial port debug connector cable you can collect the debug information with hyper terminal software.

SCENARIO 2: I am having an issue with a controller its taking lot of time to boot and detect all the drives connected how can I solve this.?

There are many possibilities that might cause this problem. One of the reason might be you are using bad drives that cannot be repaired. In those cases you replace the disks with working ones.

Another reason might be slots you connected your controller to a slot which might not be supported.

Try to connect with other types of slots.

One more probable reason is if you have flashed the firmware for different OEM's on the same hardware.

To get rid of this the flash utilities will be having option to erase all the previous and EEPROM and boot block entry option. Use that option to rectify the problem.

SCENARIO 3: I am using tape drive series 700X, even the vendor information on the Tape drive says 700X, but the POST information while booting the server is showing as 500X what could be the problem?

First you should make sure your hardware is of which series, you can find out this in the product website.

Generally you can see this because in most of the testing companies they use same hardware to test different series of same hardware type. What they do is they flash the different series firmware. You can always flash back to exact hardware type.

29. Which are the SAN topologies?

SAN can be connected in 3 types which are mentioned below :

Point to Point topology

FC Arbitrated Loop (FC :Fibre Channel)

Switched Fabric

30. Which are the 4 types of SAN architecture types

a. Core-edge

- b.Full-Mesh
- c.Partial-Mesh
- d.Cascade

31.Which command is used in linux to know the driver version of any hardware device?

dmesg

32.How many minimum drives are required to create R5 (RAID 5) ?

You need to have at least 3 disk drives to create R5.

33.Can you name some of the states of RAID array?

There are states of RAID arrays that represent the status of the RAID arrays which are given below

- a. Online
- b. Degraded
- c. Rebuilding
- d. Failed

34.Name the features of SCSI-3 standard?

QAS: Quick arbitration and selection
Domain Validation
CRC: Cyclic redundancy check

35.Can we assign a hot spare to R0 (RAID 0) array?

No, since R0 is not redundant array, failure of any disks results in failure of the entire array so we cannot rebuild the hot spare for the R0 array.

36.Can you name some of the available tape media types?

There are many types of tape media available to back up the data some of them are

DLT: digital linear tape - technology for tape backup/archive of networks and servers; DLT technology addresses midrange to high-end tape backup requirements.

LTO: linear tape open; a new standard tape format developed by HP, IBM, and Seagate.

AIT: advanced intelligent tape; a helical scan technology developed by Sony for tape backup/archive of networks and servers, specifically addressing midrange to high-end backup requirements.

37.What is HA?

HA High Availability is a technology to achieve failover with very less latency. Its a practical requirement of data centers these days when customers expect the servers to be running 24 hours on all 7 days around the whole 365 days a year - usually referred as 24x7x365. So to achieve this, a redundant infrastructure is created to make sure if one database server or if one app server fails there is a replica Database or Appserver ready to take-over the operations. End customer never experiences any outage when there is a HA network infrastructure.

38.What is virtualization?

Virtualization is logical representation of physical devices. It is the technique of managing and presenting storage devices and resources functionally, regardless of their physical layout or location. Virtualization is the pooling of physical storage from multiple network storage devices into what appears to be a single storage device that is managed from a central console. Storage virtualization is commonly used in a storage area network (SAN). The management of storage devices can be tedious and time-consuming. Storage virtualization helps the storage administrator perform the tasks of backup, archiving, and recovery more easily, and in less time, by disguising the actual complexity of the SAN.

39.Describe in brief the composition of FC Frame?

Start of the Frame locator

Frame header (includes destination id and source id, 24 bytes/6 words)

Data Payload (encapsulate SCSI instruction can be 0-2112 bytes in length)

CRC (error checking, 4 bytes)

End of Frame (1 byte)

40. What is storage virtualization?

Storage virtualization is amalgamation of multiple n/w storage devices into single storage unit.

41. What are the protocols used in physical/datalink and network layer of SAN?

- a) Ethernet
- b) SCSI
- c) Fibre Channel

42. What are the types of disk array used in SAN?

- a) JBOD
- b) RAID

43. What are different types of protocols used in transportation and session layers of SAN?

- a) Fibre Channel Protocol (FCP)
- b) Internet SCSI (iSCSI)
- c) Fibre Channel IP (FCIP)

44. What is the type of Encoding used in Fibre Channel?

8b/10b, as the encoding technique is able to detect almost all the bit errors

45.How many classes of service are available in Fibre Channel?

7 Classes of service are available in Fibre Channel

Class-1: Dedicated connection between two communicators with acknowledgement of frame delivery.

