



# Linux Administration

## Linux Interview Question and answers

Find Me on Google+

Thursday, December 15, 2011

### SAN and NAS Interview questions

#### What is a SAN?

**Ans:** SAN is short for Storage Area Network. It is a high-speed network of storage elements, similar in form and function to a LAN that establishes direct and indirect connections between multiple servers and multiple storage elements. The SAN is an extension of the server's storage bus

#### Q. What does a SAN do?

**Ans:** SANs create connectivity. SANs offer a method of attaching storage that improves data reliability, availability and performance

SAN overcomes traditional network bottlenecks by connecting in three ways:

- Server-to-storage (direct attached storage)
- Server-to-server (network attached storage)
- Storage-to-storage (SAN Attached Storage)

#### Q. Name some of the SAN topologies and Explain each of them ?

#### Followers



with Google Friend Connect

#### Members (9)



Already a member? [Sign in](#)

#### About Me



**g+** Prashanth P.Nair

Follow 126

View my complete profile

#### Blog Archive

- ▶ 2013 (1)
- ▶ 2012 (5)
- ▼ 2011 (9)
  - ▼ December (9)
    - MYSQL, Po stgres
    - SAN and NAS Interview

**Ans:** Point-to-point, arbitrated loop, and switched fabric topologies

### a) Point-to-Point

A point-to-point connection is the simplest topology. It is used when there are exactly two nodes and future expansion is not predicted. There is no sharing of the media, which allows the devices to use the total bandwidth of the link. A simple link initialization is needed before communications can begin.

### b) Arbitrary Loop

Our second topology is Fiber Channel Arbitrated Loop (FC-AL). FC-AL is more useful for storage applications. It is a loop of up to 126 nodes (NL\_Ports) that is managed as a shared bus. Traffic flows in one direction, carrying data frames and primitives around the loop with a total bandwidth of 400 MBps (or 200 MBps for a loop based on 2 Gbps technology).

### c) Switched Fabric Loop

It applies to switches and directors that support the FC-SW standard, that is, it is not limited to switches as its name suggests. A Fibre Channel fabric is one or more fabric switches in a single, sometimes extended, configuration. Switched fabrics provide full bandwidth per port compared to the shared bandwidth per port in arbitrated loop implementations.

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

**Ans:** 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.

### Q .What is FCP?

Total Pageviews

33183

Interview  
questions

Linux  
Kernel  
,ISO  
images

BIND -DNS

NFS,NIS,S  
AMBA,P  
OSTFIX,  
SENDM  
AIL,QM  
AIL

Perl and  
Shell  
Script

Linux File  
System

Process  
managme  
nt and  
DISK  
managme  
nt

Linux  
Basics

**Ans:**The Fibre Channel Protocol (FCP) is the interface protocol of SCSI on Fibre Channel. It is a gigabit speed network technology primarily used for Storage Networking. Fibre Channel is standardized in the T11 Technical Committee of the InterNational Committee for Information Technology Standards (INCITS), an American National Standard Institute (ANSI) accredited standards committee. It started for use primarily in the supercomputer field, but has become the standard connection type for storage area networks in enterprise storage. Despite its name, Fibre Channel signaling can run on both twisted-pair copper wire and fiber optic cables.

### **Q.What is iSCSI ?**

**Ans:** Internet SCSI (iSCSI) is a transport protocol that carries SCSI commands from an initiator to a target. It is a data storage networking protocol that transports standard Small Computer System Interface (SCSI) requests over the standard Transmission Control Protocol/Internet Protocol (TCP/IP) networking technology. iSCSI enables the implementation of IP-based storage area networks (SANs), enabling customers to use the same networking technologies — for both storage and data networks. As it uses TCP/IP, iSCSI is also well suited to run over almost any physical network. By eliminating the need for a second network technology just for storage, iSCSI has the potential to lower the costs of deploying networked storage.

