Comptia.Selftestengine.SG0-001.v2013-12-24.by.James.260q

Number: SG0-001
Passing Score: 800
Time Limit: 120 min
File Version: 14.5

http://www.gratisexam.com/

Exam Code: SG0-001

Exam Name: CompTIA Storage+ Powered by SNIA

Exam A

QUESTION 1
In the image of the motherboard given below, identify the area to which you connect a hard disk drive.

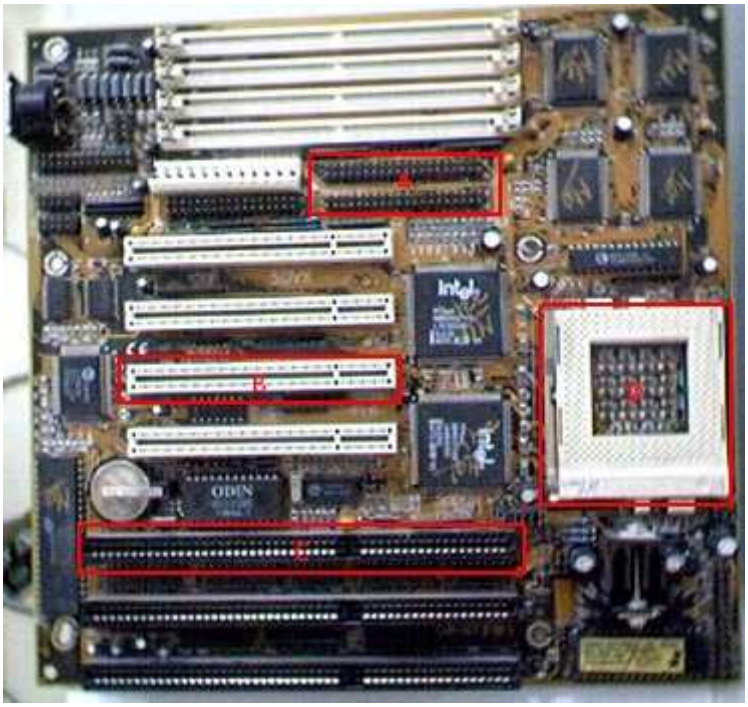

A.
B.
C.
D.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Hard disk drive uses 40 or 80 wire cables to connect to a motherboard that has an integrated IDE controller. A hard disk drive can also connect to a 40-pin connector on an IDE controller adapter using 40 or 80 wire cables.

Hard disk drives are magnetic storage devices that contain several disks, known as platters. These platters are attached to the spindle motor. The read/write heads are responsible for reading and writing to the platters. The read/write heads are attached to the head actuator, which is in charge of moving the heads around the platters.

What is a motherboard?
A motherboard is the physical arrangement in a computer that contains the computer's basic circuitry and components. On the typical motherboard, the circuitry is imprinted or affixed to the surface of a firm planar surface and usually manufactured in a single step. The most common motherboard design in desktop computers today is the AT, based on the IBM AT motherboard. A more recent motherboard specification, ATX, improves on the AT design. In both the AT and ATX designs, the computer components included in the motherboard are as follows:

- The microprocessor
- Memory
- Basic Input/Output System (BIOS)
- Expansion slot
- Interconnecting circuitry

The structure of a motherboard is shown in the figure below:

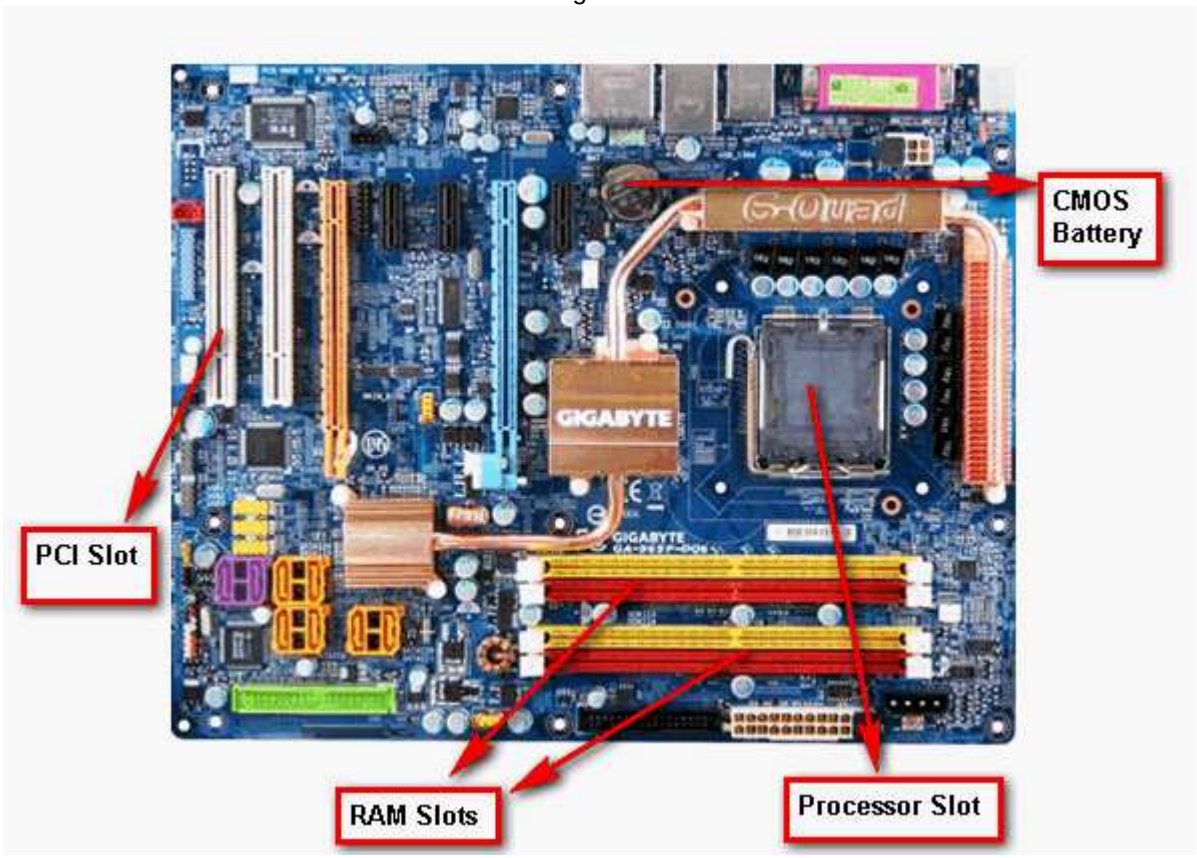

QUESTION 2
Which of the following are the hard disk components that describe the physical geometry of a newer hard disk drive?

Each correct answer represents a complete solution. Choose three.

A. Sectors per track
B. Number of cylinders
C. Number of heads
D. Landing zone

Correct Answer: ABC
Section: (none)
Explanation

Explanation/Reference:
The geometry of a hard disk is the organization of data on the platters. Geometry determines how and where data is stored on the surface of each platter, and thus the maximum storage capacity of the drive. The geometry of a hard disk is the actual physical number of heads, cylinders, and sectors used by the disk.

Answer option D is incorrect. A landing zone defines an unused cylinder as a parking space for the R/W heads. This is found in older hard disk drives.

Reference: http://www.pcguide.com/ref/hdd/geom/geomPhysical-c.html

QUESTION 3
Which of the following terms refers to the information that can be accessed on a disk drive by all the heads

without moving them?

A. Track
B. Sector
C. Cylinder
D. Platter

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
A hard disk comprises one or more circular platters, of which either or both surfaces are coated with a magnetic substance used for recording the data. For each surface, there is a read-write head that can process the recorded data. The surfaces are usually divided into concentric rings, called tracks, and these in turn are divided into sectors as shown in the image:

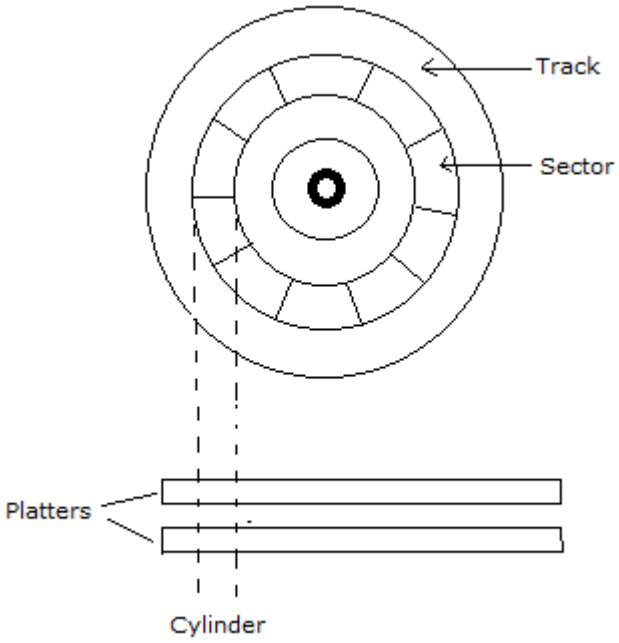

When the head for one surface is on a track, the heads for the other surfaces are also on the corresponding tracks. All the corresponding tracks taken together are called a cylinder.

QUESTION 4
Which of the following is an IP storage protocol that requires iSNS for discovery and management?

A. FIP
B. iFCP
C. IDE
D. iSCSI

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Internet Fibre Channel Protocol (iFCP) is a gateway-to-gateway protocol that provides fibre channel fabric

services to fibre channel devices over a TCP/IP network. iFCP allows existing fibre channel devices to be networked and interconnected over an IP based network at wire speeds. iFCP requires iSNS for discovery and management. iSNS is a protocol and mechanism for intelligent discovery of storage devices in an IP network.

Answer option A is incorrect. FCoE Initialization Protocol (FIP) is a protocol that enables the discovery and instantiation, and maintenance of FCoE devices.

Answer option C is incorrect. Integrated Drive Electronics (IDE) is a type of hardware interface widely used to connect hard disks, CD-ROMs, and tape drives to a PC.

Answer option D is incorrect. iSCSI is a SCSI transparent protocol for mapping of Block-oriented storage data over TCP/IP networks.

QUESTION 5
Which of the following drives is considered as a "dumb" drive, as the drive controller has all the intelligence?

A. SATA
B. PATA
C. ATA
D. SCSI

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
The SCSI drive is considered as a "dumb" drive, as it has a SCSI controller built in that has all the intelligence. The drive controller bears the responsibility for transferring data between the SCSI bus and the computer's input/output bus.

Answer options C, B, and A are incorrect. These drives are not considered as "dumb" drives.

QUESTION 6
Which standards are the responsibilities of the Internet Engineering Task Force?

Each correct answer represents a complete solution. Choose two.

A. IP over Fibre Channel
B. SNMP
C. NDMP
D. 802.3

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
Internet Engineering Task Force (IETF) is an open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and coordination of the

operation and management of the Internet. It works towards introducing procedures for new technologies on the Internet. IETF specifications are released in Requests for Comments (RFCs). IETF is responsible for a wide variety of TCP/IP (Transmission Control Protocol/Internet Protocol) standards, such as SNMP, policy for QoS (quality of service), IP over Fibre Channel, etc.

Answer options C and D are incorrect. NDMP stands for Network Data Management Protocol. It is a protocol invented by the NetApp and Legato companies. It is meant to transport data between NAS devices and backup devices. The 802.3 standard belongs to IEEE.

QUESTION 7
Why are terminators used with SCSI buses?

A. To reflect signals
B. To attach new disks
C. To remove devices
D. To absorb signals

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
SCSI bus requires correct electrical termination at both ends in order to function properly. Terminators are basically electrical resistors that absorb signals that reach the end of the bus. Without them, signals could reflect back onto the bus, leading to a confusing mix of actual and reflected signals. Bus terminator with top cover removed is shown below:

Reference: http://www.pcguide.com/ref/hdd/if/scsi/cablesTermination-c.html

QUESTION 8
Which transport method requires correct electrical termination on both ends?

A. ATAPI
B. PCI bus
C. SCSI bus
D. ISA

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
SCSI bus requires correct electrical termination at both ends in order to function properly. Proper termination make sure that the signal travelling down the SCSI bus doesn't reflect back, a situation that cause a variety of problems including "ghosted" SCSI devices, data errors, and other anomalies.

Reference: http://www.windowsitpro.com/article/windows-client/scsi-termination

QUESTION 9
How much data can be stored in a sector of AS/400 systems?

A. 512 bytes
B. 1024 bytes
C. 528 bytes
D. 520 bytes

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
AS/400 systems utilize a sector size of 520 bytes, whereas SCSI and ATA typically utilize a sector size of 512 bytes.

Answer options B and C are incorrect. AS/400 systems do not use a sector size of 1024 and 528 bytes.

QUESTION 10
Which of the following quantities is determined by the clock rate?

A. Storage capacity
B. Link rate
C. Bandwidth
D. Throughput

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
The link rate is determined by clock rate, and can be expressed in bits/second, e.g., 1Gb/s. It is not affected by distance, media, or latency.

Answer option C is incorrect. Bandwidth is determined by encoding.

Answer option D is incorrect. Throughput is the measurement of the usable data traffic in one direction only.

Answer option A is incorrect. The storage capacity is determined by the storage limit of a particular device.

QUESTION 11
How many heads per platter are required in a hard disk platter?

A. 2
B. 5
C. 4
D. 1

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
A hard disk platter (or disk) is a component of a hard disk drive. It is the circular disk on which the magnetic data is stored. The rigid nature of the platters in a hard drive is what gives them their name. Hard drives typically have several platters that are mounted on the same spindle. A platter can store information on both

sides, requiring two heads per platter.

QUESTION 12
Which of the following is an open protocol used to control data backup and recovery communications between primary and secondary storage in a heterogeneous network environment?

A. iFCP
B. FCIP
C. NDMP
D. iSCSI

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Network Data Management Protocol (NDMP) is an open protocol used to control data backup and recovery communications between primary and secondary storage in a heterogeneous network environment.

NDMP specifies a common architecture for the backup of network file servers and enables the creation of a common agent that a centralized program can use to back up data on file servers running on different platforms. By separating the data path from the control path, NDMP minimizes demands on network resources and enables localized backups and disaster recovery. With NDMP, heterogeneous network file servers can communicate directly to a network-attached tape device for backup or recovery operations.

Answer option D is incorrect. iSCSI is an abbreviation of Internet Small Computer System Interface, an Internet Protocol (IP)-based storage networking standard for linking data storage facilities. By carrying SCSI commands over IP networks, iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances. iSCSI can be used to transmit data over local area networks (LANs), wide area networks (WANs), or the Internet, and can enable location-independent data storage and retrieval.

The protocol allows clients (called initiators) to send SCSI commands (CDBs) to SCSI storage devices (targets) on remote servers. It is a popular storage area network (SAN) protocol, allowing organizations to consolidate storage into data center storage arrays while providing hosts (such as database and web servers) with the illusion of locally-attached disks. Unlike traditional Fiber Channel, which requires special-purpose cabling, iSCSI can be run over long distances using existing network infrastructure.

Answer option A is incorrect. Internet Fibre Channel Protocol (iFCP) is a gateway to gateway network protocol standard that provides Fibre Channel fabric functionality to fibre channel devices over an IP network. The iFCP protocol enables the implementation of fibre channel functionality over an IP network, within which the fibre channel switching and routing infrastructure is replaced by IP components and technology. The primary objective of iFCP is to allow existing fibre channel devices to be networked and interconnected over an IP based network at wire speeds. The iFCP protocol layer's main function is to transport Fibre Channel frame images between Fibre Channel ports attached both locally and remotely.

Answer option B is incorrect. Fibre Channel over IP (FCIP) is an IP-based storage networking technology, developed by the IETF and defined in RFC 3821. FCIP is also known as Fibre Channel tunneling or storage tunneling. FCIP enables the transmission of Fibre Channel information by tunneling data between storage area network facilities over IP networks. This capacity facilitates data sharing over a geographically distributed enterprise. FCIP is a key technology that is expected to help bring about rapid development of the storage area network market by increasing the capabilities and performance of storage data transmission.

Reference: http://searchstorage.techtarget.com/sDefinition/0,,sid5_gci837192,00.html

QUESTION 13
Which of the following allows for almost six times more data storage than on a DVD?

A. Floppy drive

B. CD drive
C. Switch
D. Blu-ray disk

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Blu-ray Disc, also known as Blu-ray or BD, is an optical disc storage medium. It is designed to supersede the standard DVD format for storing high-definition video, PlayStation 3 games, and other data, with up to 25 GB per single layered, and 50 GB per dual layered disc. Blu-ray Disc was developed by the Blu-ray Disc Association, a group representing makers of consumer electronics, computer hardware, and motion pictures. The disc has the same physical dimensions as standard DVDs and CDs. It is called Blu-ray Disc because of the blue-violet laser which is used to read the disc. Blu-ray uses a shorter wavelength, a 405 nm blue-violet laser, and allows for almost six times more data storage than on a DVD.

QUESTION 14
Which of the following cables support transmission speed of 1000Mbps?

Each correct answer represents a complete solution. Choose two.

A. CAT 5
B. CAT 5e
C. CAT 6
D. CAT 3

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
Cat 5e and Cat 6 cables support transmission speeds of 1000Mbps.

Answer options D and A are incorrect. Cat 3 and Cat 5 cables do not support transmission speeds of 1000Mbps.

QUESTION 15
Which of the following connectors use the push-pull mechanism to make the connection?

Each correct answer represents a complete solution. Choose two.

A. MT-RJ
B. SC
C. ST
D. LC

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
The subscriber connector (SC) and the Lucent connector (LC) use the push-pull mechanism to make the connection.

A subscriber connector (SC) is a fiber-optic connector used with multimode fiber. It is a square shaped

connector used for terminating fiber optic cables. SC connectors have a push-pull latching mechanism to provide quick insertion and removal while also ensuring a positive connection.


SC

Answer options C and A are incorrect. The ST and MT-RJ connectors do not use the push-pull mechanism to make the connection.

Straight tip (ST) connector:
A straight tip (ST) connector is a fiber-optic connector used with multimode fiber. An ST connector has a 2.5mm shaft and bayonet locking ring, and allows quick connect and disconnect of 125 micron multi-mode fiber.


ST

MT-RJ connector:
The MT-RJ connector is the most recent type of small form factor fiber optic connector. The MT-RJ fiber connector is less than half the size of the SC duplex connector and transceiver, so it doubles the port density of fiber-optic LAN equipment. The connector is a 2-fiber connector and takes up no more room than an RJ-45 jack.

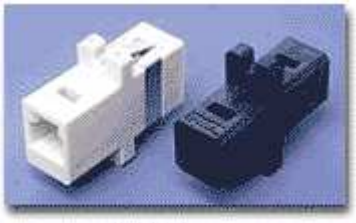
MT-RJ

Reference: http://en.wikipedia.org/wiki/Optical_fiber_connector

QUESTION 16
What advantage does Cat 5 cables have over fiber-optic cables?

A. Flexibility
B. Support for higher bandwidth
C. Immunity from EMI
D. Faster transmission of data

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Among the given options, the advantage that Cat 5 cables have over fiber-optic cables is flexibility. Category 5 cable is a twisted pair high signal integrity cable type. It is also known as Cat 5. Most cables of this type are unshielded, but some are shielded, used for high signal integrity. Cat 5 cables can transfer data at speeds of 100 Megabits per second. A Cat 5 cable typically has three twists per inch of each twisted pair of 24 gauge copper wires within it. The twisting of the cable can reduce electrical interference and crosstalk. Cat 5 cables are used in structured cabling for computer networks, such as Fast Ethernet.

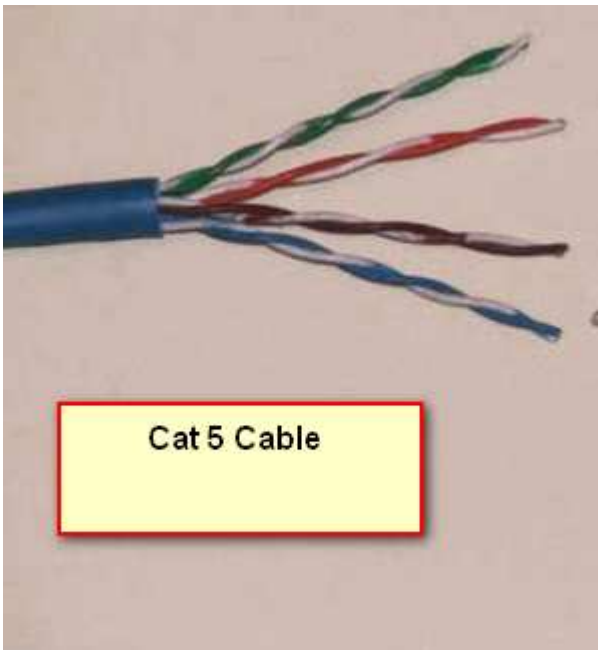

Note: The maximum segment length that a 1000BASET network supports is 100 meters. Type 1000BASET network uses 5-level encoding and CAT5 UTP media. It can provide data transmission speeds of up to 1000 megabits per second.

Answer options D, B, and C are incorrect. These are the advantages of fiber-optic cables over Cat 5 cables.

Fiber-optic cable:
A fiber-optic cable is used for high-speed, high-capacity data transmission. It uses optical fibers to carry digital data signals in the form of modulated pulses of light. The prime features of a fiber-optic cable are as follows:

It supports greater signal bandwidth transmission.
It is immune to electromagnetic interference (EMI).
It can transmit undistorted signals over great distances.

QUESTION 17
You have installed a small network at an accounting office. All the computers have gigabit Ethernet NICs. Your router works at 1000 Mbps. Yet, when you measure the Internet speed, you see it is at a max of 100 Mbps. What should you do to fix this?

A. Replace the router.
B. Replace the cable with fiber-optic.
C. Replace the cable with Category 6 cable.
D. Replace the cable with Category 5 cable.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Category 6 (Cat 6) cable gets speeds of up to 1000 Mbps (1 gigabit).

Category 6, popularly known as Cat 6, is a cable standard for networking. It has a maximum allowed length of 100 meters (330 ft) . It has several improved features such as compatibility and crosstalk and system noise reduction due to which it does not hinder the speed and performance.

Answer option D is incorrect. Category 5 cables can only get speeds of up to 100 Mbps.

Answer option A is incorrect. As the router supports the required speed, it is not a router issue.

Answer option B is incorrect. Fiber-optic is very expensive and not necessary in a small business network. And the current Ethernet NICs will not work with fiber-optic cable.

Reference: http://en.wikipedia.org/wiki/Category_6_cable

QUESTION 18
Which of the following statements are true about a Cat 5 cable?

Each correct answer represents a complete solution. Choose three.

A. It has three twists per inch of each twisted pair.
B. It is capable of transferring data at 100Mbps.
C. It is an optical fiber cable.
D. It uses a 24-gauge copper wire.

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
Category 5 cable is a twisted pair high signal integrity cable type. It is also known as Cat 5. Most cables of this type are unshielded, but some are shielded, used for high signal integrity. Cat 5 cables can transfer data at speeds of 100 Megabits per second. A Cat 5 cable typically has three twists per inch of each twisted pair of 24 gauge copper wires within it. The twisting of the cable can reduce electrical interference and crosstalk. Cat 5 cables are used in structured cabling for computer networks, such as Fast Ethernet.

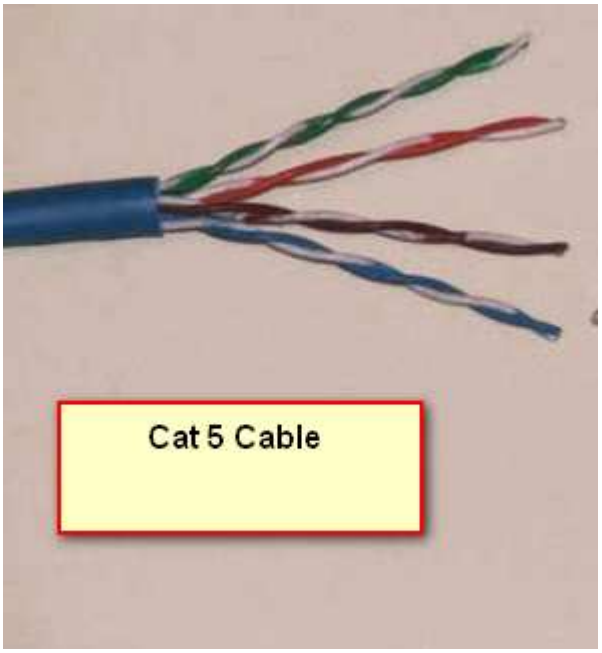

Note: The maximum segment length that a 1000BASET network supports is 100 meters. Type 1000BASET network uses 5-level encoding and CAT5 UTP media. It can provide data transmission speeds of up to 1000 megabits per second.

Answer option C is incorrect. A Cat 5 cable is not an optical fiber cable.

QUESTION 19
What is the maximum data transfer rate supported by SAS-2?