In class 1 service, a dedicated connection source and destination is established through the fabric for the duration of the transmission. It provides acknowledged service. This class of service ensures that the frames are received by the destination device in the same order in which they are sent, and reserves full bandwidth for the connection between the two devices. It does not provide for a good utilization of the available bandwidth, since it is blocking another possible contender for the same device. Because of this blocking and necessary dedicated connection, class 1 is rarely used.

Class-2: connection less but provides acknowledgement

Class 2 is a connectionless, acknowledged service. Class 2 makes better use of available bandwidth since it allows the fabric to multiplex several messages on a frame-by-frame basis. As frames travel through the fabric they can take different routes, so class 2 service does not guarantee in-order delivery. Class 2 relies on upper layer protocols to take care of frame sequence. The use of acknowledgments reduces available bandwidth, which needs to be considered in large-scale busy networks.

Class-3: connection less and provides no notification of delivery

There is no dedicated connection in class 3 and the received frames are not acknowledged. Class 3 is also called datagram connectionless service. It optimizes the use of fabric resources, but it is now upper layer protocol to ensure that all frames are received in the proper order, and to request to the source device the retransmission of missing frames. Class 3 is a commonly used class of service in Fibre Channel networks.

Class-4: allows fractional bandwidth for virtual circuits

Class 4 is a connection-oriented service like class 1, but the main difference is that it allocates only a fraction of available bandwidth of path through the fabric that connects two N_Ports. Virtual Circuits (VCs) are established between two N_Ports with guaranteed Quality of Service (QoS), including bandwidth and latency. Like class 1, class 4 guarantees in-order delivery frame delivery and provides acknowledgment of delivered frames, but now the fabric is responsible for multiplexing frames of different VCs. Class 4 service is mainly intended for multimedia applications

such as video and for applications that allocate an established bandwidth by department within the enterprise. Class 4 was added in the FC-PH-2 standard.

Class -5: Class 5 is called isochronous service, and it is intended for applications that require immediate delivery of the data as it arrives, with no buffering. It is not clearly defined yet. It is not included in the FC-PH documents.

Class-6: Provides multicast, dedicated connection with acknowledgment

Class 6 is a variant of class 1, known as multicast class of service. It provides dedicated connections for a reliable multicast. An N_Port may request a class 6 connection for one or more destinations. A multicast server in the fabric will establish the connections and get acknowledgment from the destination ports, and send it back to the originator. Once a connection is established, it should be retained and guaranteed by the fabric until the initiator ends the connection. Class 6 was designed for applications like audio and video requiring multicast functionality. It appears in the FC-PH-3 standard.

Class-F: used for switch to switch communication in the fabric.

Class F service is defined in the FC-SW and FC-SW-2 standard for use by switches communicating through ISLs. It is a connectionless service with notification of non-delivery between E_Ports used for control, coordination, and configuration of the fabric. Class F is similar to class 2; the main difference is that Class 2 deals with N_Ports sending data frames, while Class F is used by E_ports for control and management of the fabric.

46.What are the main constrains of SCSI in storage networking?

- a)Deployment distance (max. of 25 mts)
- b)Number of devices that can be interconnected (16)

47.What is a Fabric?

Interconnection of Fibre Channel Switches

48.What are the services provided by Fabric to all the nodes?

- a)Fabric Login
- b)SNS

- c)Fabric Address Notification
- d)Registered state change notification
- e)Broadcast Servers

49.What is the difference between LUN and WWN?

LUN: unique number that is assigned to each storage device or partition of the storage that the storage can support.

WWN: 64bit address that is hard coded into a fibre channel HBA and this is used to identify individual port (N_Port or F_Port) in the fabric.

50.What are the different topologies in Fibre Channel?

- a)Point-to-Point
- b)Arbitrary Loop
- c)Switched Fabric Loop

51.What are the layers of Fibre Channel Protocol?

- a)FC Physical Media
- b)FC Encoder and Decoder
- c)FC Framing and Flow control
- d)FC Common Services
- e)FC Upper Level Protocol Mapping

52.What is zoning?

Fabric management service that can be used to create logical subsets of devices within a SAN. This enables portioning of resources for management and access control purpose.

53.What are the two major classification of zoning?