### **Q.What is FCIP ?**

**Ans:** Fibre Channel over IP (FCIP) is also known as Fibre Channel tunneling or storage tunneling. It is a method to allow the transmission of Fibre Channel information to be tunnelled through the IP network. FCIP encapsulates Fibre Channel block data and subsequently transports it over a TCP socket. TCP/IP services are utilized to establish connectivity between remote SANs. Any congestion control and management, as well as data error and data loss recovery, is handled by TCP/IP services, and does not affect FC fabric services. The major point with FCIP is that it does not replace FC with IP, it

simply allows  
deployments of FC fabrics using IP tunneling

### **Q. What is iFCP**

**Ans:** Internet Fibre Channel Protocol (iFCP) is a mechanism for transmitting data to and from Fibre Channel storage devices in a SAN, or on the Internet using TCP/IP. iFCP gives the ability to incorporate already existing SCSI and Fibre Channel networks into the Internet. iFCP is able to be used in tandem with existing Fibre Channel protocols, such as FCIP, or it can replace them. Whereas FCIP is a tunneled solution, iFCP is an FCP routed solution. iFCP is a gateway-to-gateway protocol, and does not simply encapsulate FC block data. Gateway devices are used as the medium between the FC initiators and targets. As these gateways can either replace or be used in tandem with existing FC fabrics, iFCP could be used to help migration from a Fibre Channel SAN to an IP SAN, or allow a combination of both

### **Q. What is FICON ?**

**Ans:** FICON is a protocol that uses Fibre Channel as its physical medium. FICON channels are capable of data rates up to 200 MBps full duplex, they extend the channel distance (up to 100 km), increase the number of control unit images per link, increase the number of device addresses per control unit link, and retain the topology and switch management characteristics of ESCON.

### **Q. What is FICON address ?**

**Ans :** FICON generates the 24-bit FC port address field in yet another way. When communication is required from the FICON channel port to the FICON CU port, the FICON channel (using FC-SB-2 and FC-FS protocol information) will provide both the address of its port, the source port address identifier (S\_ID), and the

address of the CU port, the destination port address identifier (D\_ID) when the communication is from the channel N\_Port to the CU N\_Port.

### **Q. What is FSPF ?**

**Ans:** FSPF keeps track of the links on all switches in the fabric and associates a cost with each link. The cost is always calculated as being directly proportional to the number of hops. The protocol computes paths from a switch to all other switches in the fabric by adding the cost of all links traversed by the path, and choosing the path that minimizes the cost.

### **Q. How FSPF works**

**Ans:** The collection of link states (including cost) of all switches in a fabric constitutes

the topology database (or link state database). The topology database is kept in all switches in the fabric, and they are maintained and synchronized to each other. There is an initial database synchronization, and an update mechanism.

**82** Introduction to Storage Area Networks .The initial database synchronization is used when a switch is initialized, or when an ISL comes up. The update mechanism is used when there is a link state change. This ensures consistency among all switches in the fabric.

### **Q. What is Network Attached Storage (NAS) ?**

**Ans:** Network Attached Storage (NAS) is basically a LAN-attached file server that serves files using a network protocol such as Network File System (NFS). NAS is a term 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 engine that implements the file services (using access protocols such as NFS or CIFS), and one or more devices, on which data is stored. NAS elements may be attached to any type of network. From a SAN perspective, a SAN-attached NAS engine is treated just like any other server, but a NAS does not provide any of the activities that a server in a server-centric system typically provides, such as

e-mail, authentication, or file management.

### **Q. How is Fiber Channel Different from iSCSI?**

**Ans:** 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.

### **Q. What is Frames?**

**Ans:** Fibre Channel places a restriction on the length of the data field of a frame at 528

transmission words, which is 2112 bytes. (See Table 3-2 on page 52.) Larger

amounts of data must be transmitted in several frames. This larger unit that

consists of multiple frames is called a sequence. An entire transaction between

two ports is made up of sequences administered by an even larger unit called an

exchange.