A. 6Gbps
B. 3Gbps
C. 2Gbps
D. 12Gbps

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
The speed of the first-generation SAS (SAS 1) link is 3Gbps. The second-generation SAS (SAS 2) link speed doubles the transfer rate to 6Gbps. SAS 2 eliminates the distinction between fan out and edge expanders by replacing them with self-configuring expanders. SAS 2 enables zoning for enhanced resource deployment flexibility, security, and data traffic management. SAS 2 is also backward compatible with SAS 1.

QUESTION 20
Which of the following types of connectors are used for unshielded twisted-pair connections?

Each correct answer represents a complete solution. Choose two.

A. RJ-11
B. DB-9
C. DB-25
D. RJ-45

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
RJ-45 and RJ-11 connectors are used for unshielded twisted-pair connections. The RJ-45 connector is similar to an RJ-11 telephone connector, but is larger in size, because it has eight conductors; an RJ-11 connector only has four conductors. An RJ-11 connector is the same connector that is used on a telephone wire, while an RJ-45 connector is used for copper wire Ethernet connections.

Answer options B and C are incorrect. The DB-9 and DB-25 connectors are not used for unshielded twisted-pair connections.

DB-9 connector:
A DB-9 connector is used by the serial interface. It is a 9-pin, two rows, male connector, on the back of a PC used for connecting mouse, modem, and other serial devices to the personal computer. It is also known as serial port.

DB-25 connector:
A DB-25 connector is a 25-pin, two row female connector, on the back of a PC used for connecting printers, scanners etc. It is also called parallel port. The parallel cable interface uses a DB-25 connector.

QUESTION 21
Which of the following ports allows any device to connect to any port, with selection of the actual port type being automatic?

A. U_Port
B. E_Port
C. L_Port

D. NL_Port

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Universal Port is a port that can operate as a G_Port, E_Port, F_Port or FL_Port. Most switches and related fabric devices contain Universal Ports to allow any device to connect to any port, with selection of the actual port type being automatic.

Answer option B is incorrect. The "Expansion" port within a Fibre Channel switch connects to another Fibre Channel switch or bridge device via an inter-switch link. E_Ports are used to link Fibre channel switches to form a multi-switch fabric.

Answer option C is incorrect. A "Loop" port is capable of performing arbitrated loop functions and protocols. NL_Ports and FL_Ports are examples of loop-capable ports.

Answer option D is incorrect. A "Node Loop" port is capable of arbitrated loop functions and protocols. An NL_Port connects via an arbitrated loop to other NL_Port and at most a single FL_Port. NL_Ports handle creation, detection, and flow of message units to and from the connected systems. NL_Ports are end ports in virtual point-to-point links through a fabric.

QUESTION 22
Which of the following ports can determine operating mode at switch port initialization?

A. N_Port
B. G_Port
C. E_Port
D. F_Port

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
A "Generic" Port can operate as either an E_Port or an F_Port. A G_Port can determine operating mode at switch port initialization, F_Port when an N_Port attachment is determined, E_Port when an E_Port attachment is determined.

Answer option C is incorrect. The "Expansion" port within a Fibre Channel switch connects to another Fibre Channel switch or bridge device via an inter-switch link. E_Ports are used to link Fibre channel switches to form a multi-switch fabric.

Answer option D is incorrect. The "Fabric" port within a Fibre Channel fabric switch provides a point-to-point link attachment to a single N_Port. F_Ports are intermediate ports in virtual point-to-point links between end ports, for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch.

Answer option A is incorrect. A "Node" port connects via a point-to-point link to either a single N_Port or a single F_Port. N_Ports handle creation, detection, and flow of message units to and from the connected systems. N_Ports are end ports in virtual point-to-point links through a fabric, for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch.

QUESTION 23
Which of the following are the methods used for trunking of switches?

Each correct answer represents a complete solution. Choose two.

A. ISL
B. 802.1Q
C. STP
D. FTP

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
ISL and 802.1Q are the two methods used for trunking of switches. Inter-Switch Link (ISL) is a trunking method developed by Cisco to use for Ethernet and Token Ring trunk connections. Most of the Cisco switches and routers that support trunking also support ISL except some older switches such as Catalyst 4000 switches. ISL encapsulates the original frame by adding a 26-byte header and a 4-byte trailer.

802.1Q is a trunking method developed by IEEE. It inserts a 4-byte field into the original Ethernet frame and recomputes the FCS. Unlike the ISL trunking method, it allows trunks between different vendors' devices. This trunk supports two types of frames, i.e., tagged and untagged. Frames that are untagged do not carry any VLAN identification information. Untagged frames are simple Ethernet frames.

Answer options C and D are incorrect. FTP and STP are protocols, and they are not the methods used for trunking.

QUESTION 24
Which of the following ports provides a point-to-point link attachment to a single N_Port?

A. E_Port
B. F_Port
C. NL_Port
D. L_Port

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
The "Fabric" port within a Fibre Channel fabric switch provides a point-to-point link attachment to a single N_Port. F_Ports are intermediate ports in virtual point-to-point links between end ports, for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch.

Answer option A is incorrect. The "Expansion" port within a Fibre Channel switch connects to another Fibre Channel switch or bridge device via an inter-switch link. E_Ports are used to link Fibre channel switches to form a multi-switch fabric.

Answer option D is incorrect. A "Loop" port is capable of performing arbitrated loop functions and protocols. NL_Ports and FL_Ports are examples of loop-capable ports.

Answer option C is incorrect. A "Node Loop" port is capable of arbitrated loop functions and protocols. An NL_Port connects via an arbitrated loop to other NL_Port and at most a single FL_Port. NL_Ports handle creation, detection, and flow of message units to and from the connected systems. NL_Ports are end ports in virtual point-to-point links through a fabric.

QUESTION 25
What is the approximate maximum network bandwidth for Fibre Channel network in 8Gb/s SAN?

A. 800MB/s Full Duplex
B. 1600MB/s Full Duplex

C. 200MB/s Full Duplex
D. 400MB/s Full Duplex

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
The approximate maximum network bandwidth for Fibre Channel network in 8Gb/s SAN is 1600MB/s Full Duplex.

Answer option C is incorrect. The approximate maximum network bandwidth for Fibre Channel network in 1Gb/s SAN is 200MB/s Full Duplex.

Answer option D is incorrect. The approximate maximum network bandwidth for Fibre Channel network in 2Gb/s SAN is 400MB/s Full Duplex.

Answer option A is incorrect. The approximate maximum network bandwidth for Fibre Channel network in 4Gb/s SAN is 800MB/s Full Duplex.

QUESTION 26
Which of the following technologies provides adequate humidity control to office buildings?

A. HVAC
B. IDE
C. NDMP
D. ISL

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
HVAC stands for the closely related functions of "Heating, Ventilating, and Air Conditioning". It is the technology of indoor or automotive environmental comfort. The HVAC system design is a major sub-discipline of mechanical engineering, based on the principles of thermodynamics, fluid mechanics, and heat transfer. HVAC is particularly important in the design of medium to large industrial and office buildings.

QUESTION 27
Which of the following statements are true regarding passive cooling?

Each correct answer represents a complete solution. Choose three.

A. It either shields buildings from direct heat gain or transfers excess heat outside.
B. It uses the passive solar technology for the purpose of cooling.
C. It uses solar panels to collect heat and electrically driven pumps or fans to transport the heat.
D. It contains various elements such as overhangs, awnings, and eaves that shade from high angle summer sun while allowing winter sun to enter the building.

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
Passive cooling uses the passive solar technology for the purpose of cooling. It either shields buildings from direct heat gain or transfers excess heat outside. It contains various elements such as overhangs, awnings, and eaves that shade from high angle summer sun while allowing winter sun to enter the building. The

processes of ventilation and conduction are used for transferring excess heat. A radiant heat barrier can block up to 95% of radiant heat transfer through the roof.

Answer option C is incorrect. Active cooling uses solar panels to collect heat and electrically driven pumps or fans to transport the heat.

QUESTION 28
You need to implement the 8U server rack in your organization. What is the height of the 8U server rack?

A. 21 inches
B. 14 inches
C. 31.5 inches
D. 26.25 inches

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
The 8U server rack has a height of 14 inches.

Answer option A is incorrect. The 12U server rack has a height of 21 inches.

Answer option D is incorrect. The 15U server rack has a height of 26.25 inches.

Answer option C is incorrect. The 18U server rack has a height of 31.5 inches.

QUESTION 29
Which of the following is a globally unique 64-bit identifier assigned to each Fibre Channel port?

A. WWNN
B. WWPN
C. NPIV
D. BBC

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
In SAN, a World Wide Port Name, WWPN, or WWpN, is a World Wide Name assigned to a port in a Fibre Channel fabric. It performs a function equivalent to the MAC address in the Ethernet protocol, as it is supposed to be a unique identifier in the network. WWPN is a globally unique 64-bit identifier. It is assigned to each Fibre Channel port.

Answer option A is incorrect. Worldwide node name is a Node_Name that is worldwide unique. A World Wide Node Name, WWNN, or WWnN, is a World Wide Name assigned to a node (an endpoint, a device) in a Fibre Channel fabric. It is valid for the same WWNN to be seen on many different ports (different addresses) on the network, identifying the ports as multiple network interfaces of a single network node. WWNN is a globally unique 64-bit identifier. It is assigned to each Fibre Channel node process.

Answer option C is incorrect. N_Port ID Virtualization, or NPIV, is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port.

Answer option D is incorrect. Buffer credit, also called buffer-to-buffer credit (BBC), is used as a flow control method by Fibre Channel technology and represents the number of frames a port can store.

Reference: http://www.knowledgetransfer.net/dictionary/Storage/en/World_Wide_Port_Name.htm

QUESTION 30
Which of the following storage industry terms is defined as the point-to-point physical connection from one element of a Fibre Channel fabric to the next?

A. Worldwide node name
B. Link
C. Oversubscription
D. Buffer-to-buffer credit

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
A link is defined as a physical connection (electrical or optical) between two nodes of a network. It can also be defined as the point-to-point physical connection from one element of a Fibre Channel fabric to the next.

Answer option C is incorrect. In a SAN (storage area network) switching environment, oversubscription is the practice of connecting multiple devices to the same switch port in order to optimize switch use. Each SAN port can support a particular communication speed and a Fibre Channel switch may offer 1 Gb, 2 Gb, or 4 Gb FC ports. However, because ports are rarely run at their maximum speed for a prolonged period, multiple slower devices may fan in to a single port to take advantage of unused capacity. For instance, a single storage server may not be able to sustain 4 Gbps to a switch port, so two 2 Gb servers or four 1 Gb servers may all be aggregated to that 4 Gb switch port.

Answer option D is incorrect. Buffer credit, also known as buffer-to-buffer credit (BBC) is used as a flow control method by Fibre Channel technology and represents the number of frames a port can store. It is used to permit data communication in a Fibre Channel storage area network (SAN) where there are long spans of fiber optic cable.

In buffer-credit flow control, the source and destination sets the number of unacknowledged frames (buffer credits) permitted to accumulate before the source stops sending data. This is a power of 2, such as 4, 8, or 16. In some SANs it can range higher, but it is rare to have more than 256 buffer credits because of cost constraints. A counter at the source keeps track of the number of buffer credits. Each time a frame is sent by the source, the counter increments by 1. Each time the destination receives a frame, it sends an acknowledgment back to the source, which decrements the counter by 1. If the number of buffer credits reaches the maximum, the source stops transmission until it receives the next acknowledgement from the destination. This helps in preventing loss of frames that may result if the source races too far ahead of the destination.

Answer option A is incorrect. Worldwide node name is a Node_Name that is worldwide unique. A World Wide Node Name, WWNN, or WWnN, is a World Wide Name assigned to a node (an endpoint, a device) in a Fibre Channel fabric. It is valid for the same WWNN to be seen on many different ports (different addresses) on the network, identifying the ports as multiple network interfaces of a single network node.

Reference: http://www.snia.org/education/dictionary/l

QUESTION 31
Which of the following is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port?

A. Worldwide node name
B. Worldwide port name
C. N-port ID Virtualization
D. Buffer-to-buffer credit

Correct Answer: C
Section: (none)

Explanation

Explanation/Reference:
NPIV is defined as the ability for a single physical Fibre Channel node or switch for supporting more than one Nx_Port on a single point-to-point link.

N_Port ID Virtualization, or NPIV, is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port. It allows multiple Fibre Channel initiators to occupy a single physical port, easing hardware requirements in Storage Area Network design, especially where virtual SANs are called for.

Answer option D is incorrect. Buffer credit, also known as buffer-to-buffer credit (BBC) is used as a flow control method by Fibre Channel technology and represents the number of frames a port can store. It is used to permit data communication in a Fibre Channel storage area network (SAN) where there are long spans of fiber optic cable.

In buffer-credit flow control, the source and destination sets the number of unacknowledged frames (buffer credits) permitted to accumulate before the source stops sending data. This is a power of 2, such as 4, 8, or 16. In some SANs it can range higher, but it is rare to have more than 256 buffer credits because of cost constraints. A counter at the source keeps track of the number of buffer credits. Each time a frame is sent by the source, the counter increments by 1. Each time the destination receives a frame, it sends an acknowledgment back to the source, which decrements the counter by 1. If the number of buffer credits reaches the maximum, the source stops transmission until it receives the next acknowledgement from the destination. This helps in preventing loss of frames that may result if the source races too far ahead of the destination.

Answer option A is incorrect. Worldwide node name is a Node_Name that is worldwide unique. A World Wide Node Name, WWNN, or WWnN, is a World Wide Name assigned to a node (an endpoint, a device) in a Fibre Channel fabric. It is valid for the same WWNN to be seen on many different ports (different addresses) on the network, identifying the ports as multiple network interfaces of a single network node.

Answer option B is incorrect. In SAN, a World Wide Port Name, WWPN, or WWpN, is a World Wide Name assigned to a port in a Fibre Channel fabric. It performs a function equivalent to the MAC address in the Ethernet protocol, as it is supposed to be a unique identifier in the network.

QUESTION 32
Which of the following is the system component that originates an I/O command over an I/O bus or network?

A. Alias
B. Target
C. Initiator
D. SAN Fabric

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Initiator is the system component that originates an I/O command over an I/O bus or network. It helps in building and issuing a Command Descriptor Block (CDB), forwarding the CDB to a SCSI Target device, handling any interaction associated with that particular command (such as receiving and acknowledging data for a READ command), and checking/accepting a response (Status) from the SCSI Target device.

I/O adapters, network interface cards, and intelligent controller device I/O bus control ASICs are typical initiators.

Answer option B is incorrect. Target is the system component that receives a SCSI I/O command.

Answer option A is incorrect. Alias is an alternative name for an entity, sometimes used for creating names that are more easily human readable.

Answer option D is incorrect. A SAN fabric is a hardware device that connects workstations and servers to storage devices in a SAN network.

QUESTION 33
Which of the following uses the Fibre Channel switching technology to connect a server to a storage device?

A. SAN fabric
B. Initiator
C. Target
D. Oversubscription

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
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. An organization can get more advantageous features, such as data replication, data sharing, and centralized backups by consolidating storage devices into the same SAN fabric.

Answer option B is incorrect. Initiator is the system component that originates an I/O command over an I/O bus or network. It helps in building and issuing a Command Descriptor Block (CDB), forwarding the CDB to a SCSI Target device, handling any interaction associated with that particular command (such as receiving and acknowledging data for a READ command), and checking/accepting a response (Status) from the SCSI Target device.

I/O adapters, network interface cards, and intelligent controller device I/O bus control ASICs are typical initiators.

Answer option C is incorrect. A target is the system component that receives a SCSI I/O command. The iSCSI specification refers to a storage resource located on an iSCSI server as a target. An iSCSI target is often a dedicated network-connected hard disk storage device, but may also be a general-purpose computer. Since with initiators, software for providing an iSCSI target is available for most mainstream operating systems.

Answer option D is incorrect. In a SAN (storage area network) switching environment, oversubscription is the practice of connecting multiple devices to the same switch port in order to optimize switch use. Each SAN port can support a particular communication speed and a Fibre Channel switch may offer 1 Gb, 2 Gb, or 4 Gb FC ports. However, because ports are rarely run at their maximum speed for a prolonged period, multiple slower devices may fan in to a single port to take advantage of unused capacity. For instance, a single storage server may not be able to sustain 4 Gbps to a switch port, so two 2 Gb servers or four 1 Gb servers may all be aggregated to that 4 Gb switch port.

QUESTION 34
Which of the following is unique for each device in a SCSI chain?

A. I/O address
B. LUN
C. SCSI Identification
D. Master/Slave setting

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

The SCSI ID is unique for each device in a SCSI chain. Each device on a SCSI bus must have a SCSI ID different from that of other devices. The IDs assigned to the devices on a SCSI bus are important to the performance of the computer because IDs dictate the priority in which a device gets access to the bus. Narrow SCSI buses that support eight devices use ID 0 to 7, and wide SCSI buses that support 16 devices use ID 0 to 15. SCSI IDs of wide SCSI buses are divided into two groups, 0 to 7 and 8 to 15. 0 to 7 has priority over 8 to 15. Higher ID numbers in the group have higher priority within that group. Therefore, SCSI ID 7 has top priority. The host adapter generally uses highest ID (7 or 15) and hard disk drives typically use IDs 0 and 1.

Answer option A is incorrect. I/O address is a communication port between a device and the CPU. The CPU needs a memory address, known as Input/Output (I/O) address, to communicate with any peripheral device. I/O address is a hexadecimal number that the CPU uses to identify a device. I/O address allows the CPU to send instructions to devices installed on the bus slot of a computer.

Resources such as I/O addresses, IRQs, and DMAs are configurable aspects of communication between devices inside a PC. Whenever a component, such as a sound card or internal modem is installed in a PC, its I/O address, IRQ, and DMA channels must be correctly configured.

Answer option B is incorrect. A logical unit number (LUN) is a sub classification within a SCSI ID. It enables multiple devices to use a single ID.

Answer option D is incorrect. Master/Slave setting is used on IDE drives. The Master/Slave setting is accomplished by jumper setting of pins on the hard disk/CD-ROM drive.

SCSI
Small Computer System Interface (SCSI) is the second most popular drive interface in use today after the Integrated Drive Electronics (IDE) interface. SCSI is faster than IDE and supports more devices. SCSI devices, such as hard disk drive and CD-ROM drive, are better suited in a network environment in which many users access shared drives simultaneously. SCSI has three standards: SCSI-1, SCSI-2, and SCSI-3.

Reference: TechNet, Contents: "Managing File Systems and Drives"

QUESTION 35
Which of the following is a configuration management element of fabric that helps in restricting storage access for the entire host?

A. Domain ID
B. Zone
C. SCSI ID
D. NPIV

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Zone is a configuration management element of fabric that helps in restricting storage access for the entire host. In a storage area network (SAN), zoning is the allocation of resources for device load balancing and for selectively allowing access to data only to certain users.

Zones, zone aliases, and zone sets permit logical grouping of ports and storage devices within a storage area network. The following section describes zoning concepts and elements.

In a storage area network, a zone is a logical grouping of ports for forming a virtual private storage network.

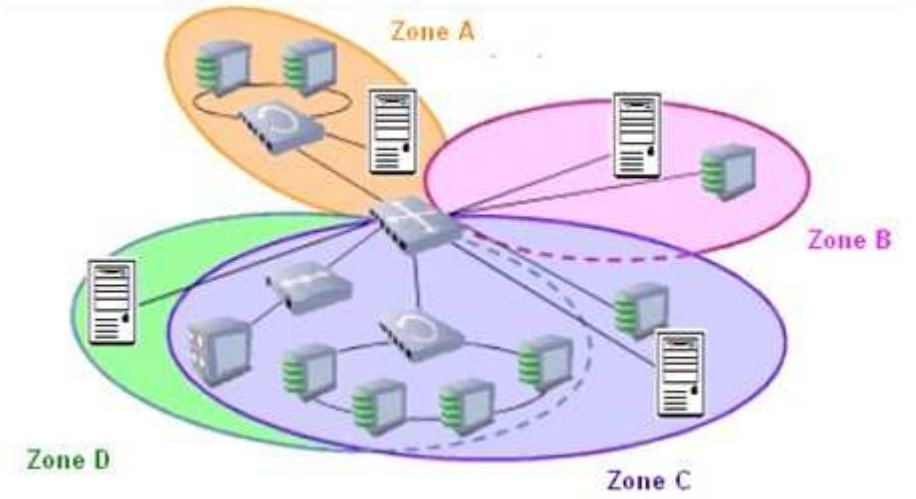

A zone that belongs to a single SAN can be grouped into a zone set that can be activated or deactivated as a single entity across all switches in the fabric.

A zone set includes one or more zones, and a zone can be a member of more than one zone set.

The ports and devices that are in a zone are called zone members. A zone includes one or more zone members.

A zone alias is a collection of zone members. A zone alias can be added to one or more zones.

Answer option A is incorrect. A domain id is a number that uniquely identifies a switch in a fabric and is the highest or most significant hierarchical level in the three-level address hierarchy.

In general, each switch is a single domain. The domain ID is an 8-bit identifier with a range of 0-255. Typically, zero (0) is reserved and one (1) is the default setting for new switches.

Answer option D is incorrect. N_Port ID Virtualization, or NPIV, is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port. It allows multiple Fibre Channel initiators to occupy a single physical port, easing hardware requirements in Storage Area Network design, especially where virtual SANs are called for.

Answer option C is incorrect. A SCSI ID is a unique identifier that is assigned to each SCSI device on the bus.

Reference: http://www.emcstorageinfo.com/2007/11/san-zoning-in-details.html

QUESTION 36
Which of the following enables the transmission of Fibre Channel information by tunneling data between storage area network facilities over IP networks?

A. FCP
B. SCSI
C. FCoE
D. FCIP

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Fibre Channel over IP (FCIP) is an IP-based storage networking technology, developed by the IETF and

defined in RFC 3821. FCIP is also known as Fibre Channel tunneling or storage tunneling. FCIP enables the transmission of Fibre Channel information by tunneling data between storage area network facilities over IP networks. This capacity facilitates data sharing over a geographically distributed enterprise. FCIP is a key technology that is expected to help bring about rapid development of the storage area network market by increasing the capabilities and performance of storage data transmission.

Answer option B is incorrect. SCSI stands for Small Computer System Interface. It is not an interface standard; however, it is a set of standards for physically connecting and transferring data between computers and peripheral devices. SCSI is an intelligent, peripheral, buffered, peer to peer interface. It hides the complexity of physical format. Every device attaches to the SCSI bus in a similar manner.

Up to 8 or 16 devices can be attached to a single bus. There can be any number of hosts and peripheral devices, but there should be at least one host. SCSI is most commonly used for hard disks and tape drives, but it can connect a wide range of other devices, including scanners and CD drives.

Answer option A is incorrect. Fibre Channel Protocol (FCP) is a transport protocol (similar to TCP used in IP networks), which predominantly transports SCSI commands over Fibre Channel networks.

Answer option C is incorrect. Fibre Channel over Ethernet (FCoE) is an encapsulation of Fibre Channel frames over Ethernet networks. It allows Fibre Channel to use 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol. FCoE maps Fibre Channel over selected full duplex IEEE 802.3 networks for providing I/O consolidation over Ethernet and reducing network complexity in the datacenter. The FCoE protocol specification replaces the FC0 and FC1 layers of the Fibre Channel stack with Ethernet.

QUESTION 37
Which of the following uniquely identifies a switch in a fabric and is the highest or most significant hierarchical level in the three-level address hierarchy?

A. NPIV
B. Zone alias
C. SCSI id
D. Domain id

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
A domain id is a number that uniquely identifies a switch in a fabric and is the highest or most significant hierarchical level in the three-level address hierarchy.

In general, each switch is a single domain. The domain ID is an 8-bit identifier with a range of 0-255. Typically, zero (0) is reserved and one (1) is the default setting for new switches.

Answer option C is incorrect. A SCSI ID is a unique identifier that is assigned to each SCSI device on the bus. It helps in determining the SCSI device's priority. There are 16 IDs, ranging from 0 to 15.

Answer option B is incorrect. A zone alias is a collection of zone members. A zone alias can be added to one or more zones.

Answer option A is incorrect. N_Port ID Virtualization, or NPIV, is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port.

QUESTION 38
Which of the following protocols is used to encapsulate Fibre Channel packets within an IP datagram permitting Fibre Channel SANs, Fibre Channel attached storage arrays, or Fibre Channel attached tape libraries to communicate across an IP network?