Two types of zoning are

- a) Software Zoning
- b) Hardware Zoning

54. What are different levels of zoning?

- a) Port Level zoning
- b) WWN Level zoning
- c) Device Level zoning
- d) Protocol Level zoning
- e) LUN Level zoning

55. What are the 3 prominent characteristics of SAS Protocol?

- a) Native Command Queuing (NCQ)
- b) Port Multiplier
- c) Port Selector

56. What are the 5 states of Arbitrarily Loop in FC?

- a) Loop Initialization
- b) Loop Monitoring
- c) Loop arbitration
- d) Open Loop
- e) Close Loop

57. How does FC Switch maintain the addresses?

FC Switch uses simple name server (SNS) to maintain the mapping table

58. What is the purpose of disk array?

Probability of unavailability of data stored on the disk array due to single point failure is totally eliminated.

59.What is disk array?

Set of high performance storage disks that can store several terabytes of data. Single disk array can support multiple points of connection to the network.

60.What is virtualization?

A technique of hiding the physical characteristics of computer resources from the way in which other system application or end user interact with those resources. Aggregation, spanning or concatenation of the combined multiple resources into larger resource pools.

61.What is Multipath I/O?

Fault tolerant technique where, there is more than one physical path between the CPU in the computer systems and its main storage devices through the buses, controllers, switches and other bridge devices connecting them.

62.What is RAID?

Technology that groups several physical drives in a computer into an array that you can define as one or more logical drive. Each logical drive appears to the operating system as single drive. This grouping enhances the performance of the logical drive beyond the physical capability of the drives.

63.What is stripe-unit-size?

It is data distribution scheme that complement s the way operating system request data. Granularity at which data is stored on one drive of the array before subsequent data is stored on the next drive of the array. Stripe unit size should be close to the size of the system I/O request.

64.What is LUN Masking?

A method used to create an exclusive storage area and access control. And this can be achieved by storage device control program.

65.What is the smallest unit of information transfer in FC?

Frame

66.How is the capacity of the HDD calculated?

Number of Heads X Number of Cylinders X Sectors per track X Sector Size

67.What is bad block reallocation?

A bad sector is remapped or reallocated to good spare block and this information is stored in the internal table on the hard disk drive. The bad blocks are identified during the media test of the HDD as well as during various types of read write operations performed during the I/O tests. Apart from the new generation of HDD comes with a technology called BGMS (background media scan) which continuously scans the HDD media for defects and maps them when the drive is idle (this is performed after the HDD is attached to the system).

68.What are two types of recording techniques on the tapes?

- a)Linear Recording
- b)Helical Scan Recording.

69.What is snapshot?

A snapshot of data object contains an image of data at a particular point of time.

70.What is HSM?

Hierarchical storage management - An application that attempts to match the priority of data with the cost of storage.

71.What is hot-swapping?

Devices are allowed to be removed and inserted into a system without turning off the system.

72.What is Hot-Sparing?

A spare device is available to be inserted into the subsystem operation without having to remove and replace a device.

73.What are different types of backup system?

- a)Offline
- b)Online
- c)Near Line

74.What is the different between mirroring, Routing and multipathing?

Redundancy Functions Relationships Role

Mirroring Generates 2 ios to 2 storage targets Creates 2 copies of data

Routing Determined by switches independent of SCSI Recreates n/w route after a failure

Multipathing Two initiator to one target Selects the LUN initiator pair to use

75.Name few types of Tape storage?

- a) Digital Linear Tape
- b) Advanced Intelligent Tape
- c) Linear Tape Open

76. What is a sequence in FC?

Group of one or more frames that encompasses one or more "information units" of a upper layer protocol.

Example:

It requires

- i) One sequence to transfer the command
- ii) One or more sequence to transfer the data
- iii) Once sequence to transfer the status.

77. What is Exchange in FC?

Exchange is to establish a relationship between 2 N_PORTS and then these two ports transfer data via one or more sequence within this relationship.

Example: Exchange exist to transfer the command, data and the status of one SCSI task

78. Why do we need Login in FC?

Port Login: To exchange service parameters between N_Ports and N_Ports

Process Login: To establish the SCSI operating environment between two N_PORTS

Fabric Login: Similar to port login, FLOGI is an extended link service command that sets up a session between two participants. With FLOGU a session is created between an N_Port or NL_Port and the switch.

79. What are the different types of clusters?

- a) High availability clusters
- b) High Performance Clusters

c)Load Balancing Clusters.

80.What are three levels of management in storage?

- a)Storage Level Management
- b)Network Level Management
- c)Enterprise Level Management

81.What are the key activities in SAN management?

- a)Monitoring
- b)Configuring
- c)Controlling
- d)Troubleshooting
- e)Diagnosing

82.What is the difference between HBA and NIC?

HBA => Host bus adapters are used in storage based traffic while NIC (Network Interface Cards) are used in IP based LAN traffic.