A frame consists of the following elements:

\_ SOF delimiter

\_ Frame header

\_ Optional headers and payload (data field)

\_ CRC field

\_ EOF delimiter

### **Q. What is Loop address ?**

**Ans :**An NL\_Port, like an N\_Port, has a 24-bit port address. If no switch connection exists, the two upper bytes of this port address are zeroes (x'00 00') and referred to as a private loop. The devices on the loop have no connection with the outside world. If the loop is attached to a fabric and an NL\_Port supports a fabric login, the upper two bytes are assigned a positive value by the switch. We call this mode a public loop.

### **Q. What is LUN?**

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

**Q. What is LUN Masking?**

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

**Q. What is WWN?**

**Ans:** WWN is a 64bit address that is hard coded into a fiber channel HBA and this is used to identify individual port (N\_Port or F\_Port) in the fabric.

**Q. What is metaLUN?**

**Ans:** A metaLUN is a type of LUN whose maximum capacity can be the combined capacities of all the LUNs that compose it. The metaLUN feature lets you dynamically expand the capacity of a single LUN (base LUN) into a larger unit called a metaLUN. You do this by adding LUNs to the base LUN. You can also add LUNs to a metaLUN to further increase its capacity. Like a LUN, a metaLUN can belong to Storage Group, and can participate in Snap View, Mirror View and SAN copy sessions. MetaLUNs are supported only on CX-Series storage systems. A metaLUN may include multiple sets of LUNs and each set of LUNs is called a component. The LUNs within a component are striped together and are independent of other LUNs in the metaLUN.

**Q. What is a HBA?**

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

**Q. What is SAN fabric?**

**Ans:** SAN fabric is a hardware device that connects workstations and servers to storage devices in a SAN network. It uses the Fibre Channel switching technology to connect a server to a storage device. The SAN fabric offers a high-speed dedicated network including high availability features, very low latency, and high throughput.

**. What is zoning?**

**Ans:** 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.

**Q. What are the two major classification of zoning?**

**Ans:** Two types of zoning are

- a) Software Zoning
- b) Hardware Zoning

**Q.What are different levels of zoning?**

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

**Q. What are the 3 prominent characteristics of SAS Protocol?**

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

c) Port Selector

**Q.What is the purpose of disk array?**

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

**Q. What is disk array?**

**Ans:** 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.

**Q.What are the advantages of RAID?**

**Ans:** "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.

### **Q. How is a SAN managed?**

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

Santricity

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

### **Q. Which one is the Default ID for SCSI HBA?**

**Ans:** Generally the default ID for SCSI HBA is 7.

SCSI- Small Computer System Interface

**Q. What is the highest and lowest priority of SCSI?**

**Ans:** 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.

**Q.what is SRDF ?**

**Ans:** RDF (Symmetrix Remote Data Facility) is a family of EMC products that facilitates the data replication from one Symmetrix storage array to another through a Storage Area Network or IP network.SRDF logically pairs a device or a group of devices from each array and replicates data from one to the other synchronously or asynchronously. An established pair of devices can be split, so that separate hosts can access the same data independently (maybe for backup), and then resynchronized.

In synchronous mode (SRDF/S), the primary array waits until the secondary array has acknowledged each write before the next write is accepted, ensuring that the replicated copy of the data is always as current as the primary. However, the latency due to propagation increases significantly with distance.

Asynchronous SRDF (SRDF/A) transfers changes to the secondary array in units called delta sets, which are transferred at defined intervals. Although the remote copy of the data will never be as current as the primary copy, this method can replicate data over considerable distances and with reduced bandwidth requirements and minimal impact on host performance. Other forms of SRDF exist to integrate with clustered environments and to manage multiple SRDF pairs where replication of multiple devices must be consistent (such as with the data files and log files of a database application).