A. FCP

B. NFS
C. FCIP
D. SCSI

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Fibre Channel over IP (FCIP) is an IP-based storage networking technology, developed by the IETF and defined in RFC 3821. FCIP is also known as Fibre Channel tunneling or storage tunneling. FCIP enables the transmission of Fibre Channel information by tunneling data between storage area network facilities over IP networks. This capacity facilitates data sharing over a geographically distributed enterprise. FCIP is a key technology that is expected to help bring about rapid development of the storage area network market by increasing the capabilities and performance of storage data transmission. The FCIP protocol is used to encapsulate Fibre Channel packets within an IP datagram permitting Fibre Channel SANs, Fibre Channel attached storage arrays, or Fibre Channel attached tape libraries to communicate across an IP network. This protocol also helps in supporting remote mirroring of Fibre Channel switch attached storage arrays across geographically dispersed data centers.

Answer option D is incorrect. Small Computer System Interface (SCSI) is a set of standards for physically connecting and transferring data between computers and peripheral devices. The SCSI standards define commands, protocols, and electrical and optical interfaces. SCSI is most commonly used for hard disks and tape drives, but it can connect a wide range of other devices, including scanners and CD drives.

Answer option A is incorrect. Fibre Channel Protocol (FCP) is a transport protocol (similar to TCP used in IP networks), which predominantly transports SCSI commands over Fibre Channel networks.

Answer option B is incorrect. Network File System (NFS) is a network file system protocol that allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed. NFS, like many other protocols, builds on the Open Network Computing Remote Procedure Call (ONC RPC) system. The Network File System is an open standard defined in RFCs, allowing anyone to implement the protocol.

QUESTION 39
In _____ zoning, each device is assigned to a particular zone, and this assignment does not change, whereas in _____ zoning, device assignments can be changed by the network administrator to accommodate variations in the demands on different servers in the network.

Correct Answer: hard, soft
Section: (none)
Explanation

Explanation/Reference:
Zoning can be either hard or soft. In hard zoning, each device is assigned to a particular zone, and this assignment does not change. Hard zoning is zoning which is implemented in hardware and physically blocks access to a zone from any device outside of the zone.

In soft zoning, device assignments can be changed by the network administrator to accommodate variations in the demands on different servers in the network. Soft zoning is zoning which is implemented in software and uses filtering implemented in fibre channel switches to prevent ports from being seen from outside of their

assigned zones.

The use of zoning is said to minimize the risk of data corruption, help secure data against hackers, slow the spread of viruses and worms, and minimize the time necessary for servers to reboot.

Zoning best practices
Following are the zoning best practices:

- Zoning should always be implemented, even if a user uses LUN masking. Logical Unit Number Masking or LUN masking is an authorization process that makes a Logical Unit Number available to some hosts and unavailable to other hosts. Zoning is sometimes confused with LUN masking, because it serves the same goals. LUN masking, however, works on Fibre Channel level 4 (i.e., on SCSI level), while zoning works on level 2. This allows zoning to be implemented on switches, whereas LUN masking is performed on endpoint devices such as host adapters or disk array controllers
- Zoning should use pWWN identification for both security and operational consistency.
- Make zoning aliases and names only as long as they required to be, which enables maximum scaling.
- Frame-based hardware enforcement should be used by all zones; the best method to do this is to use pWWN identification exclusively for all zoning configurations.
- Need to implement default zone.

QUESTION 40
Which of the following is a protocol that permits programs and makes requests for files and services on remote computers on the Internet?

A. FCoE
B. CIFS
C. NFS
D. DCB

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Common Internet File System (CIFS) is a protocol that permits programs and makes requests for files and services on remote computers on the Internet. The client/server programming model is used by CIFS. A client program makes a request of a server program (usually in another computer) for accessing a file or passing a message to a program that runs in the server computer. The server takes the requested action and returns a response.

Answer option C is incorrect. Network File System (NFS) is a network file system protocol that allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed. NFS, like many other protocols, builds on the Open Network Computing Remote Procedure Call (ONC RPC) system. The Network File System is an open standard defined in RFCs, allowing anyone to implement the protocol.

Answer option A is incorrect. Fibre Channel over Ethernet (FCoE) is an encapsulation of Fibre Channel frames over Ethernet networks. It allows Fibre Channel to use 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol. FCoE maps Fibre Channel over selected full duplex IEEE 802.3 networks for providing I/O consolidation over Ethernet and reducing network complexity in the datacenter. The FCoE protocol specification replaces the FC0 and FC1 layers of the Fibre Channel stack with Ethernet.

Answer option D is incorrect. Data center bridging (DCB) refers to a set of enhancements to Ethernet local area networks for use in data center environments.

QUESTION 41
Which of the following is a protocol and mechanism for intelligent discovery of storage devices in an IP network?

A. NFS
B. iSCSI
C. FCoE
D. CIFS

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
iSCSI (Internet Small Computer System Interface) is an IP-based storage networking standard for linking data storage facilities. iSCSI easily transfers data over intranets and manages storage over long distances by using SCSI commands over IP networks. iSCSI can transmit data over local area networks, wide area networks, or the Internet and can enable location-independent data storage and retrieval. This protocol allows initiators to send SCSI commands to SCSI storage devices on remote servers. It is a popular storage area network protocol that allows organizations to consolidate storage into data center storage arrays while providing hosts with the illusion of locally-attached disks. Unlike traditional Fibre Channel, which requires special-purpose cabling, iSCSI can be run over long distances using existing network infrastructure. iSCSI is a protocol and mechanism for intelligent discovery of storage devices in an IP network.

Reference: http://www.snia.org/education/dictionary/i

QUESTION 42
Which of the following are network file system protocols designed for file sharing and are typically used in heterogeneous NAS environments?

Each correct answer represents a complete solution. Choose two.

A. NFS
B. FCIP
C. DAFS
D. CIFS

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
NFS and CIFS are network file system protocols designed for file sharing and are typically used in heterogeneous NAS environments.

What is Network File System?
Network File System (NFS) is a network file system protocol that allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed. NFS, like many other protocols, builds on the Open Network Computing Remote Procedure Call (ONC RPC) system. The Network File System is an open standard defined in RFCs, allowing anyone to implement the protocol.

Common Internet File System
Common Internet File System (CIFS) is a protocol that permits programs and makes requests for files and services on remote computers on the Internet. The client/server programming model is used by CIFS. A client program makes a request of a server program (usually in another computer) for accessing a file or passing a message to a program that runs in the server computer. The server takes the requested action and returns a response.

Answer option C is incorrect. DAFS is a new NAS/SAN file system designed for file sharing.

Answer option B is incorrect. Fibre Channel over IP (FCIP) is an IP-based storage networking technology, developed by the IETF and defined in RFC 3821. FCIP is also known as Fibre Channel tunneling or storage

tunneling.

QUESTION 43
Which of the following components is a configuration management element of an IP storage solution?

A. HBA
B. VSAN
C. ISL
D. VLAN

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Virtual local area network (VLAN) is a virtual subnet that is created by switches and supported routers that are VLAN enabled. VLAN is created by tagging the data frames that a switch receives from hosts. Each port on the switch behaves in the same way as an IP subnet and might require routing to communicate with hosts on other VLANs. VLANs can be used to isolate hosts and segments and to control broadcast traffic. VLAN is a configuration management element of an IP storage solution.

Answer option A is incorrect. HBA is an I/O adapter that connects a host computer bus to a Fibre Channel or SCSI medium.

Answer option B is incorrect. A virtual storage area network (VSAN) is a collection of ports from a set of connected Fibre Channel switches, which form a virtual fabric.

Answer option C is incorrect. Inter-Switch Link (ISL) is a Cisco proprietary protocol that maintains VLAN information as traffic flows between switches and routers.

QUESTION 44
Which of the following classes of service helps in assigning only a fraction of the total bandwidth and different QoS parameters for each connection?

A. Class 3
B. Class 4
C. Class 1
D. Class 2

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Class 4 service is connection-oriented providing a virtual circuit between a pair of nodes. It is also known as Fractional Bandwidth, as it helps in assigning only a fraction of the total bandwidth and different QoS parameters for each connection. Class 4 service provides a quality of service with guaranteed bandwidth.

Answer option C is incorrect. Class 1 service helps in establishing a dedicated connection through the fabric before transferring data, much like a virtual private network.

Answer option D is incorrect. Class 2 service does not require a dedicated channel. In Class 2 service, the receiving device sends an acknowledgment message in order to tell the sending device that the data has been received.

Answer option A is incorrect. Class 3 service also does not require a dedicated channel. It is a connection-less and unacknowledged service.

Reference: http://searchstorage.techtarget.com/tip/Service-classes-in-Fibre-Channel-SANs

QUESTION 45
Which of the following is a neighbor discovery protocol that helps in defining a method for network access devices using Ethernet connectivity for advertising information about them to peer devices on the same physical LAN and store information about the network?

A. LLDP
B. FCoE
C. iSCSI
D. NFS

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
The Link Layer Discovery Protocol (LLDP) is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on an IEEE 802 local area network. LLDP performs functions similar to several proprietary protocols, such as Cisco Discovery Protocol, Extreme Discovery Protocol, Nortel Discovery Protocol (also known as SONMP), and Microsoft's Link Layer Topology Discovery (LLTD). It is a neighbor discovery protocol that helps in defining a method for network access devices using Ethernet connectivity for advertising information about them to peer devices on the same physical LAN and store information about the network.

Answer option B is incorrect. Fibre Channel over Ethernet (FCoE) is an encapsulation of Fibre Channel frames over Ethernet networks. It allows Fibre Channel to use 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol. FCoE maps Fibre Channel over selected full duplex IEEE 802.3 networks for providing I/O consolidation over Ethernet and reducing network complexity in the datacenter. The FCoE protocol specification replaces the FC0 and FC1 layers of the Fibre Channel stack with Ethernet.

Answer option C is incorrect. iSCSI (Internet Small Computer System Interface) is an IP-based storage networking standard for linking data storage facilities. iSCSI easily transfers data over intranets and manages storage over long distances by using SCSI commands over IP networks. iSCSI can transmit data over local area networks, wide area networks, or the Internet and can enable location-independent data storage and retrieval. This protocol allows initiators to send SCSI commands to SCSI storage devices on remote servers. It is a popular storage area network protocol that allows organizations to consolidate storage into data center storage arrays while providing hosts with the illusion of locally-attached disks. Unlike traditional Fibre Channel, which requires special-purpose cabling, iSCSI can be run over long distances using existing network infrastructure.

Answer option D is incorrect. Network File System (NFS) is a network file system protocol that allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed. NFS, like many other protocols, builds on the Open Network Computing Remote Procedure Call (ONC RPC) system. The Network File System is an open standard defined in RFCs, allowing anyone to implement the protocol.

Reference: http://www.extremenetworks.com/libraries/products/LLDP_TB.pdf

QUESTION 46
Which of the following is a mechanism used for managing traffic in a network by specifying message or packet priority or delivery acknowledgement?

A. Multipathing
B. Class of service
C. Data center bridging
D. Buffer-to-buffer

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:
A class of service is a mechanism used for managing traffic in a network by specifying message or packet priority or delivery acknowledgement.

Fibre Channel provides several connection strategies, known as classes of service, for connecting devices through a Fibre Channel fabric. ("Fabric" is Fibre Channel's term used for the cross-point switch at the heart of a fabric topology Fibre Channel.

Fibre Channel specifies 5 classes of service that address different delivery requirements:
- Class 1 service helps in establishing a dedicated connection through the fabric before transferring data, much like a virtual private network.
- Class 2 service does not require a dedicated channel. In Class 2 service, the receiving device sends an acknowledgment message in order to tell the sending device the data has been received.
- Class 3 service also does not require a dedicated channel. It is a connection-less and unacknowledged service.
- Class 4 service is connection-oriented providing a virtual circuit between a pair of nodes. It is also known as Fractional Bandwidth, as it helps in assigning only a fraction of the total bandwidth and different QoS parameters for each connection.
- Class 6 service is multi-cast (one-to-many) transmission through the fabric.

Answer option C is incorrect. Data center bridging (DCB) refers to a set of enhancements to Ethernet local area networks for use in data center environments.

Answer option D is incorrect. Buffer credit, also known as buffer-to-buffer credit (BBC), is used as a flow control method by Fibre Channel technology and represents the number of frames a port can store.

Answer option A is incorrect. Multipathing is a software technology employed at the operating system device driver level. Multipathing helps in creating a pseudo device in order to facilitate the sharing and balancing of I/O operations across all of the available I/O paths.

Reference: http://www.snia.org/education/dictionary/c

QUESTION 47
Which of the following tasks are performed by the FC traceroute tool?

Each correct answer represents a complete solution. Choose two.

A. Rescanning network devices
B. Checking end-to-end connectivity
C. Tracing the route followed by data traffic
D. Computing inter-switch (hop-to-hop) latency

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:
The FC traceroute tool is used to perform the following tasks:
- Tracing the route followed by data traffic
- Computing inter-switch (hop-to-hop) latency

FC traceroute helps in identifying the path taken on a hop-by-hop basis and includes a timestamp at each hop in both directions. FC traceroute is used to test the connectivity of TE ports along the path between the generating switch and the switch closest to the destination.

Answer option B is incorrect. The fcping feature checks the reachability of a node by checking and comparing

its end-to-end connectivity. One can invoke the fcping feature by providing the FC ID option, the destination port WWN, or the device alias information of the system. It is also known as Fibre Channel Ping Utility. The fcping utility can be used as a troubleshooting tool on Fibre Channel Storage Area Networks (SANs).

Answer option A is incorrect. The Rescan tool is used for rescanning network devices.

QUESTION 48
What are the benefits of using storage area network (SAN)?

Each correct answer represents a complete solution. Choose four.

A. Enables more effective disaster recovery processes.
B. SAN is the least expensive technology.
C. SAN architecture facilitates scalability.
D. Sharing storage usually simplifies storage administration and adds flexibility.
E. Ability to allow servers to boot from the SAN itself.

Correct Answer: ACDE
Section: (none)
Explanation

Explanation/Reference:
Following are the benefits of using storage area network (SAN):
- Sharing storage usually simplifies storage administration and adds flexibility.
- Ability to allow servers to boot from the SAN itself.
- Enables more effective disaster recovery processes.
- Enhances storage resource utilization
- Helps in decreasing network traffic due to backups.

Answer option B is incorrect. SAN technology is very expensive, as Fibre channel technology tends to be pricier and maintenance needs a higher degree of skill.

Reference: http://www.peterindia.net/SANOverview.html

QUESTION 49
Which of the following statements are true about Direct-attached storage (DAS)?

Each correct answer represents a complete solution. Choose three.

A. This storage system does not share data or unused resources with other servers.
B. The main protocols used for DAS connections are ATA, SATA, eSATA, SCSI, SAS, and Fibre Channel.
C. It is built on top of SAN or NAS technology.
D. It is a storage topology, which refers to a digital storage system directly attached to a server or workstation, without a storage network in between.

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
Direct-attached storage (DAS) is a storage topology which refers to a digital storage system directly attached to a server or workstation, without a storage network in between. A typical DAS system is made of a data storage device (for example enclosures holding a number of hard disk drives) connected directly to a computer through a host bus adapter (HBA). There is no networking device (like hub, switch, or router) working between the workstation and data storage device. The main protocols used for DAS connections are ATA, SATA, eSATA, SCSI, SAS, and Fibre Channel. The main disadvantage of DAS include that it does not share data or unused resources with other servers.

Answer option C is incorrect. NAS is built on top of SAN or DAS technology.

Reference: http://matteocappelli.files.wordpress.com/2011/02/storage_basics.pdf

QUESTION 50
What are the benefits of using a storage area network?

Each correct answer represents a complete solution. Choose three.

A. SAN technology is very inexpensive.
B. Enables more effective disaster recovery processes.
C. Increases the storage space usability.
D. Ability to permit servers to boot from the SAN itself.

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:
Following are the benefits of using storage area network:
- Ability to permit servers to boot from the SAN itself. This helps in quick and easy replacement of faulty servers.
- SAN increases the storage space usability as different servers can access a shared storage array. Also, clustered servers that are placed in different buildings can access the same storage.
- Enables more effective disaster recovery processes.

Answer option A is incorrect. SAN is a very expensive technology, as Fibre channel technology is likely to be pricier, and maintenance needs a higher degree of skill.

Reference: http://matteocappelli.files.wordpress.com/2011/02/storage_basics.pdf

QUESTION 51
Which of the following accesses the file system through a local area network?

A. CAD
B. NAS
C. SAN
D. DAS

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
NAS accesses the file system through a local area network. Network-attached storage (NAS) is file-level computer data storage connected to a computer network providing data access to heterogeneous network clients. NAS systems contain one or more hard disks, often arranged into logical, redundant storage containers or RAID arrays. It removes the responsibility of file serving from other servers on the network. NAS uses file-based protocols, such as NFS, SMB/CIFS, or AFP. NAS units rarely limit clients to a single protocol.

Answer options C, A, and D are incorrect. These technologies do not access the file system through a local area network.

Reference: http://searchstorage.techtarget.com/definition/network-attached-storage

QUESTION 52

What of the following Identifies are used by a SCSI address?

Each correct answer represents a part of the solution. Choose three.

A. Target ID
B. Adaptor number
C. SCSI bus
D. Logical unit number (LUN)

Correct Answer: ACD
Section: (none)
Explanation

Explanation/Reference:
In order for backup devices to be available, the system must know which physical device is mapped to a given SCSI address. When a SCSI bus is directly attached to a storage device, the SCSI address is determined by the physical SCSI connection. A SCSI address includes three identifiers that are as follows:
1. Bus
2. Target
3. Logical Unit Number

Reference: http://documentation.commvault.com/dell/release_7_0_0/books_online_1/english_us/features/san_lib/pre_install.htm#SAN Addressing Overview

QUESTION 53
Which of the following are the three most common advantages of implementing iSNS?

Each correct answer represents a complete solution. Choose three.

A. Discovery Domains and Login Control
B. Transport Layer Security
C. Storage Resource Discovery
D. State Change Notification

Correct Answer: ACD
Section: (none)
Explanation

Explanation/Reference:
An iSNS implementation provides four primary services:
- **Name Registration and Storage Resource Discovery**: iSNS implementations allow all entities in a storage network to register and query an iSNS database.
- **Discovery Domains and Login Control**: Administrators can use the Discovery Domains to divide storage nodes into manageable, non-exclusive groups. By grouping storage nodes, administrators are able to limit the login process of each host to the most appropriate subset of targets registered with the iSNS, which allows the storage network to scale by reducing the number of unnecessary logins and by limiting the amount of time each host spends establishing login relationships. Each target is able to use Login Control to delegate their access control and authorization policies to an iSNS server. Such delegation is intended to promote centralized management.
- **State Change Notification**: The State Change Notification (SCN) service allows an iSNS Server to issue notifications about each event that affects storage nodes on the managed network. Each iSNS client may register for notifications on behalf of its storage nodes, and each client is expected to respond according to its own requirements and implementation.
- **Bidirectional mappings between fibre channel and iSCSI devices**: Because the iSNS database stores naming and discovery information about both Fibre Channel and iSCSI devices, iSNS servers are able to store mappings of Fibre Channel devices to proxy iSCSI device images on the IP network. These mappings may also be made in the opposite direction, allowing iSNS servers to store mappings from iSCSI devices to

proxy WWNs.

Answer option B is incorrect. iSNS does not provide Transport Layer Security.

QUESTION 54
Which of the following correctly represents the strengths of Direct-Attached Storage?

A. RAID protection
B. High scalability
C. Inexpensive hardware
D. Unlimited distance for SCSI

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
DAS hardware includes relatively inexpensive commodity products. Also, it uses SCSI technology usually that is a very familiar technology.

Answer option B is incorrect. DAS provides low scalability.

Answer option A is incorrect. RAID does not protect it from server hardware or software failure.

Answer option D is incorrect. DAS can use limited distance for SCSI.

QUESTION 55
Which of the following statements correctly represent the limitations of NAS?

Each correct answer represents a complete solution. Choose two.

A. NAS is less scalable and reliable than DAS.
B. NAS does not allow access from multiple operating systems.
C. NAS performance is often less robust for applications that use block-level storage.
D. Storage traffic must compete with other network traffic.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:
In NAS, storage traffic must compete with other network traffic because host access NAS storage resources over LAN. For applications that use block-level storage, NAS does not perform robustly because TCP/IP is not optimized for block-level storage.

Answer option A is incorrect. NAS is more scalable and reliable than DAS.

Answer option B is incorrect. NAS allows access from multiple operating systems.

QUESTION 56
In which of the following RAID levels if one disk fails, all of the data on the array will be lost, as there is neither parity nor mirroring?

A. RAID 1
B. RAID 5
C. RAID 3

D. RAID 0

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
RAID 0 (striped disks) distributes data across multiple disks in a way that gives improved speed at any given instant. If one disk fails, however, all of the data on the array will be lost, as there is neither parity nor mirroring. In this regard, RAID 0 is non-redundant, so it should not be used for a critical system. A RAID 0 array requires a minimum of two drives.

Answer options A, B, and C are incorrect. These RAID levels have parity or mirroring.

Reference: http://www.ecs.umass.edu/ece/koren/architecture/Raid/basicRAID.html

QUESTION 57
You have attached ten hard disk drives in a RAID 5 array. What will be the total storage capacity of this array for storing files if each hard disk can store data up to 40GB?

A. 320GB
B. 400GB
C. 40GB
D. 360GB

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
In a RAID 5 array, one hard disk is dedicatedly used for keeping the parity information. In the above question, you will be able to use only nine hard disks for storage purpose as one of the hard disks will be used for parity. Each hard disk can store data up to 40GB so there will be only 360GB space left for storing files.

QUESTION 58
Which of the following RAID levels is also termed as "Bit striping with ECC"?

A. RAID 0
B. RAID 1
C. RAID 2
D. RAID 4

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
RAID 2 is also termed as "Bit Striping with ECC", as the disks are striped at the byte level. RAID is not much commercially viable as it uses non-standard disk drives.

Answer option A is incorrect. RAID 0 is known as "Block striping without parity".

Answer option B is incorrect. RAID 1 is called "Mirroring of disks".

Answer option D is incorrect. RAID 4 is termed as "Block striping with dedicated parity disk".

QUESTION 59

Which of following RAID levels is also known as "Block striping with two sets of distributed parity"?

A. RAID 6
B. RAID 5
C. RAID 0
D. RAID 1

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
RAID 6 is also known as "Block striping with two sets of distributed parity", as the data is striped at the block level and the parity information is distributed across all disks. RAID 6 provides protection against a double disk failure as it uses two parity blocks instead of one.

Answer option B is incorrect. RAID 5 uses only one block of parity.

Answer option C is incorrect. RAID 0 is known as "Block striping without parity".

Answer option D is incorrect. RAID 1 is called "Mirroring of disks".

QUESTION 60
What is the reason why RAID 2 is not commercially viable?

A. It does not provide protection to data.
B. It uses non-standard disk drives.
C. It does not provide any fault tolerance.
D. It provides only byte level striping.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
RAID 2 is called "Bit Striping with ECC", as the data is striped at the byte level. It uses ECC for providing protection to data. RAID 2 is not commercially viable as it uses non-standard disk drives.

QUESTION 61
Which of the following RAID levels use block level striping?

Each correct answer represents a complete solution. Choose four.

A. RAID 4
B. RAID 2
C. RAID 6
D. RAID 5
E. RAID 0

Correct Answer: ACDE
Section: (none)
Explanation

Explanation/Reference:
RAID 0, 4, 5, and 6 all use block level striping. In these types of RAID, a disk contains usually 512 bytes of data associated with a block command stored on one sector.

Answer option B is incorrect. RAID 2 and 3 use byte level striping.

QUESTION 62
A ___ is a logical reference of a storage subsystem that represents all or a part of a disk, or a section of a disk array in the storage subsystem.

Correct Answer: LUN
Section: (none)
Explanation

Explanation/Reference:
Logical Unit Number (LUN) refers to a single unit of storage. It can also be referred to as a volume or a logical drive depending on the host system environment. A RAID group is equivalent to a single LUN in simple systems that deliver RAID capability. A host server considers this LUN as a single simple storage unit present for access by the server. It is possible that RAID groups have one or more LUNs created for access by one or more host servers in advanced storage arrays.

The LUN number is a field within a SCSI command. A single bare disk drive does not have capability to separate the disk into Logical Units, so it will present its storage capacity as a single contiguous set of blocks. A 10 GB disk will present approximately 20,000,000 x 512 Bytes sectors to the host.

Reference: http://en.wikipedia.org/wiki/Logical_Unit_Number

QUESTION 63
Which of the following LUN types provides fault tolerance and improved I/O performance in a Windows Server environment?

A. Mirroring
B. Stripping
C. Spanning
D. Striped with parity

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Striped with Parity is a type of Logical Unit Number (LUN) that provides fault tolerance and improved I/O performance in a Windows Server environment. Striped with Parity LUN type is also known as RAID 5. It requires minimum of three disks. In case a disk fails, data of that disk is restored from the parity information on other disks. In case of failure of two (or more) disks, all data will be lost.

QUESTION 64
Which of the following processes makes a Logical Unit Number available to some hosts and unavailable to other hosts?