83.What is the measuring unit of data activity?

Gigabits per second (Gb/ps)

84.What are the basic storage policies?

- a)Security and authentication
- b)Capacity, Content and quota management
- c)Quality of Service

85.What is bypass circuitry?

A circuit that automatically removes the storage device from the data path (FC device out of FC AL loop) when signaling is lost (this signal is called port by-pass signal).

86.How many connections are possible in Fabric topology?

2^{24} (24 bit address to the port), and the largest possible fabric will have 239 interconnected switches.

87.What is one of the constrain of using storage switch?

Latency

88.What is the difference between NAS and SAN?

NAS

Cables used in the n/w

n/w protocols (TCP/IP, IPx) and file sharing protocols (CIFS & NFS)

Lower TCO

Support heterogeneous clients

Slow

SAN

High-speed connectivity such as FC

Do not use n/w protocols because data request are not made over LAN

Higher TCO

Requires special s/w to provide access to heterogeneous clients

Fast

89.What is Jitter?

Jitter refers to any deviation in timing that a bit stream suffers as it traverses the physical medium and the circuitry on-board the end devices. A certain amount of deviation from the original signaling will occur naturally as serial bit stream propagates over fibre-optic or copper cabling.

Mainly caused by electro-magnetic interference

90.What is BER/Bit error rate?

Probability that a transmitted bit will be erroneously received is the measure of number of bits (erroneous) at the output of the receiver and dividing by the total number of bits in transmission.

91.What is WWPN?

WWPN is the 16bit character that is assigned to the port, SAN volume controller uses it to uniquely identify the fibre channel HBA that is installed in the host system.

92.What is connection allegiance?

Given multiple connections are established, individual command/response pair must flow over the same connection. This connection allegiance ensures that specific read or writes commands are fulfilled without any additional overhead of monitoring multiple connections and to see whether a particular request is completed.

93.What is burst Length?

The burst length is the number of bytes that the SCSI initiator sends to the SCSI target in the FCP_DATA sequence.

94.What is NAS in detail?

NAS or Network Attached Storage

"NAS is used to refer to storage elements that connect to a network and provide file access services to computer systems. A NAS Storage Element consists of an interface or engine, which implements the file services, and one or more devices, on which data is stored. NAS elements may be attached to any type of network. When attached to SANs, NAS elements may be considered to be members of the SAS (SAN Attached Storage) class of storage elements.

A class of systems that provide file services to host computers. A host system that uses network attached storage uses a file system device driver to access data using file access protocols such as NFS or CIFS. NAS systems interpret these commands and perform the internal file and device I/O operations necessary to execute them.

Though the NAS does speed up bulk transfers, it does not offload the LAN like a SAN does. Most storage devices cannot just plug into gigabit Ethernet and be shared - this requires a specialized file server the variety of supported devices is more limited. NAS has various protocols established for such needed features as discovery, access control, and name services.

95. Briefly list the advantages of SAN?

SANs fully exploit high-performance, high connectivity network technologies

SANs expand easily to keep pace with fast growing storage needs

SANs allow any server to access any data

SANs help centralize management of storage resources

SANs reduce total cost of ownership (TCO).

iSCSI fundamentals

iSCSI is a protocol defined by the Internet Engineering Task Force (IETF) which enables SCSI commands to be encapsulated in TCP/IP traffic, thus allowing access to remote storage over low

cost IP networks.

96. What advantages would using an iSCSI Storage Area Network (SAN) give to your organization over using Direct Attached Storage (DAS) or a Fibre Channel SAN?

- iSCSI is cost effective, allowing use of low cost Ethernet rather than expensive Fibre architecture.
- Traditionally expensive SCSI controllers and SCSI disks no longer need to be used in each server, reducing overall cost.
- Many iSCSI arrays enable the use of cheaper SATA disks without losing hardware RAID functionality.
- The iSCSI storage protocol is endorsed by Microsoft, IBM and Cisco, therefore it is an industry standard.
- Administrative/Maintenance costs are reduced.
- Increased utilisation of storage resources.
- Expansion of storage space without downtime.
- Easy server upgrades without the need for data migration.
- Improved data backup/redundancy.



0 comments:

Post a Comment

Enter your comment...

Comment as:

[Publish](#)

[Preview](#)

IoT/M2M Cloud Platform

Learn How to Quickly Deploy IoT. Ayla is Internet-Connected Products



[Older Post](#)



Copyright 2009 - [Storage](#)

Blogspot Theme designed by: Ray Creations, Ray Hosting.