#### **Q. EMC control center 5.1**

**Ans:** The EMC Control Center Web Console uses [Port : 10799] data stored in the Repository to monitor your storage-attached network and manageControlCenter alerts remotely through a Web browser

#### **Q. What Is Emc Power path?**

**Ans:** EMC PowerPath is a server-resident software solution that enhances performance and information availability. It integrates multiple path I/O capabilities, automatic load balancing, and path failover functions into one comprehensive package for use on open server platforms connected to Symmetrix enterprise storage systems. PowerPath enables you to do more work in a shorter time so you can serve

more customers, run more applications, and exploit more business opportunities.

**Q. Define RAID? Which one you feel is good choice?**

**Ans:** 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.

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

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

## Q. What are the advantages of SAN?

**Ans:** 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.

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

**Ans:** 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.

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

**Ans:** If we consider any small business following are essentials components of SAN

- Fabric Switch
- FC Controllers
- JBOD's

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

**Ans:** 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.

**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.

**Q. What is the most critical component in SAN?**

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

**Q.Can you name some of the states of RAID array?**

**Ans:** 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

**Q. Name the features of SCSI-3 standard?**

\

**Ans:** QAS: Quick arbitration and selection

Domain Validation

CRC: Cyclic redundancy check

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

**Ans:** 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.

**Q. What is storage virtualization?**

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

### Q. What is virtualization?

**Ans:** 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.

### Q.What is Multipath I/O?

**Ans:** 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.

### Q. What is RAID?

**Ans:** 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.

### Q. What is stripe-unit-size?

**Ans:** It is a data distribution scheme that complements the way an operating system requests data. It is the 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.

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

**Ans:** Frame

### Q. What is the difference between mirroring, Routing and multipathing?

**Ans:** Redundancy Functions Relationships Role

Mirroring Generates 2 I/Os 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

**Q. Briefly list the advantages of SAN?**

**Ans:** 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.

**Q.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?**

**Ans:** 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.

Posted by [Prashanth P.Nair](#) at 3:30 PM



Recommend this on Google

## 2 comments:



**Ajesh kumar k.s** December 6, 2013 at 11:19 AM

Nice post

[Reply](#)



**Sunil Sony** March 17, 2015 at 2:01 AM

gud..rilly vry helpful sir :) nd i want to contact u in mail or in phone if u dnt mind

[Reply](#)

Enter your comment...

Comment as:

[Publish](#)

[Preview](#)

[Newer Post](#)

[Home](#)

[Older Post](#)

Subscribe to: [Post Comments \(Atom\)](#)

## Popular Posts

### Nagios Interview questions

Q.What is Nagios and how it Works ?. Ans: Nagios is an open source System and Network Monitoring application.Nagios runs on a serve...

### SAN and NAS Interview questions

What is a SAN? Ans: SAN is short for Storage Area Network. It is a high-speed network of storage elements, similar in form and funct...

### NFS,NIS,SAMBA,POSTFIX,SENDMAIL,QMAIL

Q. What is NFS? Ans: NFS stands for Network File System. NFS was originally developed by Sun Microsystems in the 1980's. NFS al...

### Open LDAP interview Questions and answers

Q. What is "LDIF" ?. Ans: The LDAP Interchange Format (LDIF) is a standard text file format for storing LDAP configuration inf...



### BIND -DNS

Q. What is BIND? What are the key Components of DNS ? Ans: BIND stands for Berkeley Internet Name Domain which is...

### Perl and Shell Script

Q. Can you explain the exit status of shell and commands under Linux / UNIX operating system? Ans . All UNIX and Linux command has a...

### Linux File System

Q. What command is used to list the contents of directory? Ans: ls ls -l Q.How to list a directory conta...

### Linux Kernel ,ISO images

Q.What is Kernel? Explain the task it performs. Ans: Kernel is used in UNIX like systems and is considered to be the heart of the o...

### Linux Basics

Q. What is the default Window system / Windows manager used in Linux? Ans :X.org Q. What is LILO? Ans: LILO stands for...

### Web Server

Q.How web Servers works? Ans :L ets say that you sitting at you computer ,surfing web http://www.eduinkerala.com So you type that...

Simple template. Template images by [luoman](#). Powered by [Blogger](#).