A. Multiplexing
B. Shoe-shining
C. Disk Mirroring
D. LUN Masking

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Logical Unit Number Masking or LUN Masking is an authorization process that makes a Logical Unit Number available to some hosts and unavailable to other hosts. LUN Masking is implemented primarily at the HBA (Host Bus Adapter) level. Some storage controllers also support LUN Masking. LUN Masking is important because Windows-based servers attempt to write volume labels to all available Logical Unit Numbers.

The mask is created by assigning the HBAs unique worldwide name, typically the World Wide Port Name, to each Logical Unit.

Answer option C is incorrect. Disk mirroring is the replication of logical disk volumes onto separate physical hard disks in real time to ensure continuous availability.

Answer option B is incorrect. Shoe-shining is a disadvantage with read/write tape drives. During shoe-shining, the data transfer rate falls below the minimum threshold at which the tape drive heads were designed to transfer data to or from a continuously running tape. In this situation, the modern fast-running tape drive is unable to stop the tape instantly. Instead, the drive must decelerate and stop the tape, rewind it a short distance, restart it, position back to the point at which streaming stopped, and then resume the operation. If the condition repeats, the resulting back-and-forth tape motion resembles that of shining shoes with a cloth. Shoe-shining decreases the attainable data transfer rate, and drive and tape life.

In early tape drives, non-continuous data transfer was normal and unavoidable - computer processing power and memory available were usually insufficient to provide a constant stream. So tape drives were typically designed for so called start-stop operation. Early drives used very large spools, which necessarily had high inertia and did not start and stop moving easily. To provide high start, stop, and seeking performance, several feet of loose tape was played out and pulled by a suction fan down into two deep open channels on either side of the tape head and capstans. The long thin loops of tape hanging in these vacuum columns had far less inertia than the two reels and could be rapidly started, stopped and repositioned.

Later, most tape drives of the 1980s introduced the use of an internal data buffer to somewhat reduce start-stop situations. These drives are often referred to as tape streamers. The tape was stopped only when the buffer contained no data to be written, or when it was full of data during reading.

Most recently, drives no longer operate at a single fixed linear speed, but have several speeds. Internally, they implement algorithms that dynamically match the tape speed level to the computer's data rate. Example speed levels could be 50 percent, 75 percent and 100 percent of full speed. A computer that streams data slower than the lowest speed level (e.g. at 49 percent) will still cause shoe-shining.

Answer option A is incorrect. Multiplexing is the process of alternately writing or interleaving multiple clients-data to a single tape. Multiplexing sends data from multiple clients to a single tape drive. This is useful for low end clients with slow throughput. It allow backups of multiple clients to send data to a single tape drive simultaneously.

QUESTION 65
Which of the following virtualization techniques is known as asymmetric virtualization device?

A. Host-based LDM
B. Host-based LVM
C. In-band appliance
D. Out-of-band appliance

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Answer options B and A are incorrect. Host-based virtualization requires additional software running on the host, as a privileged task or process. Most modern operating systems have some form of logical volume manager built-in, such as, LVM (Logical Volume Manger) in UNIX/Linux and Logical Disk Manager (LDM) in Windows.

Answer option C is incorrect. In-band appliances (symmetric virtualization devices) actually sit in the data path between the host and storage. All I/O requests and their data pass through the device. Hosts perform I/O to the virtualization device and never interact with the actual storage device. The virtualization device in turn performs I/O to the storage device. Caching of data, statistics about data usage, replications services, data migration and thin provisioning are all easily implemented in an in-band device.

QUESTION 66
Which of the following are host-based virtualization techniques?

Each correct answer represents a complete solution. Choose two.

A. Out-of-band appliance
B. LDM
C. LVM
D. In-band appliance

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
Host-based virtualization technique requires additional software running on the host, as a privileged task or process. Most modern operating systems have some form of logical volume manager built-in, such as, LVM (Logical Volume Manger) in UNIX/Linux and Logical Disk Manager (LDM) in Windows.

Answer options D and A are incorrect. These two are network based storage virtualization techniques.

QUESTION 67
Which of the following sentences describes the advantages of virtualization?

Each correct answer represents a complete solution. Choose two.

A. It allows file sharing between dissimilar applications.
B. It allows an operator to manage more number of devices.
C. It improves data integrity.
D. It provides higher storage asset utilization.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Storage virtualization is a storage technique used in a storage area network (SAN) for the management of storage devices. The various ways used to configure storage virtualization are as follows:
- Host-based
- Storage device-based
- Network-based

Storage virtualization helps to achieve location independence by abstracting the physical location of the data. The physical storage resources are aggregated into storage pools, from which the logical storage is created. It presents to the user a logical space for data storage and transparently handles the process of mapping it to the actual physical location. This is implemented in modern disk arrays, using vendor proprietary solutions. Storage virtualization provides higher storage asset utilization and an operator to manage more number of devices.

QUESTION 68
Which of the following helps in avoiding any compromise of critical data within the SAN?

A. Overlay of VSAN
B. IPSec
C. CHAP
D. SLAP

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Overlay of VSAN helps in avoiding any compromise of critical data within the SAN. Zoning allows a user to overlay a security map stating which devices (namely hosts) can see which targets are reducing the risk of data loss.

Answer option B is incorrect. Internet Protocol Security (IPsec) is a method of securing data. It secures traffic by using encryption and digital signing. It enhances the security of data as if an IPsec packet is captured. Its contents cannot be read. IPsec also provides sender verification that ensures the certainty of the datagram's origin to the receiver.

The IPsec protocol has three major components:
1. Authentication Header (AH)
2. Encapsulating Security Payload (ESP)
3. Internet Key Exchange (IKE)

Answer option C is incorrect. Challenge Handshake Authentication Protocol (CHAP) is an authentication protocol that uses a secure form of encrypted authentication. Using CHAP, network dial-up connections are able to securely connect to almost all PPP servers.

Answer option D is incorrect. SLAP is an authentication technique for Fiber Channel switches, which make use of digital certificates to authenticate switch ports and switches. SLAP has been designed to prevent the unauthorized addition of switches into a Fiber Channel network.

QUESTION 69
The command-line interface (CLI) is used when the router is accessed by the user from the console or via a remote network connection. Which of the following are the CLI interfaces?

Each correct answer represents a complete solution. Choose three.

A. SSH
B. Telnet
C. RDP
D. Console

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
The following are the different types of CLI interfaces:
- **Telnet**: It is the least secure CLI method. It is used when there is a requirement of management access and the connecting client does not use SSH client software. Interfaces that are allocated to the trust zone are telnet enabled by default.
- **SSH**: It is the most common network-based CLI method because this method has encrypted and secure management traffic.
- **Console**: It is the most secure CLI method. In this method, a direct physical access to the device is required by the means of a console cable.

Answer option C is incorrect. RDP provides access to the user interface of a remote device.

QUESTION 70
Which of the following is the default port for the SNMP protocol?

A. 22
B. 21
C. 80
D. 161

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Simple Network Management Protocol (SNMP) is a part of the TCP/IP protocol suite, which allows users to manage the network. SNMP is used to keep track of what is being used on the network and how the object is behaving. Port 161 is the default port for the SNMP protocol.

Answer options C, A, and B are incorrect. Ports 80, 22, and 21 are the default ports for HTTP, SFTP, and FTP protocols, respectively.

QUESTION 71
Which of the following automates the immigration of data from one storage class to another based on how old the data is and how frequently users access it?

A. ILM
B. HSM
C. CAS
D. SRM

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
HSM tools are used to automate the immigration of data from one storage class to another based on how old the data is and how frequently users access it.

Answer option C is incorrect. Content Address Storage (CAS) is a system that store data in such a way that once the data is stored, it cannot be modified or changed. It is not a backup solution, but it is used to provide a non-changing copy of data.

Answer option A is incorrect. Information Life Management (ILM) refers to the creation and management of a storage infrastructure. It includes processes and technologies required to manage data for its lifetime. ILM determines the priority and value of the data, including considerations such as data placement, deletion, repurposing, and regulatory archiving.

A successful information lifecycle management strategy must be business-centric by tying closely with key processes, applications, and initiatives of the business. It should be centrally managed, providing an integrated view into all information assets of the business, both structured and unstructured.

Answer option D is incorrect. SRM applications help a SAN administrator track and control use of SAN resources.

QUESTION 72
_____ _____ helps in protection against cache data loss that occurs when a storage processor fails.

Correct Answer: Cache mirroring
Section: (none)
Explanation

Explanation/Reference:
Cache mirroring helps in protecting against cache data loss that occurs when a storage processor fails. As a prerequisite to cache mirroring, the storage system should have redundant storage processors with failover capability. If one controller fails, the other controller is assumed to have the ownership of the failed processor's storage and continues where it left off.

A complex operation, cache mirroring is used to transfer large amounts of rapidly changing data between controllers. Many storage systems are implemented such that the cache data path between controllers is the same as the data path (for example, SCSI bus or Fibre Channel loop) to the disk drives.

QUESTION 73
Which of the following closely resembles snapshot and requires the same storage capacity as the primary volume?

A. Clone
B. Asynchronous replication
C. Mirroring
D. Synchronous replication

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Clones closely resemble snapshots, as they are bit-identical copies of full volumes, and thus require the same storage capacity as the primary volume.

Snapshot:
Snapshot is a virtual representation of a volume at a given point and is ideal for short-term backup and recovery purposes. Snapshot uses less space since only modified data is copied from the primary.

Snapshot is a point-in-time copy of a defined collection of data. It is space-efficient, any-point-in-time copies of changed bits or volumes, achieving time, cost and capacity efficient. Snapshot can be of files, LUNs, file systems, or any other type of container supported by the system it can be read-only or read-write copies of files, LUNs, file systems or any other type of container.

Reference: http://www.dellstorage.com/data-protection/data-protection-and-recovery/snapshots-clones-mirrors-replicas.aspx

QUESTION 74
Which of the following is a technique in which data should be committed to stable storage at both the primary site and the secondary site before the write is acknowledged to the host?

A. Clone
B. Snapshot
C. Asynchronous replication
D. Synchronous replication

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Synchronous replication is a technique in which data should be committed to stable storage at both the primary site and the secondary site before the write is acknowledged to the host.

In synchronous replication, data is copied over a storage area network (SAN), local area network (LAN), or wide area network (WAN) in order to have multiple up-to-date copies of the data.

Synchronous replication task is to write data to the primary and secondary sites at the same time such that the data remains current between sites. Synchronous replication is pricier than other forms of replication. This technique is often used for disaster recovery purposes and is preferred for applications with low recovery time objectives (RTOs) that cannot tolerate data loss.

Answer option C is incorrect. Asynchronous replication is a technique in which data should be committed to storage at only the primary site and not the secondary site before the write is acknowledged to the host.

Answer option B is incorrect. Snapshot is a virtual representation of a volume at a given point and is ideal for short-term backup and recovery purposes. Snapshot uses less space since only modified data is copied from the primary.

Snapshot is a point-in-time copy of a defined collection of data. It is space-efficient, any-point-in-time copies of changed bits or volumes, achieving time, cost and capacity efficient. Snapshot can be of files, LUNs, file systems, or any other type of container supported by the system it can be read-only or read-write copies of files, LUNs, file systems or any other type of container.

Answer option A is incorrect. Clones are closely similar to snapshots, as they are bit-identical copies of full volumes, and thus need the same storage capacity as the primary volume.

Reference: http://www.snia.org/education/dictionary/s#synchronous_replication

QUESTION 75
Snapshot is a point-in-time copy of a defined collection of data. What are the advantages of using snapshots?

Each correct answer represents a complete solution. Choose three.

A. It provides minimal use of SAN bandwidth.
B. It provides minimal impact on application performance.
C. Snapshot data set is 100% up to date.
D. It is a point-in-time copy that protects against corrupted data.

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
Following are the advantages of snapshots:
- It provides minimal impact on application performance.
- It provides minimal use of SAN bandwidth.
- It is a point-in-time copy that protects against corrupted data.

Answer option C is incorrect. Snapshot data set is never 100% up to date.

Snapshot:
Snapshot is a virtual representation of a volume at a given point and is ideal for short-term backup and recovery purposes. Snapshot uses less space since only modified data is copied from the primary.

Snapshot is a point-in-time copy of a defined collection of data. It is space-efficient, any-point-in-time copies of changed bits or volumes, achieving time, cost and capacity efficient. Snapshot can be of files, LUNs, file systems, or any other type of container supported by the system it can be read-only or read-write copies of files, LUNs, file systems or any other type of container.

QUESTION 76
Which of the following statements about incremental backups are true?

Each correct answer represents a complete solution. Choose two.

A. It backs up the entire database, including the transaction log.
B. It backs up only the files changed since the most recent backup and clears the archive bit.
C. It is the slowest method for taking a data backup.
D. It is the fastest method of backing up data.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
An incremental backup is the fastest method of backing up data. It backs up only the files changed since the most recent backup and clears the archive bit. In an incremental backup, data restoration is slower than the other backup methods. Restoring data from an incremental backup requires the last full backup and all subsequent incremental backups. Incremental backups must be restored in the same order as they were created.

QUESTION 77
Which of the following is defined as the maximum acceptable time period prior to a failure or disaster during which changes to data may be lost as a result of recovery?

A. RTO
B. MPIO
C. LVM
D. RPO

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Recovery Point Objective (RPO) is defined as the maximum acceptable time period prior to a failure or disaster during which changes to data may be lost as a result of recovery.

Data modified prior to the failure or disaster by at least this time period is preserved by recovery. Zero is a valid value and is correspondent to a "zero data loss" requirement.

Organizations determine RPO according to "acceptable loss" in a disaster situation. If the RPO of a company is 5 hours, it means that more than 5 hours of production data cannot be lost. In other words, it can be said that the data must be restored within 5 hours of the disaster.

Answer option A is incorrect. Recovery Time Objective (RTO) is defined as the maximum acceptable time period needed to bring one or more applications and associated data back from an outage to a correct operational state.

Answer option B is incorrect. In computer storage, multipath I/O is a fault-tolerance and performance enhancement technique whereby there is more than one physical path between the CPU in a computer system and its mass storage devices through the buses, controllers, switches, and bridge devices connecting them.

Answer option C is incorrect. Logical volume management is the method used for providing a higher level-view of the disk storage on a computer than the traditional view of disks and partitions. Due to its flexibility, it is used by the system administrators for allocating storage to applications and users.

Although storage volumes need some upgrading of file system tools, they can be resized and moved around at the user's will under the control of the logical volume manager. The logical volume manager also facilitates in the management of storage volumes in user-defined groups. This helps the system administrators to deal with sensibly named volume groups rather than physical disk names.

The following are the features of the LVM (Logical Volume Manager):
- It resizes volume groups online by absorbing new physical volumes (PV) or by ejecting existing ones.
- It resizes logical volumes (LV) online by concatenating extents onto them or by truncating extents from them.
- It creates read-only snapshots of logical volumes (LVM1).
- It creates read-write snapshots of logical volumes (LVM2).
- It stripes whole or parts of logical volumes across multiple PVs, in a fashion similar to RAID 0.
- It mirrors whole or parts of logical volumes, in a fashion similar to RAID 1.
- It moves online logical volumes between PVs.
- It splits or merges volume groups in situ (as long as no logical volumes span the split). This can be useful when migrating whole logical volumes to or from offline storage.

Reference: http://www.snia.org/education/dictionary/r

QUESTION 78
Which of the following are necessary and measurable criteria for evaluating the right backup and disaster recovery solutions?

Each correct answer represents a complete solution. Choose two.

A. RTO
B. HSM
C. NPIV
D. RPO

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Following are necessary and measurable criteria for evaluating the right backup and disaster recovery solutions:
- Recovery Point Objective (RPO): Recovery Point Objective (RPO) is defined as the maximum acceptable time period prior to a failure or disaster during which changes to data may be lost as a result of recovery.Data modified prior to the failure or disaster by at least this time period is preserved by recovery. Zero is a valid value and is correspondent to a "zero data loss" requirement. Organizations determine RPO according to "acceptable loss" in a disaster situation. If the RPO of a company is 5 hours, it means that more than 5 hours of production data cannot be lost. In other words, it can be said that the data must be restored within 5 hours of the disaster.
- Recovery Time Objective (RTO): The Recovery Time Objective is the duration of time and a service level within which a business process must be restored after a disaster in order to avoid unacceptable consequences associated with a break in business continuity. It includes the time for trying to fix the problem without a recovery, the recovery itself, tests, and the communication to the users. Decision time for users' representative is not included.The business continuity timeline usually runs parallel with an incident management timeline and may start at the same or different points.

Answer options B and C are incorrect. These are not necessary and measurable criteria for evaluating the right backup and disaster recovery solutions.

Reference: http://www.snia.org/education/dictionary/r

QUESTION 79
Using which of the following can data be sent electronically via a remote backup service?

A. VTL
B. Remote vaulting
C. E-vaulting
D. Vaulting

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Data can also be sent electronically via a remote backup service, which is known as electronic vaulting or e-vaulting. Electronic vaulting helps in reducing the risk of exposure of business critical data by getting the data off-site, off-line, and out-of-reach. This helps in ensuring that most current data is available for rapid recovery whenever and wherever the user requires it. It ensures automatic backup which are performed consistently, and securely.

Electronic vaulting helps in reducing the risk of exposure of business critical data by getting the data off-site, off-line, and out-of-reach. This helps in ensuring that most current data is available for rapid recovery whenever and wherever the user requires it. It ensures automatic backup which are performed consistently, and securely.

Electronic vaulting is the best way of storing and accessing key information from back up data, off line, off- site and out of reach. Using electronic vaulting, safe, efficient storage and retrieval of data in complex, distributed, multi-vendor environments are possible.

Sending backups off-site ensures systems and servers can be reloaded with the latest data in the event of a natural disaster, accidental error, or system crash. Sending backups off-site also ensures that there is a copy of pertinent data that is not stored on-site. Off-site backup services are convenient for companies that backup pertinent data on a daily basis (classified and unclassified).

Answer option D is incorrect. Vaulting is the strategy of sending critical data out of the main location (off the main site) as part of a disaster recovery plan. The storage of off-site data is also known as vaulting, as backups are stored in purpose built vaults. Using removable storage media such as magnetic tape or optical storage, data is usually transported off-site.

Answer option B is incorrect. This is an invalid option.

Answer option A is incorrect. A virtual tape library (VTL) is a data storage virtualization technology used typically for backup and recovery purposes.

QUESTION 80
In which of the following backup targets is data initially copied to backup storage on a disk storage system and then periodically copied again to a tape storage system?

A. Disk-to-disk-to-tape
B. Virtual tape library
C. Disk-to-disk
D. Disk-to-tape

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Disk-to-disk-to-tape (D2D2T) is an approach to computer storage backup and archiving. In D2D2T scheme, data is initially copied to backup storage on disk storage system and then periodically copied again to a tape storage system.

A D2D2T scheme permits the administrator to automate daily backups on disk for implementing fast restores

and then move data to tape when an administrator has time. The tape is also used to transfer more mature data offsite for disaster recovery protection and to comply with regulatory policies for long-term data retention at a relatively inexpensive cost.

Answer option B is incorrect. A virtual tape library (VTL) is a data storage virtualization technology used typically for backup and recovery purposes. A VTL is used to present a storage component (usually hard disk storage) as tape libraries or tape drives for use with existing backup software.

Virtualizing the disk storage as tape permits integration of VTLs with existing backup software and existing backup and recovery processes and policies. The benefits of such virtualization include storage consolidation and faster data restore processes.

By backing up data to disks instead of tapes, VTL often enhances performance of both backup and recovery operations. Restore processes are found to be faster than backup regardless of implementations. In some cases, the data stored on the VTL's disk array is exported to other media, such as physical tapes, for disaster recovery purposes (scheme called disk-to-disk-to-tape, or D2D2T).

Answer option C is incorrect. The term "disk-to-disk", or "D2D", generally refers to disk-to-disk backup. With D2D, a computer hard disk is backed up to another hard disk rather than to a tape or floppy disk. D2D is often confused with virtual tape, but it differs in that it enables multiple backup and recovery operations to simultaneously access the disk directly by using a true file system.

Answer option D is incorrect. The term "disk-to-tape" refers to disk-to-tape backup. In this process, a computer hard disk is backed up to a tape.

In addition, tape is quite simply a pain to work with, mainly if a cartridge must be retrieved, loaded, and scanned in its entirety for recovering one file and tapes can be lost or stolen, too.

When there is necessity of restoring data, disk's big benefits over tape is that it is random-access rather than sequential access. That means that if a user only needs one file or a few files back, it will be faster and easier to search and recover from disk.

One of the first backup applications for disk was to emulate a tape drive and using a VTL, there is no need to change your software or processes - they just run a lot faster.

Reference: http://searchstorage.techtarget.com/definition/disk-to-disk-to-tape

QUESTION 81
How many sets of backups are defined in Grandfather-father-son backup rotation scheme?

A. 4
B. 3
C. 2
D. 5

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Grandfather-father-son backup refers to the most common rotation scheme for rotating backup media. Originally designed for tape backup, it works well for any hierarchical backup strategy. The basic method is to define three sets of backups, such as daily, weekly, and monthly. The daily or son backups are rotated on a daily basis with one graduating to father status each week. The weekly or father backups are rotated on a weekly basis with one graduating to grandfather status each month. In addition, quarterly, biannual, and/or annual backups can also be separately retained. Often one or more of the graduated backups are removed from the site for safekeeping and disaster recovery purposes.

Reference: http://en.wikipedia.org/wiki/Backup_rotation_scheme

QUESTION 82
Offsite tape storage helps in serving which of the following purposes?

Each correct answer represents a complete solution. Choose two.

A. Storage of back up media
B. Misconfiguration issue
C. Interoperability issue
D. Disaster recovery

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Offsite tape storage helps in serving two purposes for these computer reliant organizations. Firstly, disaster recovery, when an organization is in the throes of a major technical or environmental disaster, it ensures that any data that has been recently produced can be reproduced at the earliest possible convenience. This is generally obtained by requesting from offsite tape storage that can be restored to ensure that the organization can be up and running as soon as possible.

Secondly, storage of backup media and computer backups are important in situations where corporate liability and culpability are concerned.

QUESTION 83
Which of the following backup implementation methods is a process of backing up of server data to a shared, central storage device without sending the data over the local area network (LAN)?

A. Serverless
B. Server-based
C. Server-free
D. LAN-free

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
A LAN-free backup is a process of backing up of server data to a shared, central storage device without sending the data over the local area network (LAN). It is usually achieved by using a storage area network (SAN). The goal of LAN-free backup is to reduce the load on LAN and reduce the time it takes to complete the backup. It helps in offering an alternative way of backup than a simple data copy to network-attached storage (NAS) over LAN.

It comes in different flavors:
- With backup server: In addition to a shared storage device (usually a traditional tape library), there exists a central server arbitrating access to device (for all the other SAN servers). The central server, however, does not handle data stream itself.
- Without backup server: The storage facility (usually a virtual tape library, or VTL) is smart enough to handle multiple data accesses without intermediate component.

The advantages of LAN-free backup contain shorter backup and recovery times and less disruption to other systems and applications.

Answer option B is incorrect. Server-based backup uses a server with large disk capacity for backing up data on other servers and workstations in the network. The data on this server is then copied to portable disks or tape for storage offsite. This solution helps in permitting for server-to-server replication. Using a second server

placed at an offsite location and specialized software for replicating the data from one server to another over the Internet permits a user to copy the data from the backup server to the offsite server.

The benefit of using server-based backup accrues to those organizations that have huge amount of data. The speed of the backup increases because server-based disk drives is much faster than tape or portable disks. Also, the disk sizes on the server are much larger than what is currently present on portable drives. Server-based backup also permits for the copy process to take place during the day when other servers are being used for productive work.

Answer option A is incorrect. Serverless backup is a process of offloading backup procedures from a server.

Answer option C is incorrect. In server-free backup, the backup task is achieved without the use of the server.

Reference: http://en.wikipedia.org/wiki/LAN-free_backup

QUESTION 84
Which of the following backup methods takes the maximum amount of time while restoring data?

A. Shadow copy
B. Full backup
C. Incremental backup
D. Differential backup

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Incremental backups take the longest to restore, but they are the fastest to perform as compared to full and differential backups. An incremental backup backs up files that are created or changed since the last full or incremental backup. An incremental backup provides a faster method of backing up data than most other backup methods. Restoring data from an incremental backup requires the last full backup and all subsequent incremental backups. Incremental backups must be restored in the same order as they were created.

If any incremental backup in the incremental backup set is damaged or becomes corrupt, the data backed up after corruption cannot be restored.

QUESTION 85
Which of the following backup technologies relies on the Application Programming Interface (API) for the Extensible Storage Engine (ESE)?

A. Legacy streaming backup
B. Copy backup
C. Tape backup
D. Disk-to-disk backup

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
The Legacy Streaming backup technology relies on the Application Programming Interface (API) for the Extensible Storage Engine (ESE).

Answer option B is incorrect. A copy backup is similar to a full backup except for the archive bit. It does not rely on the Application Programming Interface (API) for the Extensible Storage Engine (ESE).

Answer option C is incorrect. In computers, tape backup is the ability to periodically copy the contents of all or a

designated amount of data from its usual storage device to a tape cartridge device so that, in the event of a hard disk crash or comparable failure, the data will not be lost. Tape backup can be done manually or, with appropriate software, be programmed to happen automatically.

Answer option D is incorrect. The term "disk-to-disk" or "D2D" generally refers to disk-to-disk backup. With D2D, a computer hard disk is backed up to another hard disk rather than to a tape or floppy. D2D is often confused with virtual tape, but differs in that it enables multiple backup and recovery operations to simultaneously access the disk directly by using a true file system.

QUESTION 86
Which of the following encryption methods is used in conjunction with file system-level encryption in order to provide a more secure implementation?

A. Disk encryption
B. Host encryption
C. Network encryption
D. Tape encryption

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Disk encryption uses disk encryption software or hardware for encrypting every bit of data that goes on a disk or disk volume. Disk encryption is used to prevent unauthorized access to data storage. The term "full disk encryption" (or whole disk encryption) is often used for signifying that everything on a disk is encrypted, including the programs that can encrypt bootable operating system partitions. But they must still leave the master boot record (MBR), and thus part of the disk, unencrypted. There are, however, hardware-based full disk encryption systems that can truly encrypt the entire boot disk, including the MBR.

Disk encryption is sometimes used in conjunction with file system-level encryption in order to provide a more secure implementation. Since disk encryption generally uses the same key for encrypting the whole volume, all data is decryptable when the system runs. However, some disk encryption solutions use multiple keys for encrypting different partitions.

Answer option D is incorrect. Tape encryption is one of the encryption methods used for data security. In tape encryption, encryption is implemented in the tape drive, which encrypts the data before it is written to the cartridge. The tape drive first compresses the data to be written and then encrypts it when tape compression is enabled. This relates to the fact that there is no loss of capacity with Tape Encryption technology. The encrypted data will be less compressed if the encryption solution encrypts the data first and then tries to compress the data. The tape drive requires a key which helps in encrypting the data; this key is in encrypted form so as to make the Tape Encryption solution more secure. The encryption key is offered by Encryption Key Manager.

Answer option C is incorrect. Network encryption (sometimes called network layer, or network level encryption) is a network security process that helps in applying crypto services at the network transfer layer - above the data link level, but below the application level.

Network encryption is implemented through Internet Protocol Security (IPSec). It is a set of open Internet Engineering Task Force (IETF) standards that is used in conjunction for creating a framework for private communication over IP networks. IPSec works through the network architecture in which end users and applications don't need to be altered in any way. Encrypted packets appear to be identical to unencrypted packets and can be easily routed through any IP network.

Answer option B is incorrect. Host encryption is a host based encryption solution that helps in providing compliance and maintaining data confidentiality for data from the host to disk on storage. The main advantage of host encryption is that it helps in protecting information in storage environment in the event it becomes compromised through unauthorized access or disk removal.

Host encryption allows users to choose the LUNs or volumes that include sensitive data and needs to be encrypted. There is no requirement of encrypting the entire environment or array, so management is minimized. Host encryption secures the data from the host to disk on storage where it resides. Following are the features of host encryption:

- Security compliance
- Cost-effective encryption
- Support for high availability
- Ability to scale
- Ease of deployment

QUESTION 87
In which of the following is encryption implemented in the tape drive, which encrypts the data before it is written to the cartridge?

A. Host encryption
B. Tape encryption
C. Disk encryption
D. Network encryption

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Tape encryption is one of the encryption methods used for data security. In tape encryption, encryption is implemented in the tape drive, which encrypts the data before it is written to the cartridge. The tape drive first compresses the data to be written and then encrypts it when tape compression is enabled. This relates to the fact that there is no loss of capacity with Tape Encryption technology. The encrypted data will be less compressed if the encryption solution encrypts the data first and then tries to compress the data. The tape drive requires a key which helps in encrypting the data; this key is in encrypted form so as to make the Tape Encryption solution more secure. The encryption key is offered by Encryption Key Manager.

Answer option A is incorrect. Host encryption is a host based encryption solution that helps in providing compliance and maintaining data confidentiality for data from the host to disk on storage. The main advantage of host encryption is that it helps in protecting information in storage environment in the event it becomes compromised through unauthorized access or disk removal.

Host encryption allows users to choose the LUNs or volumes that include sensitive data and needs to be encrypted. There is no requirement of encrypting the entire environment or array, so management is minimized. Host encryption secures the data from the host to disk on storage where it resides. Following are the features of host encryption:

- Security compliance
- Cost-effective encryption
- Support for high availability
- Ability to scale
- Ease of deployment

Answer option D is incorrect. Network encryption (sometimes called network layer, or network level encryption) is a network security process that helps in applying crypto services at the network transfer layer - above the data link level, but below the application level.

Network encryption is implemented through Internet Protocol Security (IPSec). It is a set of open Internet Engineering Task Force (IETF) standards that is used in conjunction for creating a framework for private communication over IP networks. IPSec works through the network architecture in which end users and applications don't need to be altered in any way. Encrypted packets appear to be identical to unencrypted packets and can be easily routed through any IP network.

Answer option C is incorrect. Disk encryption uses disk encryption software or hardware for encrypting every bit of data that goes on a disk or disk volume. Disk encryption is used to prevent unauthorized access to data

storage. The term "full disk encryption" (or whole disk encryption) is often used for signifying that everything on a disk is encrypted, including the programs that can encrypt bootable operating system partitions. But they must still leave the master boot record (MBR), and thus part of the disk, unencrypted. There are, however, hardware-based full disk encryption systems that can truly encrypt the entire boot disk, including the MBR.

Disk encryption is sometimes used in conjunction with file system-level encryption in order to provide a more secure implementation. Since disk encryption generally uses the same key for encrypting the whole volume, all data is decryptable when the system runs. However, some disk encryption solutions use multiple keys for encrypting different partitions.

QUESTION 88
Which of the following is a term in computer security that refers to the ability of people to physically gain access to a computer system?

A. ACL
B. Physical access
C. Host encryption
D. Encryption key

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Physical access is a term in computer security. It refers to the ability of people to physically gain access to a computer system. According to Gregory White, "Given physical access to an office, the knowledgeable attacker will quickly be able to find the information needed to gain access to the organization's computer systems and network."

Physical access opens up a variety of avenues for hacking. Physical access also permits hardware key loggers to be installed. An intruder may be able to boot from a CD or other external media and then read unencrypted data on the hard drive. They may also exploit a lack of access control in the boot loader. One could also use a rogue device for accessing a poorly secured wireless network; if the signal were sufficiently strong, one might not even need to breach the perimeter.

Answer option A is incorrect. An access control list (ACL) is an ordered list of access control entries (ACEs). Each ACE identifies a trustee and specifies a set of access rights allowed, denied, or audited for that trustee. A security descriptor of an object contains two ACL types. They are as follows:
- Discretionary Access Control List (DACL): It identifies a specified trustee that is allowed or denied access to a securable object.
- System Access Control List (SACL): It enables an administrator to log attempts for accessing a secured object.

Answer option D is incorrect. Encryption key is a piece of information (a parameter) that determines the functional output of a cryptographic algorithm or cipher.

Answer option C is incorrect. Host encryption is a host based encryption solution that helps in providing compliance and maintaining data confidentiality for data from the host to disk on storage.

Reference: http://en.wikipedia.org/wiki/Physical_access

QUESTION 89
Which of the following is a persistent list, commonly composed of Access Control Entries (ACEs), that is used to enumerate the rights of principals (users and groups of users and/or groups) for accessing resources?

A. Access control list
B. Vaulting
C. CIFS

D. Physical access

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
ACL is a persistent list, commonly composed of Access Control Entries (ACEs) that is used to enumerate the rights of principals (users and groups of users and/or groups) for accessing resources.

An access control list (ACL) is an ordered list of access control entries (ACEs). Each ACE identifies a trustee and specifies a set of access rights allowed, denied, or audited for that trustee. A security descriptor of an object contains two ACL types. They are as follows:

- Discretionary Access Control List (DACL): It identifies a specified trustee that is allowed or denied access to a securable object.
- System Access Control List (SACL): It enables an administrator to log attempts for accessing a secured object.

Answer option D is incorrect. Physical access is a term in computer security. It refers to the ability of people to physically gain access to a computer system. According to Gregory White, "Given physical access to an office, the knowledgeable attacker will quickly be able to find the information needed to gain access to the organization's computer systems and network."

Physical access opens up a variety of avenues for hacking. Physical access also permits hardware key loggers to be installed. An intruder may be able to boot from a CD or other external media and then read unencrypted data on the hard drive. They may also exploit a lack of access control in the boot loader. One could also use a rogue device for accessing a poorly secured wireless network; if the signal were sufficiently strong, one might not even need to breach the perimeter.

Answer option C is incorrect. The Common Internet File System (CIFS) also known as Server Message Block (SMB) is a network protocol used in sharing files on a LAN. The protocol allows a client to access servers, files, and printers on the LAN. CIFS has a client and a server part. The server is probably more applicable as it could expose the repository as a CIFS server. Repository would show up in the Network Neighbourhood and users would be able to map the repository as a normal Windows drive letter.

Answer option B is incorrect. Vaulting is the strategy of sending critical data out of the main location (off the main site) as part of a disaster recovery plan. The storage of off-site data is also known as vaulting, as backups are stored in purpose built vaults. Using removable storage media such as magnetic tape or optical storage, data is usually transported off-site.

Reference: http://www.snia.org/education/dictionary/a

QUESTION 90
Which of the following actions can a user perform that has Write permission on a file?

Each correct answer represents a complete solution. Choose two.

A. Access the file's content
B. View the file
C. Moving the file to another folder
D. Delete the file

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
The Write permission on a file permits a user for viewing or accessing the file's contents.

Read, Write, Modify, and Full Access permissions
Read, Write, Modify, and Full Access are fundamental permissions that can be assigned to files and folders. These permissions may be summarized as follows:

Read: Permits viewing and listing of files and subfolders - Permits viewing or accessing of the file's contents
Write: Permits adding of files and subfolders - Permits writing to a file
Modify: Permits reading and writing of files and subfolders; allows deletion of the folder - Permits reading and writing of the file; allows deletion of the file
Full Access: Permits reading, writing, changing, and deleting of files and subfolders - Permits reading, writing, changing and deleting of the file

Answer options C and D are incorrect. To delete a file or to move the file to another folder, user requires Full Access permission.

QUESTION 91
You work as a Storage Administrator for uCertify.com. An employee of your organization wants permission on a folder to view some information. You are required to provide him appropriate permissions to enable him to view that information. Which of the following steps will you take to accomplish the task?

A. Provide him the write permission on the file.
B. Provide him the read permission on the file.
C. Provide him the read permission on the folder.
D. Provide him the write permission on the folder.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
While providing permissions, it's important to assign only the required permissions. From a permissions standpoint, all computers or applications must adhere to a principle of 'least privilege' such that authorized users will not have access beyond the permissions required to perform their authorized job functions.

According to the question, the employee needs to view some information. In order to enable him to perform his task, you should provide him the read permission on the file. The read permission permits viewing or accessing of the file's content.

Answer options D and C are incorrect. Providing folder permission to the employee will give him access to other files of that folder.

Answer option A is incorrect. Providing the write permission will permit him writing to the file.

QUESTION 92
Input/Output Operations Per Second (IOPS) is still the most common metric in use when it comes to measuring a storage system's overall performance. What are the factors that help in calculating the IOPS capability of an individual storage system?

Each correct answer represents a complete solution. Choose three.

A. Average peak time
B. Rotational speed
C. Average latency
D. Average seek time

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:
Input/Output Operations Per Second (IOPS) is still the most common metric in use when it comes to measuring a storage system's overall performance. There are a number of factors that help in calculating the IOPS capability of an individual storage system.

Every disk in a storage system has a maximum theoretical IOPS value that depends on a formula. Disk performance and IOPS are based on three key factors:

- Rotational speed: It is measured in revolutions per minute (RPM). A higher rotational speed is related with a higher performing disk. The rotational speed value is not used directly in calculations, but it is highly important.
- Average latency: It is the time it takes for the sector of the disk being accessed to rotate into position under a read/write head.
- Average seek time: It is defined as the time (in ms) it takes for the hard drive's read/write head to position itself over the track being read or written. There are both read and write seek times; the average of the two values is taken.

In order to calculate the IOPS range, use this formula:

Average IOPS: Divide 1 by the sum of the average latency in ms and the average seek time in ms (1 / (average latency in ms + average seek time in ms).

QUESTION 93
Replication is the process of copying data from one system to another across a network. What are the benefits offered by replication?

Each correct answer represents a complete solution. Choose three.

A. Increased reliability
B. Improved system throughput
C. Host platform dependent
D. Increased availability

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
Following are the benefits of replication:

- Increased availability: One of the important advantage of replication is that it helps in marking and tolerating failures in the network gracefully. It helps in providing alternate copies of a replicated data when the primary copy is unavailable.
- Increased reliability: Many applications need extremely high reliability of their data stored in files. Replication is very beneficial for such applications as it permits the existence of multiple copies of their files.
- Improved system throughput: Replication improves system throughput by enabling multiple client's requests for access to the same file to be serviced in parallel by different servers.

Other benefits of replication include minimal client network load, host platform independency, network independency, etc.

Reference: http://www.snia.org/sites/default/education/tutorials/2011/spring/virtualization/PeglarRob-Virtualization_I.pdf

QUESTION 94
Which of the following refers to a failed attempt to read or write a piece of data in the cache?

A. Cache miss
B. Destaging

C. Queue depth
D. Cache

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
A cache miss refers to a failed attempt to read or write a piece of data in the cache, which results in a main memory access with much longer latency.

Answer option B is incorrect. A destaging is a process of updating data or parity in the disks from the write cache.

Answer option C is incorrect. Queue depth is something that most storage administrators need to overlook as this is the reason for most performance issues. Queue depth is defined as the number of commands that the hba can send / receive in a single chunk - per LUN. From a host-hba point (initiator), it is defined as the number of commands that can be queued (or stored) and then sent to storage. From a storage point (target), it is defined as the number of commands it can accept in one-shot, again, per lun.

Performance degradation starts quickly once user reaches max queue depth on target and it starts backing off the queues and it will hit hard on user's response times. Ideally, if users have few hosts and luns, the default is good. If users have the luxury to dedicate ports for high performance applications, ensure that the queue depth is configured appropriately so that users get the best of performance.

Answer option D is incorrect. Cache is a component that transparently stores data so that future requests for that data can be served faster.

QUESTION 95
Which of the following RAID levels are included in RAID 10?

Each correct answer represents a part of the solution. Choose two.

A. RAID 1
B. RAID 5
C. RAID 2
D. RAID 0

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
RAID-10 is a combination of RAID-1 and RAID-0. It is implemented at hardware level rather than on the operating system. It is used to connect mirrored disk pairs to form a RAID-0 array. Data is written on the striped set of disk array as in RAID-0 and then it is mirrored as in RAID-1. Although expensive, RAID-10 provides better fault tolerance as well as input/output performance. The diagrammatic representation of RAID 10 is shown below, where A, B, C, D, E and F represent blocks:

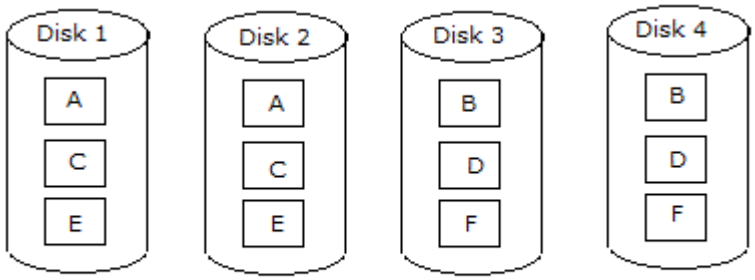

RAID 10

Reference: http://www.ecs.umass.edu/ece/koren/architecture/Raid/basicRAID.html

QUESTION 96
Which of the following performance measurements are taken into consideration for determining the selection of a physical device in the tiered storage model?

Each correct answer represents a part of the solution. Choose two.

A. Age of data
B. RAID capabilities
C. I/Os per second
D. Time to access

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
I/Os per second, RAID capabilities, and performance criteria for the physical devices are taken into consideration for deciding which physical device would go inside which tier in the tiered storage model.

Answer options A and D are incorrect. Age of data and time to access are taken into consideration when trying to decide what data goes in which tier in the tiered storage model.

Tiered storage
Tiered storage is the assignment of different categories of data to different types of storage media for reducing total storage cost. Categories may be based on levels of protection required, performance requirements, frequency of use, and other considerations. For examples:
- Disk and Tape: Two separate storage tiers identified by differences in all four defining attributes.
- Old technology disk and new technology disk: Two separate storage tiers identified by differences in one or more of the attributes.
- High performing disk storage and less expensive, slower disk of the same capacity and function: Two separate tiers.
- Identical Enterprise class disk configured to utilize different functions such as RAID level or replication: A separate storage tier for each set of unique functions.

QUESTION 97
_____ _____ is defined as the number of commands that the HBA can send/receive in a single chunk per LUN.

Correct Answer: Queue depth
Section: (none)
Explanation

Explanation/Reference:

Queue depth is something that most storage administrators do not need to overlook as this is the reason for most performance issues. Queue depth is defined as the number of commands that the hba can send / receive in a single chunk - per LUN. From a host-hba point (initiator), it is defined as the number of commands that can be queued (or stored) and then sent to storage. From a storage point (target), it is defined as the number of commands it can accept in one-shot, again, per lun.

Performance degradation starts quickly once user reaches max queue depth on target and it starts backing off the queues and it will hit hard on user's response times. Ideally, if users have few hosts and luns, the default is good. If users have the luxury to dedicate ports for high performance applications, ensure that the queue depth is configured appropriately so that users get the best of performance.

QUESTION 98
Which of the following is the automated progression or demotion of data across different tiers (types) of storage devices and media?

A. Queue depth
B. Automated tiered storage
C. Cache
D. Hierachical storage management

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Automated tiered storage is the automated progression or demotion of data across different tiers (types) of storage devices and media. This movement of data is automatic to the different types of disk according to performance and capacity requirements.

Answer option D is incorrect. Hierarchical storage management (HSM) is a data storage technique which automatically moves data between high-cost and low-cost storage media. HSM systems exist because high-speed storage devices, such as hard disk drive arrays, are more expensive (per byte stored) than slower devices, such as optical discs and magnetic tape drives. Instead, HSM systems store the bulk of the enterprise's data on slower devices, and then copy data to faster disk drives when needed. In effect, HSM turns the fast disk drives into caches for the slower mass storage devices. The HSM system monitors the way data is used and makes best guesses as to which data can safely be moved to slower devices and which data should stay on the fast devices.

Answer option A is incorrect. Queue depth is something that most storage administrators need to overlook as this is the reason for most performance issues. Queue depth is defined as the number of commands that the hba can send / receive in a single chunk - per LUN. From a host-hba point (initiator), it is defined as the number of commands that can be queued (or stored) and then sent to storage. From a storage point (target), it is defined as the number of commands it can accept in one-shot, again, per lun.

Performance degradation starts quickly once user reaches max queue depth on target and it starts backing off the queues and it will hit hard on user's response times. Ideally, if users have few hosts and luns, the default is good. If users have the luxury to dedicate ports for high performance applications, ensure that the queue depth is configured appropriately so that users get the best of performance.

Answer option C is incorrect. Cache is a component that transparently stores data so that future requests for that data can be served faster. The data that is stored within a cache (values) that have been computed earlier or duplicates of original values that are stored elsewhere. If requested data is contained in the cache (cache hit), this request can be served by simply reading the cache, which is comparatively faster. Otherwise (cache miss), the data has to be recomputed or fetched from its original storage location, which is comparatively slower. Hence the greater the number of requests that can be served from the cache, the faster the overall system performance becomes. To be cost efficient and to enable an efficient use of data, caches are relatively small.

QUESTION 99

Which of the following protocols is designed for use in real-time aircraft control system?

A. FC-SB
B. FC-AV
C. FC-AE
D. FC-SP

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Avionics is technically specialized market for Fibre channel. Fibre channel provides increased throughput and reliability to avionics applications. The Fibre Channel Avionics Environment (FC-AE) protocol has been designed for use in real time aircraft control system.

Answer option B is incorrect. The Fibre Channel Audio Video (FC-AV) protocol maps digital A/V formats onto Fibre Channel.

Answer option A is incorrect. FC-SB (Fibre Channel Single Byte) is the industry-standard command protocol for ESCON (Enterprise Systems Connection), a type of data connection, over Fibre Channel.

Answer option D is incorrect. The Fibre Channel Security Protocol (FC-SP) defines various protocols used for implementing security in a fibre channel fabric.

QUESTION 100
Which of the following is a computer system monitor tool used for collecting and showing operating system storage input and output statistics?

A. sysmon
B. perfmon
C. inputstat
D. iostat

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Iostat (input/output statistics) is a computer system monitor tool used for collecting and showing operating system storage input and output statistics. It is often used for identifying performance issues with local disks or networked file systems such as NFS.

Answer option A is incorrect. Sysmon is a network monitoring tool that is designed to give high performance and accurate network monitoring. Currently, supported protocols contain SMTP, IMAP, HTTP, TCP, UDP, NNTP, and PING tests.

Answer option B is incorrect. PerfMon is a powerful troubleshooting tool used for generating a log or multiple logs that can be graphically analyzed to identify problem areas in system performance.

Answer option C is incorrect. This is an invalid option.

Exam B

QUESTION 1
Which of the following is the system component that receives a SCSI I/O command?

A. Connection
B. Initiator
C. Target
D. SAN Fabric

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
A target is the system component that receives a SCSI I/O command. The iSCSI specification refers to a storage resource located on an iSCSI server as a target. An iSCSI target is often a dedicated network-connected hard disk storage device, but may also be a general-purpose computer. Since with initiators, software for providing an iSCSI target is available for most mainstream operating systems.

Answer option B is incorrect. Initiator is the system component that originates an I/O command over an I/O bus or network. It helps in building and issuing a Command Descriptor Block (CDB), forwarding the CDB to a SCSI Target device, handling any interaction associated with that particular command (such as receiving and acknowledging data for a READ command), and checking/accepting a response (Status) from the SCSI Target device.

I/O adapters, network interface cards, and intelligent controller device I/O bus control ASICs are typical initiators.

Answer option D is incorrect. 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. An organization can get more advantageous features, such as data replication, data sharing, and centralized backups by consolidating storage devices into the same SAN fabric.

Answer option A is incorrect. A connection is a communication path between the initiator and target using a TCP/IP connection. In iSCSI, one or more connection makes up a session. A connection helps in carrying control messages, SCSI commands, parameters, and data within iSCSI PDUs.

QUESTION 2
Which of the following encryption methods uses disk encryption software or hardware for encrypting every bit of data that goes on a disk or disk volume?

A. Network encryption
B. Host encryption
C. Disk encryption
D. Tape encryption

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Disk encryption uses disk encryption software or hardware for encrypting every bit of data that goes on a disk or disk volume. Disk encryption is used to prevent unauthorized access to data storage. The term "full disk encryption" (or whole disk encryption) is often used for signifying that everything on a disk is encrypted, including the programs that can encrypt bootable operating system partitions. But they must still leave the

master boot record (MBR), and thus part of the disk, unencrypted. There are, however, hardware-based full disk encryption systems that can truly encrypt the entire boot disk, including the MBR.

Disk encryption is sometimes used in conjunction with file system-level encryption in order to provide a more secure implementation. Since disk encryption generally uses the same key for encrypting the whole volume, all data is decryptable when the system runs. However, some disk encryption solutions use multiple keys for encrypting different partitions.

Answer option D is incorrect. Tape encryption is one of the encryption methods used for data security. In tape encryption, encryption is implemented in the tape drive, which encrypts the data before it is written to the cartridge. The tape drive first compresses the data to be written and then encrypts it when tape compression is enabled. This relates to the fact that there is no loss of capacity with Tape Encryption technology. The encrypted data will be less compressed if the encryption solution encrypts the data first and then tries to compress the data. The tape drive requires a key which helps in encrypting the data; this key is in encrypted form so as to make the Tape Encryption solution more secure. The encryption key is offered by Encryption Key Manager.

Answer option A is incorrect. Network encryption (sometimes called network layer, or network level encryption) is a network security process that helps in applying crypto services at the network transfer layer - above the data link level, but below the application level.

Network encryption is implemented through Internet Protocol Security (IPSec). It is a set of open Internet Engineering Task Force (IETF) standards that is used in conjunction for creating a framework for private communication over IP networks. IPSec works through the network architecture in which end users and applications don't need to be altered in any way. Encrypted packets appear to be identical to unencrypted packets and can be easily routed through any IP network.

Answer option B is incorrect. Host encryption is a host based encryption solution that helps in providing compliance and maintaining data confidentiality for data from the host to disk on storage. The main advantage of host encryption is that it helps in protecting information in storage environment in the event it becomes compromised through unauthorized access or disk removal.

Host encryption allows users to choose the LUNs or volumes that include sensitive data and needs to be encrypted. There is no requirement of encrypting the entire environment or array, so management is minimized. Host encryption secures the data from the host to disk on storage where it resides. Following are the features of host encryption:

- Security compliance
- Cost-effective encryption
- Support for high availability
- Ability to scale
- Ease of deployment

QUESTION 3
_____ is defined as a point-in-time copy of a defined collection of data.

Correct Answer: Snapshot
Section: (none)
Explanation

Explanation/Reference:
Snapshot is a virtual representation of a volume at a given point and is ideal for short-term backup and recovery purposes. Snapshot uses less space since only modified data is copied from the primary.

Snapshot is a point-in-time copy of a defined collection of data. It is space-efficient, any-point-in-time copies of changed bits or volumes, achieving time, cost and capacity efficient. Snapshot can be of files, LUNs, file systems, or any other type of container supported by the system it can be read-only or read-write copies of files, LUNs, file systems or any other type of container.

What is snapshot replication?
Snapshot replication copies data exactly as they appear at one time. In snapshot replication, Subscribers are

updated with a complete refresh of the data set instead of individual transactions. Therefore, if you insert a record and then remove it, snapshot replication will not replicate these transactions, as there is no change in data. Because snapshot replication replicates an entire data set at one time, it can be scheduled to run during off-peak hours.

Snapshot replication is used when:
- Data is mostly static and does not change often. Even if it changes, it makes more sense to publish an entirely new copy to Subscribers.
- It is acceptable to have copies of data that are out of date for a period of time.
- Replicating small volumes of data in which an entire refresh of the data is reasonable.

Reference: http://www.snia.org/education/dictionary/s

QUESTION 4
Which of the following provides the ability to a single physical Fibre Channel node or switch to support more than one Nx_Port on a single point-to-point link?

A. SCSI ID
B. Domain ID
C. Zone alias
D. NPIV

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
N_Port ID Virtualization, or NPIV, is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port. It allows multiple Fibre Channel initiators to occupy a single physical port, easing hardware requirements in Storage Area Network design, especially where virtual SANs are called for. NPIV provides the ability to a single physical Fibre Channel node or switch to support more than one Nx_Port on a single point-to-point link.

Answer option A is incorrect. A SCSI ID is a unique identifier that is assigned to each SCSI device on the bus.

Answer option B is incorrect. A Domain ID is a number that uniquely identifies a switch in a fabric and is the highest or most significant hierarchical level in the three-level address hierarchy.

Answer option C is incorrect. A zone alias is a collection of zone members. A zone alias can be added to one or more zones.

Reference: http://www.snia.org/education/dictionary/n

QUESTION 5
Which of the following storage industry terms is used as a flow control method by Fibre Channel technology and represents the number of frames a port can store?

A. Buffer-to-buffer credit
B. Oversubscription
C. Link
D. Worldwide node name

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Buffer credit, also known as buffer-to-buffer credit (BBC) is used as a flow control method by Fibre Channel

technology and represents the number of frames a port can store. It is used to permit data communication in a Fibre Channel storage area network (SAN) where there are long spans of fiber optic cable.

In buffer-credit flow control, the source and destination sets the number of unacknowledged frames (buffer credits) permitted to accumulate before the source stops sending data. This is a power of 2, such as 4, 8, or 16. In some SANs it can range higher, but it is rare to have more than 256 buffer credits because of cost constraints. A counter at the source keeps track of the number of buffer credits. Each time a frame is sent by the source, the counter increments by 1. Each time the destination receives a frame, it sends an acknowledgment back to the source, which decrements the counter by 1. If the number of buffer credits reaches the maximum, the source stops transmission until it receives the next acknowledgement from the destination. This helps in preventing loss of frames that may result if the source races too far ahead of the destination.

Answer option C is incorrect. A link is defined as a physical connection (electrical or optical) between two nodes of a network. It can also be defined as the point-to-point physical connection from one element of a Fibre Channel fabric to the next.

Answer option B is incorrect. In a SAN (storage area network) switching environment, oversubscription is the practice of connecting multiple devices to the same switch port in order to optimize switch use. Each SAN port can support a particular communication speed and a Fibre Channel switch may offer 1 Gb, 2 Gb, or 4 Gb FC ports. However, because ports are rarely run at their maximum speed for a prolonged period, multiple slower devices may fan in to a single port to take advantage of unused capacity. For instance, a single storage server may not be able to sustain 4 Gbps to a switch port, so two 2 Gb servers or four 1 Gb servers may all be aggregated to that 4 Gb switch port.

Answer option D is incorrect. Worldwide node name is a Node_Name that is worldwide unique. A World Wide Node Name, WWNN, or WWnN, is a World Wide Name assigned to a node (an endpoint, a device) in a Fibre Channel fabric. It is valid for the same WWNN to be seen on many different ports (different addresses) on the network, identifying the ports as multiple network interfaces of a single network node.

Reference: http://searchstorage.techtarget.com/definition/buffer-credits

QUESTION 6
In which of the following topologies are two devices connected directly to each other?

A. Point-to-point
B. Redundant fabric
C. Single fabric
D. Arbitrated loop

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
The Point-to-point topology is also known as FC-P2P. In this design, two devices are connected directly to each other. This is the simplest topology with limited connectivity. The point-to-point topology enables a user to interconnect ports directly.

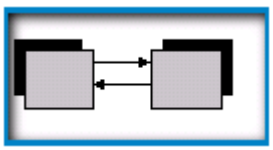

Answer option D is incorrect. Arbitrated loop, also known as FC-AL, is a Fibre Channel topology in which devices are connected in a one-way loop fashion in a ring topology. It is a serial architecture that is compatible with the SCSI, handling up to 127 ports (devices). One port may optionally connect a loop to a fabric switch port. The bandwidth on the loop is shared among all ports. Only two ports may communicate at a time on the

loop. One port wins arbitration and may open one other port in either half or full duplex mode.

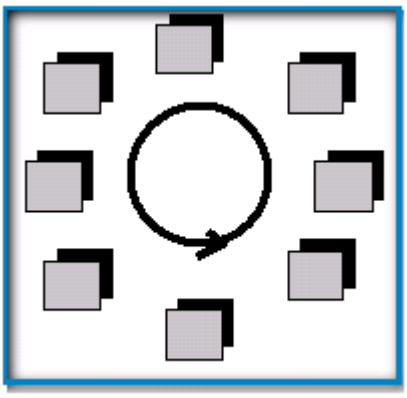

Answer option C is incorrect. The single fabric topology is used for connecting many (224) devices in a cross-point switched configuration. The advantage of this topology is that many devices can communicate at the same time; the media is not shared.

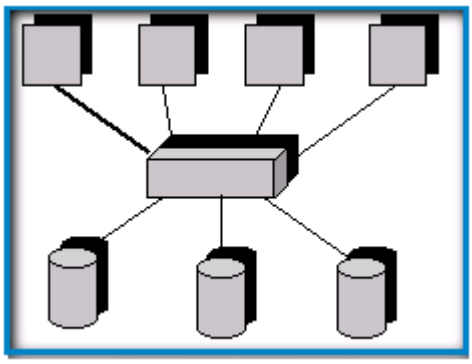

Answer option B is incorrect. The redundant fabric topology is used for connecting dual connected hosts and dual connected storage. Using redundant fabrics essentially doubles the maximum size of a SAN. When deploying redundant fabrics, it is not always important to deploy symmetric fabrics.

QUESTION 7
Which of the following connectors are used with fiber-optic cables?

Each correct answer represents a complete solution. Choose three.

A. ST
B. SC
C. F-Type
D. MT-RJ

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
The following connectors are used with fiber-optic cables:

- SC connector: A subscriber connector (SC) is a fiber-optic connector used with multimode fiber. It is a square shaped connector used for terminating fiber optic cables. SC connectors have a push-pull latching mechanism to provide quick insertion and removal while also ensuring a positive connection.

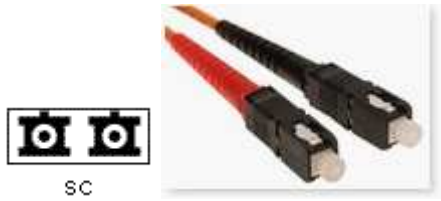

- ST connector: A straight tip (ST) connector is a fiber-optic connector used with multimode fiber. An ST connector has a 2.5mm shaft and bayonet locking ring, and allows quick connect and disconnect of 125 micron multi-mode fiber.

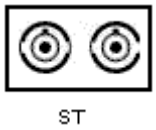

- MT-RJ connector: The MT-RJ connector is the most recent type of small form factor fiber optic connector. The MT-RJ fiber connector is less than half the size of the SC duplex connector and transceiver, so it doubles the port density of fiber-optic LAN equipment. The connector is a 2-fiber connector and takes up no more room than an RJ-45 jack.

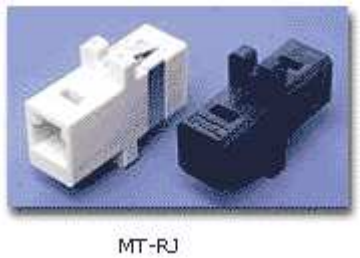

Answer option C is incorrect. An F-type connector is a threaded medium performance coaxial signal connector, which is used in TVs and VCRs. The pin of the connector is actually the center conductor of the coaxial cable. It is an inexpensive connector.

QUESTION 8
Which of the following statements are true regarding a NAS gateway?

Each correct answer represents a complete solution. Choose three.

A. It works as integrated disk storage.
B. It provides file serving to clients and servers.
C. It has clustering and failover features.
D. It provides connection to an IP network.

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:
NAS head or gateway is hardware that performs the NAS control functions. The clients always connect to the

NAS head, as it is the NAS head that is addressable on the network. NAS head is independent of the storage devices and contains an imbedded operating system that does not need a keyboard, mouse, or monitor. The NAS head includes feature like clustering, failover, and file serving.

Reference: http://www.lascon.co.uk/d004004.htm

QUESTION 9
What are the benefits of using network attached storage?

Each correct answer represents a complete solution. Choose three.

A. All applications support NAS technology.
B. NAS devices permit administrators to implement simple, low cost load-balancing, and fault tolerant systems.
C. NAS volumes can be mounted by multiple clients at a time.
D. NAS helps in removing the responsibility of file serving from other servers on the network.

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:
Following are the benefits of using network attached storage:
- NAS volume can be mounted by multiple clients at a time.
- NAS devices permit administrators to implement simple, low cost load-balancing, and fault tolerant systems.
- NAS helps in removing the responsibility of file serving from other servers on the network.
- The NAS technology can be easily implemented using the same protocols as file servers.
- In NAS, clients that are running various operating systems can share the same files.

Answer option A is incorrect. The drawback to a NAS is that all applications will not support it since they are expecting a block-level storage device, and most clustering solutions are designed to run on a SAN.

Reference: http://matteocappelli.files.wordpress.com/2011/02/storage_basics.pdf

QUESTION 10
Which of the following is a component of SMI-S that represents an object-oriented description of the entities and relationships in a business management environment, divided into core and common models?

A. CIM
B. SLP
C. ILM
D. WBEM

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
CIM (Common Information Model) is the key component of SMI-S. It provides a common definition of management information for systems, networks, applications and services, and allows for vendor extensions. It represents an object-oriented description of the entities and relationships in a business management environment, divided into core and common models.

Answer option D is incorrect. Web-based Enterprise Management (WBEM) is a set of management and internet standard technologies developed to unify the management of enterprise computing environments.

Answer option C is incorrect. ILM stands for Information Lifecycle Management. It refers to the creation and management of a storage infrastructure.

Answer option B is incorrect. Service Location Protocol (SLP) enables computers and other devices to find services in a local area network without prior configuration.

QUESTION 11
Which of the following ports connects to another Fibre Channel switch or bridge device via an inter-switch link?

A. F_Port
B. N_Port
C. E_Port
D. G_Port

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
The "Expansion" port within a Fibre Channel switch connects to another Fibre Channel switch or bridge device via an inter-switch link. E_Ports are used to link Fibre channel switches to form a multi-switch fabric.

Answer option D is incorrect. A "Generic" Port can operate as either an E_Port or an F_Port. A G_Port can determine operating mode at switch port initialization, F_Port when an N_Port attachment is determined, E_Port when an E_Port attachment is determined.

Answer option A is incorrect. The "Fabric" port within a Fibre Channel fabric switch provides a point-to-point link attachment to a single N_Port. F_Ports are intermediate ports in virtual point-to-point links between end ports, for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch.

Answer option B is incorrect. A "Node" port connects via a point-to-point link to either a single N_Port or a single F_Port. N_Ports handle creation, detection, and flow of message units to and from the connected systems. N_Ports are end ports in virtual point-to-point links through a fabric, for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch.

QUESTION 12
Which of the following physically stores data and protects against errors on a fixed block disk?

A. Sector
B. Spindle
C. Head
D. Cylinder

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
A hard disk comprises one or more circular platters, of which either or both surfaces are coated with a magnetic substance used for recording the data. For each surface, there is a read-write head that can process the recorded data. The surfaces are usually divided into concentric rings, called tracks, and these in turn are divided into sectors as shown in the image:

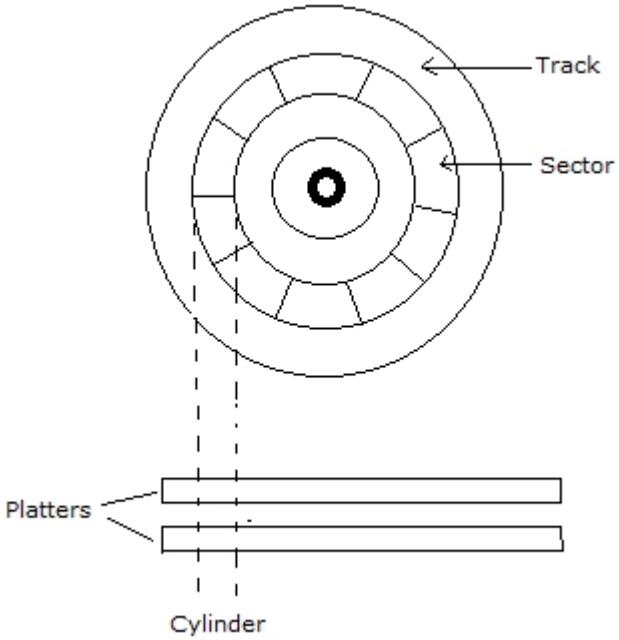

When the head for one surface is on a track, the heads for the other surfaces are also on the corresponding tracks. All the corresponding tracks taken together are called a cylinder.

A sector is the basic unit of data storage on a hard disk. Data is physically stored and protected against errors on a fixed block architecture disk in the sector. Fixed block architecture is the disk model on which SCSI is predicated.

Reference: http://www.snia.org/education/dictionary/s

QUESTION 13
Which of the following terms refers to the condition when the data transfer rate falls below the minimum threshold at which the tape drive heads were designed to transfer data?

A. Multistreaming
B. Multiplexing
C. Shoe-shining
D. Multistoring

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Shoe-shining is a disadvantage with read/write tape drives. During shoe-shining, the data transfer rate falls below the minimum threshold at which the tape drive heads were designed to transfer data to or from a continuously running tape. In this situation, the modern fast-running tape drive is unable to stop the tape instantly. Instead, the drive must decelerate and stop the tape, rewind it a short distance, restart it, position back to the point at which streaming stopped, and then resume the operation. If the condition repeats, the resulting back-and-forth tape motion resembles that of shining shoes with a cloth. Shoe-shining decreases the attainable data transfer rate, and drive and tape life.

In early tape drives, non-continuous data transfer was normal and unavoidable - computer processing power and memory available were usually insufficient to provide a constant stream. So tape drives were typically designed for so called start-stop operation. Early drives used very large spools, which necessarily had high inertia and did not start and stop moving easily. To provide high start, stop, and seeking performance, several

feet of loose tape was played out and pulled by a suction fan down into two deep open channels on either side of the tape head and capstans. The long thin loops of tape hanging in these vacuum columns had far less inertia than the two reels and could be rapidly started, stopped and repositioned.

Later, most tape drives of the 1980s introduced the use of an internal data buffer to somewhat reduce start-stop situations. These drives are often referred to as tape streamers. The tape was stopped only when the buffer contained no data to be written, or when it was full of data during reading.

Most recently, drives no longer operate at a single fixed linear speed, but have several speeds. Internally, they implement algorithms that dynamically match the tape speed level to the computer's data rate. Example speed levels could be 50 percent, 75 percent and 100 percent of full speed. A computer that streams data slower than the lowest speed level (e.g. at 49 percent) will still cause shoe-shining.

Answer options B and A are incorrect. Multiplexing is the process of alternately writing or interleaving multiple clients-data to a single tape. Multistreaming allows multiple backups (streams) of a single client to occur simultaneously to multiple tape drives.

Answer option D is incorrect. Multistoring is not a valid term for tape drives.

QUESTION 14
Which of the following is a globally unique 64-bit identifier assigned to each Fibre Channel node process?

A. WWNN
B. WWPN
C. NPIV
D. BBC

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Worldwide node name is a Node_Name that is worldwide unique. A World Wide Node Name, WWNN, or WWnN, is a World Wide Name assigned to a node (an endpoint, a device) in a Fibre Channel fabric. It is valid for the same WWNN to be seen on many different ports (different addresses) on the network, identifying the ports as multiple network interfaces of a single network node. WWNN is a globally unique 64-bit identifier. It is assigned to each Fibre Channel node process.

Answer option B is incorrect. WWPN is a globally unique 64-bit identifier assigned to each Fibre Channel port.

Answer option C is incorrect. N_Port ID Virtualization, or NPIV, is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port.

Answer option D is incorrect. Buffer credit, also called buffer-to-buffer credit (BBC), is used as a flow control method by Fibre Channel technology and represents the number of frames a port can store.

Reference: http://www.knowledgetransfer.net/dictionary/Storage/en/World_Wide_Node_Name.htm

QUESTION 15
Which of the following is a distributed service needed by the fabric to register and discover the attributes of Fibre Channel N_Ports?

A. Tracert
B. Ping
C. Nslookup
D. Name server

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:
Name server is a distributed service needed by the fabric to register and discover the attributes of Fibre Channel N_Ports. Once registered, the attributes may be viewed by requesting N_ports.

Answer option C is incorrect. NSLOOKUP is a tool for diagnosing and troubleshooting Domain Name System (DNS) problems. It performs its functions by sending queries to the DNS server and obtaining detailed responses at the command prompt. This information can be useful for diagnosing and resolving name resolution issues, verifying whether or not the resource records are added or updated correctly in a zone, and debugging other server-related problems. This tool is installed along with the TCP/IP protocol through the Control Panel.

Answer option B is incorrect. The ping command-line utility is used to test connectivity with a host on a TCP/IP-based network. This is achieved by sending out a series of packets to a specified destination host. On receiving the packets, the destination host responds with a series of replies. These replies can be used to determine whether or not the network is working properly.

Answer option A is incorrect. TRACERT is a route-tracing Windows utility that displays the path an IP packet takes to reach the destination. It shows the Fully Qualified Domain Name (FQDN) and the IP address of each gateway along the route to the remote host.

QUESTION 16
Which of the following is a counter for measuring flow of data across the storage network?

A. Buffer credit
B. ADM
C. Snapshot
D. RTO

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Buffer credits, also called buffer-to-buffer credits (BBC), are used as a flow control method by Fibre Channel technology and represent the number of frames a port can store.

Answer option B is incorrect. Automated Data Migration (ADM) tools automate the migration of data from one storage point to another based on user-defined policy.

Answer option C is incorrect. Snapshot is a virtual representation of a volume at a given point and is ideal for short-term backup and recovery purposes.

Answer option D is incorrect. Recovery Time Objective (RTO) is defined as the maximum acceptable time period needed to bring one or more applications and associated data back from an outage to a correct operational state.

QUESTION 17
You are advising a school district on disaster recovery plans. In case a disaster affects the main IT centers of the district, they will need to be able to work from an alternate location. However, budget is an issue. Which of the following is most appropriate for this client?

A. Cold site
B. Hot site
C. Off site
D. Warm site

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
A cold site provides an office space, and in some cases basic equipment. However, you will need to restore your data to that equipment in order to use it. This is a much less expensive solution than the hot site.

Answer option B is incorrect. A hot site has equipment installed, configured, and ready to use. This may make disaster recovery much faster, but will also be more expensive. And a school district can afford to be down for several hours before resuming IT operations, so the less expensive option is more appropriate.

Answer option D is incorrect. A warm site is between a hot and cold site. It has some equipment ready and connectivity ready. However, it is still significantly more expensive than a cold site, and not necessary for this scenario.

Answer option C is incorrect. Off site is not any type of backup site terminology.

QUESTION 18
You work as a Storage Administrator for uCertify Inc. You are developing a backup policy for the company and want to accomplish the following goals:

- The number of tapes required to back up data should be minimized.
- The number of tapes required to restore data should be minimized.

Which of the following backup policies will you choose?

A. A full backup once a month and an incremental backup daily
B. A differential backup every night
C. A full backup every Friday and differential backups from Monday to Thursday
D. A full backup every Friday and incremental backups from Monday to Thursday

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
In order to minimize the number of tapes required to back up data, you should perform a full backup every Friday. Performing differential backups from Monday to Thursday will reduce the number of tapes required to restore data. This is because in the event of data loss, you will have to restore only the last full backup and the last differential backup.

Differential backup
A differential backup backs up files that are created or changed since the last full backup. It requires the minimum space to back up data. A differential backup requires only the last full backup tape and the last differential backup tape to restore data. It is faster as compared to the full backup.

Full backup
A full backup backs up the entire database including the transaction log. Taking a full backup daily is impractical, as it is time consuming. Instead, a well-defined backup strategy should be implemented as a weekly full backup and a daily differential backup.

Answer option B is incorrect. Performing a full backup is necessary to restore data. Taking only a differential backup will not work.

Answer options D and A are incorrect. An incremental backup requires restoring the last full backup tape and all incremental backup tapes since the last full backup. This will increase the number of tapes to be restored.

Incremental backup

An incremental backup backs up files that are created or changed since the last full or incremental backup. An incremental backup provides a faster method of backing up data than most other backup methods. Restoring data from an incremental backup requires the last full backup and all subsequent incremental backups. Incremental backups must be restored in the same order as they were created.

If any incremental backup in the incremental backup set is damaged or becomes corrupt, the data backed up after corruption cannot be restored.

QUESTION 19
You work as an IT Technician for BlueBell Inc. John, a trainee, wants to know about fire extinguishers that are used to put out fires on computer components. Which of the following classes of fire extinguishers will you tell John to use in the event of a computer-related fire?

A. Class C
B. Class D
C. Class A
D. Class B

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Class C extinguishers are used to put out fires that are electrically energized, such as those in computer equipment. The presence of the letter "C" indicates that the extinguishing agent is non-conductive.

Answer option C is incorrect. Class A extinguishers are used to put out fires in ordinary combustibles such as wood and paper.

Answer option D is incorrect. Class B extinguishers are used to put out fires that are caused by flammable liquids such as grease, petrol, gasoline, oil, etc.

Answer option B is incorrect. Class D extinguishers are designed for use on flammable metals and are often specific for the type of metal.

QUESTION 20
Which of the following is an ordered list of access control entries (ACEs) that identifies a trustee and specifies a set of access rights allowed, denied, or audited for that trustee?

A. Network encryption
B. Physical access
C. Access control list
D. CIFS

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
An access control list (ACL) is an ordered list of access control entries (ACEs). Each ACE identifies a trustee and specifies a set of access rights allowed, denied, or audited for that trustee. A security descriptor of an object contains two ACL types. They are as follows:

- **Discretionary Access Control List (DACL)**: It identifies a specified trustee that is allowed or denied access to a securable object.
- **System Access Control List (SACL)**: It enables an administrator to log attempts for accessing a secured object.

Answer option B is incorrect. Physical access is a term in computer security. It refers to the ability of people to physically gain access to a computer system.

Answer option A is incorrect. Network encryption (sometimes called network layer, or network level encryption) is a network security process that helps in applying crypto services at the network transfer layer - above the data link level, but below the application level.

Network encryption is implemented through Internet Protocol Security (IPSec). It is a set of open Internet Engineering Task Force (IETF) standards that is used in conjunction for creating a framework for private communication over IP networks. IPSec works through the network architecture in which end users and applications don't need to be altered in any way. Encrypted packets appear to be identical to unencrypted packets and can be easily routed through any IP network.

Answer option D is incorrect. Common Internet File System (CIFS) is a protocol that permits programs and makes requests for files and services on remote computers on the Internet. The client/server programming model is used by CIFS. A client program makes a request of a server program (usually in another computer) for accessing a file or passing a message to a program that runs in the server computer. The server takes the requested action and returns a response.

QUESTION 21
Which of the following terms defines the strategy of sending critical data out of the main location (off the main site) as part of a disaster recovery plan?

A. RTO
B. D2D2T
C. VTL
D. Vaulting

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Vaulting is the strategy of sending critical data out of the main location (off the main site) as part of a disaster recovery plan. The storage of off-site data is also known as vaulting, as backups are stored in purpose built vaults. Using removable storage media such as magnetic tape or optical storage, data is usually transported off-site.

Answer option C is incorrect. A virtual tape library (VTL) is a data storage virtualization technology used typically for backup and recovery purposes. A VTL is used to present a storage component (usually hard disk storage) as tape libraries or tape drives for use with existing backup software.

Virtualizing the disk storage as tape permits integration of VTLs with existing backup software and existing backup and recovery processes and policies. The benefits of such virtualization include storage consolidation and faster data restore processes.

By backing up data to disks instead of tapes, VTL often enhances performance of both backup and recovery operations. Restore processes are found to be faster than backup regardless of implementations. In some cases, the data stored on the VTL's disk array is exported to other media, such as physical tapes, for disaster recovery purposes (scheme called disk-to-disk-to-tape, or D2D2T).

Answer option B is incorrect. Disk-to-disk-to-tape (D2D2T) is an approach to computer storage backup and archiving. In D2D2T scheme, data is initially copied to backup storage on disk storage system and then periodically copied again to a tape storage system.

A D2D2T scheme permits the administrator to automate daily backups on disk for implementing fast restores and then move data to tape when an administrator has time. The tape is also used to transfer more mature data offsite for disaster recovery protection and to comply with regulatory policies for long-term data retention at a relatively inexpensive cost.

Answer option A is incorrect. The Recovery Time Objective is the duration of time and a service level within which a business process must be restored after a disaster in order to avoid unacceptable consequences associated with a break in business continuity. It includes the time for trying to fix the problem without a recovery, the recovery itself, tests, and the communication to the users. Decision time for users' representative is not included.

The business continuity timeline usually runs parallel with an incident management timeline and may start at the same or different points.

Reference: http://en.wikipedia.org/wiki/Off-site_data_protection

QUESTION 22
Which of the following is a set of standards for physically connecting and transferring data between computers and peripheral devices?

A. FCP
B. SCSI
C. FCoE
D. FCIP

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
SCSI stands for Small Computer System Interface. It is not an interface standard; however, it is a set of standards for physically connecting and transferring data between computers and peripheral devices. SCSI is an intelligent, peripheral, buffered, peer to peer interface. It hides the complexity of physical format. Every device attaches to the SCSI bus in a similar manner.

Up to 8 or 16 devices can be attached to a single bus. There can be any number of hosts and peripheral devices, but there should be at least one host. SCSI is most commonly used for hard disks and tape drives, but it can connect a wide range of other devices, including scanners and CD drives.

Answer option A is incorrect. Fibre Channel Protocol (FCP) is a transport protocol (similar to TCP used in IP networks), which predominantly transports SCSI commands over Fibre Channel networks.

Answer option D is incorrect. Fibre Channel over IP (FCIP) is an IP-based storage networking technology, developed by the IETF and defined in RFC 3821. FCIP is also known as Fibre Channel tunneling or storage tunneling. FCIP enables the transmission of Fibre Channel information by tunneling data between storage area network facilities over IP networks. This capacity facilitates data sharing over a geographically distributed enterprise. FCIP is a key technology that is expected to help bring about rapid development of the storage area network market by increasing the capabilities and performance of storage data transmission.

Answer option C is incorrect. Fibre Channel over Ethernet (FCoE) is an encapsulation of Fibre Channel frames over Ethernet networks. It allows Fibre Channel to use 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol. FCoE maps Fibre Channel over selected full duplex IEEE 802.3 networks for providing I/O consolidation over Ethernet and reducing network complexity in the datacenter. The FCoE protocol specification replaces the FC0 and FC1 layers of the Fibre Channel stack with Ethernet.

QUESTION 23
Which of the following applications are used by ISPs, ASPs, and online media providers to push content across private networks to multiple points-of-presence?

A. Content Distribution Applications
B. Backup Applications
C. Database Applications

D. Avionic Applications

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Content Distribution Applications are used by ISPs, ASPs, and online media providers to push content across private networks to multiple points-of-presence. These applications require highly scalable bandwidth for both storage I/O and data networks. These are mostly used by video broadcast industries.

Answer option D is incorrect. Avionic applications are used for instrumentation and signal processing of avionic commands and controls.

Answer options C and B are incorrect. Database applications and backup applications are not used by ISPs, ASPs, and online media providers to push content across private networks to multiple points-of-presence.

QUESTION 24
Which of the following allows Fibre Channel to be carried over Ethernet links?

A. FCoE
B. FCIP
C. FCP
D. iFCP

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Fibre Channel over Ethernet (FCoE) is an encapsulation of Fibre Channel frames over Ethernet networks. It allows Fibre Channel to use 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol. FCoE maps Fibre Channel over selected full duplex IEEE 802.3 networks for providing I/O consolidation over Ethernet and reducing network complexity in the datacenter. The FCoE protocol specification replaces the FC0 and FC1 layers of the Fibre Channel stack with Ethernet.

Answer option B is incorrect. Fibre Channel over IP (FCIP) is an IP-based storage networking technology, developed by the IETF and defined in RFC 3821. FCIP is also known as Fibre Channel tunneling or storage tunneling. FCIP enables the transmission of Fibre Channel information by tunneling data between storage area network facilities over IP networks. This capacity facilitates data sharing over a geographically distributed enterprise. FCIP is a key technology that is expected to help bring about rapid development of the storage area network market by increasing the capabilities and performance of storage data transmission.

Answer option C is incorrect. Fibre Channel Protocol (FCP) is a transport protocol (similar to TCP used in IP networks), which predominantly transports SCSI commands over Fibre Channel networks.

Answer option D is incorrect. Internet Fibre Channel Protocol (iFCP) is a gateway to gateway network protocol standard that provides Fibre Channel fabric functionality to fibre channel devices over an IP network. The iFCP protocol enables the implementation of fibre channel functionality over an IP network, within which the fibre channel switching and routing infrastructure is replaced by IP components and technology. The primary objective of iFCP is to allow existing fibre channel devices to be networked and interconnected over an IP based network at wire speeds. The iFCP protocol layer's main function is to transport Fibre Channel frame images between Fibre Channel ports attached both locally and remotely.

QUESTION 25
Which of the following types of fire extinguishers is needed for combustible liquids' fire?

A. Class B

B. Class A
C. Class D
D. Class C

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Class B extinguishers are needed for flammable or combustible liquids' fire, such as gasoline, kerosene, grease, and oil. The numerical rating for a class B extinguisher indicates the approximate area, in square feet, in which it can extinguish fire. Carbon Dioxide extinguishers are used for class B and C fires.

Answer option B is incorrect. Class A extinguishers are needed for ordinary flammable materials' fire, such as paper, wood, cardboard, and most plastics.

Answer option D is incorrect. Class C fire extinguishers are needed for electrical fires. The electrical fire involves electrical equipment, such as appliances, wiring, and circuit breakers. Water should not be used as an extinguisher for this type of fire. Carbon Dioxide is a good choice as an extinguisher for an electrical fire on a computer or other electronic device, such as television.

Answer option C is incorrect. Class D extinguishers are needed for combustible metals' fire, such as magnesium, titanium, potassium, and sodium.

QUESTION 26
How much data can you fit on a single-layer Blu-ray Disc?

A. 25GB
B. 5GB
C. 15GB
D. 10GB

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
A single-layer disc of Blu-ray can hold 25GB of data. Blu-ray Disc format ensures that it is easily extendable, as it also includes support for multi-layer discs, which should allow the storage capacity to be increased to 100GB-200GB (25GB per layer) in the future simply by adding more layers to the discs.

QUESTION 27
Which of the following backup implementation methods uses a server with large disk capacity for backing up data on the other servers and workstations in the network?

A. Server-based
B. Server-free
C. Online backup
D. Serverless

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Server-based backup uses a server with large disk capacity for backing up data on other servers and

workstations in the network. The data on this server is then copied to portable disks or tape for storage offsite. This solution helps in permitting for server-to-server replication. Using a second server placed at an offsite location and specialized software for replicating the data from one server to another over the Internet permits a user to copy the data from the backup server to the offsite server.

The benefit of using server-based backup accrues to those organizations that have huge amount of data. The speed of the backup increases because server-based disk drives is much faster than tape or portable disks. Also, the disk sizes on the server are much larger than what is currently present on portable drives. Server-based backup also permits for the copy process to take place during the day when other servers are being used for productive work.

Answer option D is incorrect. Serverless backup is a process of offloading backup procedures from a server.

Answer option C is incorrect. Online backup uses the Internet and a secure (or encrypted) connection for backing up data on the network to a third-party provider's server storage system.

Answer option B is incorrect. In server-free backup, the backup task is achieved without the use of the server.

Reference: http://www.cpapracticeadvisor.com/article/10274074/new-trends-in-backup-is-your-disaster-recovery-plan-keeping-up?page=2

QUESTION 28
Which of the following is a unique identifier that is assigned to each SCSI device on the bus?

A. WWPN
B. SCSI ID
C. Domain ID
D. WWNN

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
A SCSI ID is a unique identifier that is assigned to each SCSI device on the bus. It helps in determining the SCSI device's priority. There are 16 IDs, ranging from 0 to 15.

Answer option C is incorrect. A domain id is a number that uniquely identifies a switch in a fabric and is the highest or most significant hierarchical level in the three-level address hierarchy.

In general, each switch is a single domain. The domain ID is an 8-bit identifier with a range of 0-255. Typically, zero (0) is reserved and one (1) is the default setting for new switches.

Answer option D is incorrect. Worldwide node name is a Node_Name that is worldwide unique. A World Wide Node Name, WWNN, or WWnN, is a World Wide Name assigned to a node (an endpoint, a device) in a Fibre Channel fabric. It is valid for the same WWNN to be seen on many different ports (different addresses) on the network, identifying the ports as multiple network interfaces of a single network node.

Answer option A is incorrect. In SAN, a World Wide Port Name, WWPN, or WWpN, is a World Wide Name assigned to a port in a Fibre Channel fabric. It performs a function equivalent to the MAC address in the Ethernet protocol, as it is supposed to be a unique identifier in the network.

QUESTION 29
Which of the following uses an RJ-45 8-pin male connector?

A. 10BaseT Ethernet cable
B. VGA
C. Serial Port

D. PSTN

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
A 10BaseT Ethernet cable uses an RJ-45 8-pin male connector to connect to the network adapter jack on a computer. An RJ-45 connector is similar to an RJ-11 telephone connector, but larger in size. An RJ-45 connector has eight conductors, while RJ-11 has only four conductors.

Answer option D is incorrect. Public Switched Telephone Network (PSTN) cable uses a RJ-11 4-pin male connector.

Answer option C is incorrect. A serial port uses a DB-9 (9-pin) male connector.

Answer option B is incorrect. A VGA/SVGA video adapter uses a DB-15 (15-pin) female connector.

QUESTION 30
Which of the following RAID levels can provide protection against two concurrent disk failures?

A. RAID 6
B. RAID 3
C. RAID 4
D. RAID 5

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
RAID 6 is also known as "Block striping with two sets of distributed parity", as the data is striped at the block level, and the parity information is distributed across all disks. RAID 6 is similar to RAID 5, but it provides protection against a double disk failure as it uses two parity blocks instead of one.

QUESTION 31
Which virtualization technique requires the addition of an agent to each host for the purpose of managing storage allocation?

A. In-band appliance
B. Out-of-band appliance
C. Host-based LDM
D. Host-based LVM

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Answer options D and C are incorrect. Host-based virtualization requires additional software running on the host, as a privileged task or process. Most modern operating systems have some form of logical volume manager built-in, such as, LVM (Logical Volume Manger) in UNIX/Linux and Logical Disk Manager (LDM) in Windows.

Answer option A is incorrect. In-band appliances (symmetric virtualization devices) actually sit in the data path between the host and storage. All I/O requests and their data pass through the device. Hosts perform I/O to the virtualization device and never interact with the actual storage device. The virtualization device in turn performs

I/O to the storage device. Caching of data, statistics about data usage, replications services, data migration and thin provisioning are all easily implemented in an in-band device.

QUESTION 32
Which of the following is a network file system access protocol primarily used by Windows clients for communicating file access requests to Windows servers?

A. CIFS
B. NDMP
C. iSCSI
D. FCIP

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
The Common Internet File System (CIFS) also known as Server Message Block (SMB) is a network protocol used in sharing files on a LAN. The protocol allows a client to access servers, files, and printers on the LAN. CIFS has a client and a server part. The server is probably more applicable as it could expose the repository as a CIFS server. Repository would show up in the Network Neighbourhood and users would be able to map the repository as a normal Windows drive letter. CIFS is a network file system access protocol primarily used by Windows clients for communicating file access requests to Windows servers.

Answer option B is incorrect. Network Data Management Protocol (NDMP) is an open protocol used to control data backup and recovery communications between primary and secondary storage in a heterogeneous network environment.

NDMP specifies a common architecture for the backup of network file servers and enables the creation of a common agent that a centralized program can use to back up data on file servers running on different platforms. By separating the data path from the control path, NDMP minimizes demands on network resources and enables localized backups and disaster recovery. With NDMP, heterogeneous network file servers can communicate directly to a network-attached tape device for backup or recovery operations.

Answer option C is incorrect. iSCSI (Internet Small Computer System Interface) is an IP-based storage networking standard for linking data storage facilities. iSCSI easily transfers data over intranets and manages storage over long distances by using SCSI commands over IP networks. iSCSI can transmit data over local area networks, wide area networks, or the Internet and can enable location-independent data storage and retrieval. This protocol allows initiators to send SCSI commands to SCSI storage devices on remote servers. It is a popular storage area network protocol that allows organizations to consolidate storage into data center storage arrays while providing hosts with the illusion of locally-attached disks. Unlike traditional Fibre Channel, which requires special-purpose cabling, iSCSI can be run over long distances using existing network infrastructure.

Answer option D is incorrect. Fibre Channel over IP (FCIP) is an IP-based storage networking technology, developed by the IETF and defined in RFC 3821. FCIP is also known as Fibre Channel tunneling or storage tunneling. FCIP enables the transmission of Fibre Channel information by tunneling data between storage area network facilities over IP networks. This capacity facilitates data sharing over a geographically distributed enterprise. FCIP is a key technology that is expected to help bring about rapid development of the storage area network market by increasing the capabilities and performance of storage data transmission.

Reference: http://www.snia.org/education/dictionary/c#common_internet_file_system

QUESTION 33
Which of the following electrical events shows that the demand of the electrical power exceeds the capability of the electrical power supply system and reduces the voltage for everyone?

A. Blackout

B. Power surge
C. Power spike
D. Brownout

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Brownout shows that the demand of the electrical power exceeds the capability of the electrical power supply system and reduces the voltage for everyone. It indicates there is enough power on the grid to prevent blackout or a total power loss but there is not enough power to meet the current electrical demand. It frequently occurs during informal weather conditions such as suddenly cold or hot spell. Brownout is quite harder for computer equipment than blackout. A brownout lasts longer than power sag and corrupts more data.

Answer option A is incorrect. A blackout indicates a complete loss of PC's electrical source. It is an event that shows a sudden drop-off power source, which can cause a wide variety of problems on a PC or a network. A blackout is not a power failure over an entire area but it can be in a section or a part of a building, city, or any other larger area. It is caused by electrical storms, traffic accidents in utility poles, or a total collapse of the power system due to demand overload.

Answer option B is incorrect. Power surge is a sharp increase in the voltage or an over voltage event. It is a short and temporary increase in voltage on the power grid and it is like a rough wave. Different types of electrical disturbance such as lightning storm, distant lightning strikes, or problems on the electrical power supply grid can cause the voltage to suddenly increase.

Answer option C is incorrect. A power spike is a sudden isolated extremely high over voltage event on an electrical line. The primary cause of the power spike is lightning strikes. Lightning carries millions of volts, and if a home or office takes a direct hit, a PC along with other devices are likely to be heavily damaged. Direct striking is a rare event but a strike within a mile can create a sudden spike in the electrical current near the strike.

QUESTION 34
Which of the following is a storage technology that allows storage devices to connect to servers directly?

A. SAN
B. SDN
C. DAS
D. NAS

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Direct-Attached Storage (DAS) is a storage technology that allows storage to connect to servers directly. Block-level access protocols are used by applications to access data from DAS, such as internal HDD of a host, tape libraries, directly connected external HDD packs, etc.

Answer option D is incorrect. Network-attached storage (NAS) is file-level computer data storage connected to a computer network providing data access to heterogeneous network clients. NAS systems contain one or more hard disks, often arranged into logical, redundant storage containers or RAID arrays. It removes the responsibility of file serving from other servers on the network. NAS uses file-based protocols, such as NFS, SMB/CIFS, or AFP. NAS units rarely limit clients to a single protocol.

Answer option A is incorrect. A storage area network (SAN) is a high-speed special-purpose network (or subnetwork) that interconnects different kinds of data storage devices with associated data servers on behalf of a larger network of users. The SAN architecture physically and logically separates storage from servers and

uses block I/O architecture.

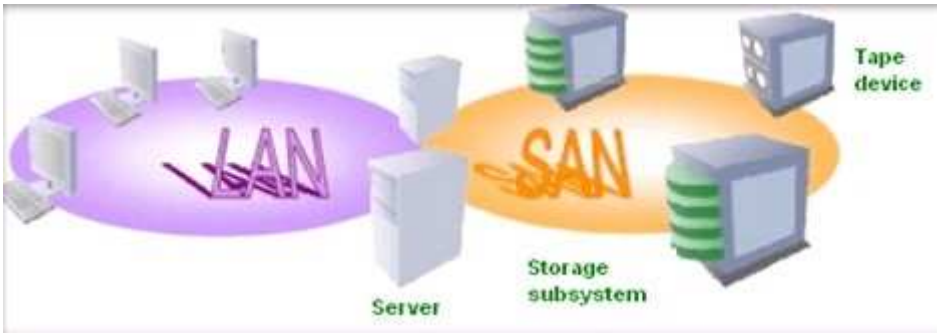

Typically, a storage area network is part of the overall network of computing resources for an enterprise. SANs support disk mirroring, backup and restore, archival and retrieval of archived data, data migration from one storage device to another, and the sharing of data among different servers in a network. SANs can incorporate subnetworks with network-attached storage (NAS) systems. Following are the features of SAN:

- Fibre Channel: SANs often utilize a Fibre Channel fabric topology - an infrastructure specially designed to handle storage communications. It provides faster and more reliable access than higher-level protocols used in NAS. A fabric is similar in concept to a network segment in a local area network. A typical Fibre Channel SAN fabric is made up of a number of Fibre Channel switches.
- Block mode: The SAN architecture physically and logically separates storage from servers and uses Block I/O architecture. A SAN is used to address data by disk block number and transfers raw disk blocks.
- File system on host: Multiple hosts can be configured for accessing the same storage devices, permitting servers and storage to scale independently. File System is managed by servers.
- FC or iSCSI protocol: SAN uses protocols such as FCIP, FCP, iFCP, iSCSI, iSNS, NDMP, SAS, and SCSI.

Answer option B is incorrect. This is an invalid option.

QUESTION 35
Which of the following components enables any-server-to-any-storage device connectivity through the use of Fibre Channel switching technology?

A. SCSI host adapter
B. Parallel ATA
C. Fibre Channel switch
D. SAN fabric

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
The SAN fabric is an area of the network that contains routers and switches. It enables any-server-to-any-storage device connectivity through the use of Fibre Channel switching technology.

Answer option C is incorrect.A Fibre Channel switch is a network switch compatible with the Fibre Channel (FC) protocol. It allows the creation of a Fibre Channel fabric that is currently the core component of most storage area networks. Fibre Channel switches implement zoning, a mechanism that disables unwanted traffic between certain fabric nodes.

Fibre Channel switches may be deployed one at a time or in larger multi-switch configurations. SAN administrators typically add new switches as their server and storage needs grow, connecting switches together via fiber optic cable using the standard device ports.

Answer option A is incorrect. A SCSI host adapter is a device used to connect one or more other SCSI devices to a computer bus.

Answer option B is incorrect. Parallel ATA (PATA) is an interface standard for the connection of storage devices, such as hard disks, solid-state drives, floppy drives, and CD-ROM drives in computers.

QUESTION 36
Which of the following methods backs up all changes made since the last full or normal backup?

A. Full backup
B. Half backup
C. Differential backup
D. Incremental backup

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
A differential backup backs up files that are created or changed since the last full backup. It requires the minimum space to back up data. A differential backup requires only the last full backup tape and the last differential backup tape to restore data. It is faster as compared to the full backup.

Differential backup contains all files that were changed since the last full backup. The advantage of a differential backup over an incremental backup is that it shortens the restore time.

Answer option A is incorrect. A full backup backs up the entire database including the transaction log. Taking a full backup daily is impractical, as it is time consuming. Instead, a well-defined backup strategy should be implemented as a weekly full backup and a daily differential backup.

Answer option D is incorrect. An incremental backup backs up files that are created or changed since the last full or incremental backup. An incremental backup provides a faster method of backing up data than most other backup methods. Restoring data from an incremental backup requires the last full backup and all subsequent incremental backups. Incremental backups must be restored in the same order as they were created.

If any incremental backup in the incremental backup set is damaged or becomes corrupt, the data backed up after corruption cannot be restored.

Answer option B is incorrect. There is no such backup method as half backup.

Reference: TechNet, Contents: "Chapter 3 - Backing Up and Restoring Databases"

QUESTION 37
Which of the following protocols is used for backups on an Ethernet-based network?

A. FDDI
B. iSNS
C. FCIP
D. NDMP

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Network Data Management Protocol (NDMP) is an open protocol used to control data backup and recovery communications between primary and secondary storage on an Ethernet-based network. With NDMP, heterogeneous network file servers can communicate directly to a network-attached tape device for backup or recovery operations.

Answer option C is incorrect. Fibre Channel over IP (FCIP) is an IP-based storage networking technology,

developed by the IETF and defined in RFC 3821. FCIP is also known as Fibre Channel tunneling or storage tunneling. FCIP enables the transmission of Fibre Channel information by tunneling data between storage area network facilities over IP networks. This capacity facilitates data sharing over a geographically distributed enterprise. FCIP is a key technology that is expected to help bring about rapid development of the storage area network market by increasing the capabilities and performance of storage data transmission.

Answer option A is incorrect. FDDI (Fiber Distributed Data Interface) is a standard for data transmission on fiber optic lines in a local area network (LAN).

Answer option B is incorrect. Internet Storage Name Service (iSNS) is a protocol that allows automated discovery, management, and configuration of iSCSI and Fibre Channel devices (using iFCP gateways) on a TCP/IP network.

Reference: http://www.ndmp.org/

QUESTION 38
Which of the following devices physically connects nodes on a multi-point bus or loop?

A. Hub
B. HBA
C. Gateway
D. Tape drive

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
A hub is a communication infrastructure device that connects nodes on a multi-point bus or loop. Hubs are used in networks for connecting devices to a bus structure. Hubs cannot aggregate bandwidth as switches can, but maintain the logical loop topology of the network by creating a "hub and spoke" physical star layout.

Reference: http://www.snia.org/education/dictionary/h

QUESTION 39
Which of the following is the duration of time and a service level within which a business process must be restored after a disaster in order to avoid unacceptable consequences associated with a break in business continuity?

A. RTA
B. RTO
C. RPO
D. RCO

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Recovery time objective (RTO) is defined as the maximum acceptable time period needed to bring one or more applications and associated data back from an outage to a correct operational state. The recovery time objective is the duration of time and a service level within which a business process must be restored after a disaster in order to avoid unacceptable consequences associated with a break in business continuity. It includes the time for trying to fix the problem without a recovery, the recovery itself, tests, and the communication to the users. Decision time for users' representative is not included. The business continuity timeline usually runs parallel with an incident management timeline and may start at the same or different points.

Answer option A is incorrect. The Recovery Time Actual (RTA) is established during an exercise, actual event, or predetermined based on the recovery methodology the technology support team develops. This is the time frame the technology support takes to deliver the recovered infrastructure to the business.

Answer option D is incorrect. The Recovery Consistency Objective (RCO) is used in Business Continuity Planning in addition to Recovery Point Objective (RPO) and Recovery Time Objective (RTO). It applies data consistency objectives to continuous data protection services.

Answer option C is incorrect. The Recovery Point Objective (RPO) describes the acceptable amount of data loss measured in time. It is the point in time to which data must be recovered as defined by the organization. The RPO is generally a definition of what an organization determines is an "acceptable loss" in a disaster situation. If the RPO of a company is 2 hours and the time it takes to get the data back into production is 5 hours, the RPO is still 2 hours. Based on this RPO, the data must be restored to within 2 hours of the disaster.

QUESTION 40
Which of the following statements are true about Network-Attached Storage (NAS)?

Each correct answer represents a complete solution. Choose three.

A. It provides storage over LAN.
B. It uses a host-independent file system.
C. It allows multiple operating systems to access the same storage resources.
D. It uses the block I/O architecture.

Correct Answer: ABC
Section: (none)
Explanation

Explanation/Reference:
Network-attached storage (NAS) is file-level computer data storage connected to a computer network providing data access to heterogeneous network clients. NAS provides access to storage over LAN and uses file I/O architecture. NAS devices include an optimized OS that implements a host-independent file system.

Network-Attached Storage (NAS) provides access to storage over the LAN. It uses a file server that is optimized serving storage. It implements a host-independent file system. NAS allows multiple operating systems to access the same storage resources. NAS can also work as a gateway to storage on the SAN

NAS systems contain one or more hard disks, often arranged into logical, redundant storage containers or RAID arrays. It removes the responsibility of file serving from other servers on the network. Following are the features of NAS:

- **TCP/IP based**: NAS uses TCP/IP Networks: Ethernet, FDDI, ATM. It connects the storage to the TCP/IP network and helps in enabling users to access storage resources using both CIFS (Microsoft based Common Internet File System) and NFS (Unix based Network File System) protocols over the familiar TCP/IP networks. The NAS appliance has its own IP address and hence can be accessed by users directly over the network for storing files as well as access them when needed. NAS generally uses multiple disks for

storing data and providing file level storage and access to data.
- **UNC addressable storage**: NAS uses UNC addressable storage.
- **File system on storage (NFS or CIFS)**: NAS uses file-based protocols, such as NFS, SMB/CIFS, or AFP. NAS units rarely limit clients to a single protocol. A NAS is used to identify data by file name and byte offsets, transfers file data or file meta-data, and handles security, user authentication, and file locking.
- **Ethernet based**: NAS systems helps in storing data as files (predominantly) and is used to support both CIFS and NFS protocols. They can be accessed easily over the commonly used TCP/IP Ethernet based networks and support number of users connecting to it simultaneously.

Answer option D is incorrect. NAS uses file I/O architecture.

QUESTION 41
Which concept includes considerations, such as data placement, deletion, repurposing, and regulatory archiving?

A. Tiers of management
B. Information lifecycle management
C. Common data usage model
D. Backup and recovery

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Information Life Management (ILM) refers to the creation and management of a storage infrastructure. It includes processes and technologies required to manage data for its lifetime. ILM determines the priority and value of the data, including considerations such as data placement, deletion, repurposing, and regulatory archiving.

A successful information lifecycle management strategy must be business-centric by tying closely with key processes, applications, and initiatives of the business. It should be centrally managed, providing an integrated view into all information assets of the business, both structured and unstructured.

Reference: http://searchstorage.techtarget.com/definition/information-life-cycle-management

QUESTION 42
With ___, a computer hard disk is backed up to another hard disk rather than to a tape or floppy disk.

Correct Answer: D2D
Section: (none)
Explanation

Explanation/Reference:
The term "disk-to-disk", or "D2D", generally refers to disk-to-disk backup. With D2D, a computer hard disk is backed up to another hard disk rather than to a tape or floppy disk. D2D is often confused with virtual tape, but it differs in that it enables multiple backup and recovery operations to simultaneously access the disk directly by using a true file system.

Advantages of disk-to-disk
- Higher speed and higher capacity, relative to tape or floppy, resulting in shorter backup and recovery windows.
- Non-linear recovery of data, enabling a specific file to be restored quicker and simpler than with tape.
- Lower total cost of ownership due to increased automation and lower hardware costs.

Reference: http://en.wikipedia.org/wiki/Disk-to-disk

QUESTION 43

Which of the following is the process of copying data from one system to another across a network?

A. Hierarchical storage management
B. Queue depth
C. Replication
D. Automated tiered storage

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Replication is the process of copying data from one system to another across a network. Replication is used for data protection in order to make sure that if something happens to the primary site, the data is still safe and available for almost immediate use at the secondary site.

Backups also serve the same function. The difference between backup and replication is that with replication, data is copied in an ongoing basis, whereas backup data is copied to tape only occasionally, and in many enterprises, sent off site even less often than that.

Answer option B is incorrect. Queue depth is something that most storage administrators need to overlook as this is the reason for most performance issues. Queue depth is defined as the number of commands that the hba can send / receive in a single chunk - per LUN. From a host-hba point (initiator), it is defined as the number of commands that can be queued (or stored) and then sent to storage. From a storage point (target), it is defined as the number of commands it can accept in one-shot, again, per lun.

Performance degradation starts quickly once user reaches max queue depth on target and it starts backing off the queues and it will hit hard on user's response times. Ideally, if users have few hosts and luns, the default is good. If users have the luxury to dedicate ports for high performance applications, ensure that the queue depth is configured appropriately so that users get the best of performance.

Answer option A is incorrect. Hierarchical storage management (HSM) is a data storage technique which automatically moves data between high-cost and low-cost storage media. HSM systems exist because high-speed storage devices, such as hard disk drive arrays, are more expensive (per byte stored) than slower devices, such as optical discs and magnetic tape drives. Instead, HSM systems store the bulk of the enterprise's data on slower devices, and then copy data to faster disk drives when needed. In effect, HSM turns the fast disk drives into caches for the slower mass storage devices. The HSM system monitors the way data is used and makes best guesses as to which data can safely be moved to slower devices and which data should stay on the fast devices.

Answer option D is incorrect. Automated tiered storage is the automated progression or demotion of data across different tiers (types) of storage devices and media. This movement of data is automatic to the different types of disk according to performance and capacity requirements.

QUESTION 44
How does an IP Storage adapter differ from a traditional Network Interface Card (NIC)?

A. IP storage adapters have larger frame size.
B. IP storage adapters have data rate of 2Gb/s.
C. IP storage adapters have link aggregation.
D. IP storage adapters include TCP/IP segmentation offload.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
IP storage adapters include protocol processing, such as TCP/IP offload engines, to reduce processing loads

on the host devices. Traditional NICs are designed to transfer packets among PCs, servers, and storage devices, such as NAS appliances. A TCP/IP offload engine removes the TCP/IP processing from the host CPU and completes TCP/IP processing and packet creation on the HBA.

Reference: http://www.soltechnology.com/articles/iSCSI-for-Storage-Networking.htm

QUESTION 45
Which of following terms is defined as the practice of connecting multiple devices to the same switch port in order to optimize switch use in a SAN (storage area network) switching environment?

A. NPIV
B. Link
C. Oversubscription
D. WWPN

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
In a SAN (storage area network) switching environment, oversubscription is the practice of connecting multiple devices to the same switch port in order to optimize switch use. Each SAN port can support a particular communication speed and a Fibre Channel switch may offer 1 Gb, 2 Gb, or 4 Gb FC ports. However, because ports are rarely run at their maximum speed for a prolonged period, multiple slower devices may fan in to a single port to take advantage of unused capacity. For instance, a single storage server may not be able to sustain 4 Gbps to a switch port, so two 2 Gb servers or four 1 Gb servers may all be aggregated to that 4 Gb switch port.

Answer option D is incorrect. In SAN, a World Wide Port Name, WWPN, or WWpN, is a World Wide Name assigned to a port in a Fibre Channel fabric. It performs a function equivalent to the MAC address in the Ethernet protocol, as it is supposed to be a unique identifier in the network.

Answer option A is incorrect. NPIV is defined as the ability for a single physical Fibre Channel node or switch for supporting more than one Nx_Port on a single point-to-point link.

N_Port ID Virtualization, or NPIV, is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port. It allows multiple Fibre Channel initiators to occupy a single physical port, easing hardware requirements in Storage Area Network design, especially where virtual SANs are called for.

Answer option B is incorrect. A link is defined as a physical connection (electrical or optical) between two nodes of a network. It can also be defined as the point-to-point physical connection from one element of a Fibre Channel fabric to the next.

Reference: http://searchstorage.techtarget.com/definition/oversubscription

QUESTION 46
Which of the following is the process or procedure to ensure that an organization follows relevant laws, regulations, and business rules?

A. Rotational scheme
B. Legal consistency
C. Legal completeness
D. Legal compliance

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Legal compliance is the process or procedure to ensure that an organization follows relevant laws, regulations, and business rules. The definition of legal compliance, especially in the context of corporate legal departments, has recently been expanded to include understanding and adhering to ethical codes within entire professions, as well.

Answer option A is incorrect. A rotation scheme is a method for effectively backing up data where multiple media (such as tapes) are used in the backup process.

Answer option B is incorrect. Legal consistency is a property that declares enterprise policies to be free of contradictions with the law.

Answer option C is incorrect. Legal completeness is a property that declares enterprise policies to cover all scenarios included or suggested by the law.

Reference: http://en.wikipedia.org/wiki/Legal_governance,_risk_management,_and_compliance

QUESTION 47
Which of the following converged storage network technologies refers to a set of enhancements to Ethernet local area networks for use in data center environments?

A. FCoE
B. DCB
C. Classes of service
D. LLDP

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Data center bridging (DCB) refers to a set of enhancements to Ethernet local area networks for use in data center environments. Specifically, DCB goals are, for selected traffic, to eliminate loss due to queue overflow and to be able to allocate bandwidth on links. Essentially, DCB enables, to some extent, the treatment of different priorities as if they were different pipes. The primary motivation was the sensitivity of Fibre Channel over Ethernet to frame loss. The higher level goal is to use a single set of Ethernet physical devices or adapters for computers to talk to a Storage Area Network, Local Area network and InfiniBand fabric.

Answer option A is incorrect. Fibre Channel over Ethernet (FCoE) is an encapsulation of Fibre Channel frames over Ethernet networks. It allows Fibre Channel to use 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol. FCoE maps Fibre Channel over selected full duplex IEEE 802.3 networks for providing I/O consolidation over Ethernet and reducing network complexity in the datacenter. The FCoE protocol specification replaces the FC0 and FC1 layers of the Fibre Channel stack with Ethernet.

Answer option D is incorrect. The Link Layer Discovery Protocol (LLDP) is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on an IEEE 802 local area network. LLDP performs functions similar to several proprietary protocols, such as Cisco Discovery Protocol, Extreme Discovery Protocol, Nortel Discovery Protocol (also known as SONMP), and Microsoft's Link Layer Topology Discovery (LLTD).

Answer option C is incorrect. Fibre Channel provides several connection strategies, known as classes of service, for connecting devices through a Fibre Channel fabric.

QUESTION 48
Which of the following statements are true regarding the IP storage networking?

Each correct answer represents a complete solution. Choose four.

A. It builds on SCSI and Ethernet technologies.

B. It takes advantage of existing Ethernet/IP knowledge base.
C. It supports up to 16 million devices.
D. It enables storage to be accessed over LAN, MAN, or WAN environments, without needing to alter storage applications.
E. It is easy to keep up the data if IP storage network breaks down.

Correct Answer: ABCD
Section: (none)
Explanation

Explanation/Reference:
IP storage networking is a general term for several approaches of using the Internet Protocol (IP) for storage. With the advent of new IP storage products and transport protocol standards, such as, iSCSI, FCIP, and iFCP, users now have more choices for accessing files over IP-based LANs. IP storage networking provides the following advantages:

- It leverages the large installed base of Ethernet-TCP/IP networks.
- It enables storage to be accessed over LAN, MAN, or WAN environments, without needing to alter storage applications.
- It takes advantage of existing Ethernet/IP knowledge base and management tools.
- It builds on SCSI and Ethernet technologies.
- It is an alternative to the Fibre Channel framework of the traditional SAN.
- It supports up to 16 million devices.

Answer option E is incorrect. It is difficult to keep up the data if IP storage network breaks down.

Reference: http://searchstorage.techtarget.com/definition/IP-storage

QUESTION 49
In a storage area network, which of the following is a logical grouping of ports for forming a virtual private storage network?

A. Domain ID
B. NPIV
C. Zone
D. SCSI ID

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
In a storage area network (SAN), zoning is the allocation of resources for device load balancing and for selectively allowing access to data only to certain users.

Zones, zone aliases, and zone sets permit logical grouping of ports and storage devices within a storage area network. The following section describes zoning concepts and elements.

In a storage area network, a zone is a logical grouping of ports for forming a virtual private storage network.

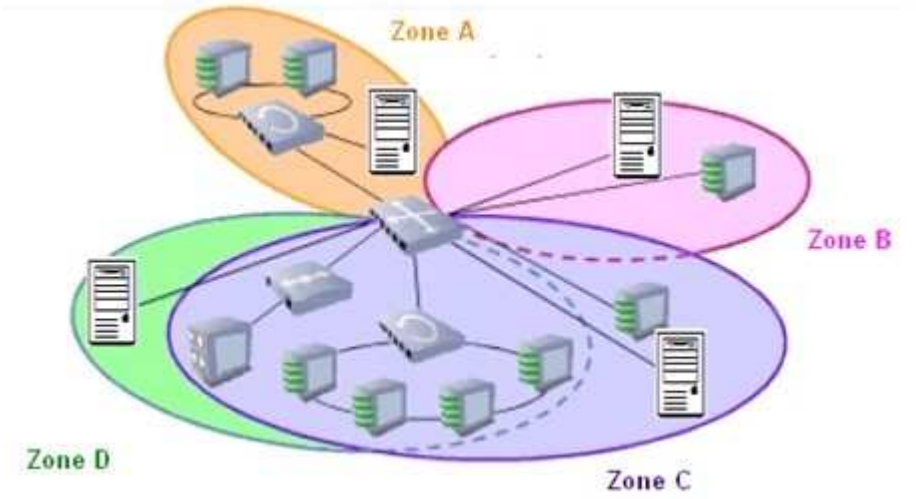

A zone that belongs to a single SAN can be grouped into a zone set that can be activated or deactivated as a single entity across all switches in the fabric.

A zone set includes one or more zones, and a zone can be a member of more than one zone set.

The ports and devices that are in a zone are called zone members. A zone includes one or more zone members.

A zone alias is a collection of zone members. A zone alias can be added to one or more zones.

Answer option B is incorrect. N_Port ID Virtualization, or NPIV, is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port. It allows multiple Fibre Channel initiators to occupy a single physical port, easing hardware requirements in Storage Area Network design, especially where virtual SANs are called for.

Answer option A is incorrect. A domain id is a number that uniquely identifies a switch in a fabric and is the highest or most significant hierarchical level in the three-level address hierarchy.

In general, each switch is a single domain. The domain ID is an 8-bit identifier with a range of 0-255. Typically, zero (0) is reserved and one (1) is the default setting for new switches.

Answer option D is incorrect. A SCSI ID is a unique identifier that is assigned to each SCSI device on the bus. It helps in determining the SCSI device's priority. There are 16 IDs, ranging from 0 to 15.

QUESTION 50
Which of the following types of errors is indicated by a parity error?

A. A vulnerability issue in the network
B. A problem with file sharing in the network
C. An error in the file system
D. A problem with data stored in the memory

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Parity error indicates a problem with data stored in the memory. Parity is a very basic check of information integrity. Each byte (eight bits) of RAM storage actually takes nine bits of information. Eight bits are used for data and the last bit, known as parity bit, is used to store the parity of data. The processor is in charge of

checking the accuracy of the parity bit. A parity error can be caused by different circumstances, but it is almost always a hardware problem.

One cause of parity error is a faulty memory. The most foolproof way to resolve this problem is to swap out each piece of memory, until the problem disappears.

QUESTION 51
Which of the following devices is used to connect several LANs of an organization to a WAN?

A. Hub
B. Router
C. Bridge
D. Repeater

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
A router is a device that routes data packets between computers in different networks. It is used to connect multiple networks, and it determines the path to be taken by each data packet to its destination computer. A router can connect dissimilar networks, such as Ethernet, FDDI, and Token Ring, and route data packets among them.

Answer options A, D, and C are incorrect. These devices cannot be used to connect several LANs to a WAN. They are used only within a LAN.

QUESTION 52
Which of the following are removable storage devices?

Each correct answer represents a complete solution. Choose three.

A. Floppy disk
B. USB flash drive
C. Hard disk
D. ZIP disk

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
Floppy disks, Zip disks, USB flash drives, memory cards, CDs, and DVDs are removable storage devices. Disks and cards that can be inserted into a computer's drives or USB ports are called removable storage devices.

Answer option C is incorrect. A hard disk drive (HDD) is a non-volatile storage device that stores digitally encoded data on rapidly rotating rigid (i.e., hard) platters with magnetic surfaces. Early HDDs had removable media; however, an HDD today is typically a sealed unit (except for a filtered vent hole to equalize air pressure) with fixed media.

Reference: http://en.wikipedia.org/wiki/Hard_disk_drive

QUESTION 53
Which of the following versions of the SNMP protocol provides the best security feature to a storage network?

A. SNMPv2

B. SNMPv1
C. SNMPv3
D. SNMPv2c

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
SNMPv3 is the current standard version of SNMP. SNMPv3 primarily added security and remote configuration enhancements to SNMP. The IETF has designated SNMPv3 as a full Internet Standard, the highest maturity level for an RFC. SNMPv3 provides important security features:

- Message integrity to ensure that a packet has not been tampered with in transit.
- Authentication to verify that the message is from a valid source.
- Encryption of packets to prevent snooping by an unauthorized source.

SNMPv3 has three security levels. The highest level is with authentication and privacy. The middle level is with authentication and no privacy and the bottom level is without authentication or privacy.

Answer option B is incorrect. SNMPv1 provides the least security; it was the first version of the SNMP protocol. SNMPv1 did not have a very good implementation of security.

Answer option A is incorrect. SNMPv2 brought far better security than SNMPv1, but it does not provide the best security feature to the storage network.

Answer option D is incorrect. SNMPv2c comprises SNMPv2 without the controversial new SNMPv2 security model. It uses the simple community-based security scheme of SNMPv1.

QUESTION 54
Which of the following virtualization techniques is known as symmetric virtualization devices?

A. In-band appliances
B. Out-of-band appliances
C. Host-based LVM
D. Host-based LDM

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
In-band appliances (symmetric virtualization devices) actually sit in the data path between the host and storage. All I/O requests and their data pass through the device. Hosts perform I/O to the virtualization device and never interact with the actual storage device. The virtualization device in turn performs I/O to the storage device. Caching of data, statistics about data usage, replications services, data migration and thin provisioning are all easily implemented in an in-band device.

Answer options C and D are incorrect. Host-based virtualization requires additional software running on the host, as a privileged task or process. Most modern operating systems have some form of logical volume manager built-in, such as, LVM (Logical Volume Manger) in UNIX/Linux and Logical Disk Manager (LDM) in Windows.

Answer option B is incorrect.

QUESTION 55
_____ storage is the assignment of different categories of data to different types of storage media for reducing total storage cost.

Correct Answer: Tiered
Section: (none)
Explanation

Explanation/Reference:
Tiered storage is the assignment of different categories of data to different types of storage media for reducing total storage cost. Categories may be based on levels of protection required, performance requirements, frequency of use, and other considerations. For examples:

- Disk and Tape: Two separate storage tiers identified by differences in all four defining attributes.
- Old technology disk and new technology disk: Two separate storage tiers identified by differences in one or more of the attributes.
- High performing disk storage and less expensive, slower disk of the same capacity and function: Two separate tiers.
- Identical Enterprise class disk configured to utilize different functions such as RAID level or replication: A separate storage tier for each set of unique functions.

QUESTION 56
Which of the following RAID levels is also known as "Striped mirroring array"?

A. RAID 0
B. RAID 1+0
C. RAID 0+1
D. RAID 1

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
RAID 0+1 is also known as "Striped mirroring array", as it strips data across multiple drives, and then mirrors the array to another array.

Answer option B is incorrect. RAID 1+0 uses pairs of devices to create two copies of each block of data.

Answer option A is incorrect. RAID 0 is known as "Block striping without parity".

Answer option D is incorrect. RAID 1 is called "Mirroring of disks".

QUESTION 57
Which of the following is the automated migration of data objects among storage devices, usually based on inactivity?

A. Hierarchical storage management
B. RAID
C. Cache
D. Queue depth

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Hierarchical storage management is defined as the automated migration of data objects among storage devices, usually based on inactivity.

Hierarchical storage management (HSM) is a data storage technique which automatically moves data between high-cost and low-cost storage media. HSM systems exist because high-speed storage devices, such as hard

disk drive arrays, are more expensive (per byte stored) than slower devices, such as optical discs and magnetic tape drives. Instead, HSM systems store the bulk of the enterprise's data on slower devices, and then copy data to faster disk drives when needed. In effect, HSM turns the fast disk drives into caches for the slower mass storage devices. The HSM system monitors the way data is used and makes best guesses as to which data can safely be moved to slower devices and which data should stay on the fast devices.

Answer option D is incorrect. Queue depth is defined as the number of commands that the HBA can send or receive in a single chunk - per LUN.

Answer option C is incorrect. Cache is a component that transparently stores data so that the future requests for that data can be served faster.

Answer option B is incorrect. RAID is described as a redundant array of inexpensive disks. It is a technology that allows computer users to achieve high levels of storage reliability from low-cost and less reliable PC-class disk-drive components, via the technique of arranging the devices into arrays for redundancy.

Reference: http://www.snia.org/education/dictionary/h

QUESTION 58
Which of the following protocols removes the need for transporting the data through the backup server itself, thus enhancing speed and removing load from the backup server?

A. NFS
B. FCoE
C. NDMP
D. CIFS

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Network Data Management Protocol (NDMP) is a protocol meant to transport data between NAS devices and backup devices. This removes the need for transporting the data through the backup server itself, thus enhancing speed and removing load from the backup server. The NDMP protocol is governed by the Network Data Management Task Force organization.

Answer option D is incorrect. Common Internet File System (CIFS) is a protocol that permits programs to make requests for files and services on remote computers on the Internet.

Answer option A is incorrect. Network File System (NFS) is a network file system protocol that allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed.

Answer option B is incorrect. Fibre Channel over Ethernet (FCoE) is an encapsulation of Fibre Channel frames over Ethernet networks.

QUESTION 59
Which of the following is a data storage virtualization technology used to present a storage component (usually hard disk storage) as tape libraries or tape drives for use with existing backup software?

A. Virtual tape library
B. Disk-to-disk
C. Disk-to-tape
D. Disk-to-disk-to-tape

Correct Answer: A
Section: (none)

Explanation

Explanation/Reference:
A virtual tape library (VTL) is a data storage virtualization technology used typically for backup and recovery purposes. A VTL is used to present a storage component (usually hard disk storage) as tape libraries or tape drives for use with existing backup software.

Virtualizing the disk storage as tape permits integration of VTLs with existing backup software and existing backup and recovery processes and policies. The benefits of such virtualization include storage consolidation and faster data restore processes.

By backing up data to disks instead of tapes, VTL often enhances performance of both backup and recovery operations. Restore processes are found to be faster than backup regardless of implementations. In some cases, the data stored on the VTL's disk array is exported to other media, such as physical tapes, for disaster recovery purposes (scheme called disk-to-disk-to-tape, or D2D2T).

Answer option B is incorrect. The term "disk-to-disk", or "D2D", generally refers to disk-to-disk backup. With D2D, a computer hard disk is backed up to another hard disk rather than to a tape or floppy disk. D2D is often confused with virtual tape, but it differs in that it enables multiple backup and recovery operations to simultaneously access the disk directly by using a true file system.

Answer option C is incorrect. The term "disk-to-tape" refers to disk-to-tape backup. In this process, a computer hard disk is backed up to a tape.

In addition, tape is quite simply a pain to work with, mainly if a cartridge must be retrieved, loaded, and scanned in its entirety for recovering one file and tapes can be lost or stolen, too.

When there is necessity of restoring data, disk's big benefits over tape is that it is random-access rather than sequential access. That means that if a user only needs one file or a few files back, it will be faster and easier to search and recover from disk.

One of the first backup applications for disk was to emulate a tape drive and using a VTL, there is no need to change your software or processes - they just run a lot faster.

Answer option D is incorrect. Disk-to-disk-to-tape (D2D2T) is an approach to computer storage backup and archiving. In D2D2T scheme, data is initially copied to backup storage on disk storage system and then periodically copied again to a tape storage system.

A D2D2T scheme permits the administrator to automate daily backups on disk for implementing fast restores and then move data to tape when an administrator has time. The tape is also used to transfer more mature data offsite for disaster recovery protection and to comply with regulatory policies for long-term data retention at a relatively inexpensive cost.

Reference: http://en.wikipedia.org/wiki/Virtual_tape_library

QUESTION 60
Which of the following is an alternative name for an entity, which is sometimes used for creating names that are more easily human readable?

A. Target
B. Initiator
C. Connection
D. Alias

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Alias is an alternative name for an entity, which is sometimes used for creating names that are more easily human readable.

Answer option C is incorrect. A connection is a communication path between the initiator and target using a TCP/IP connection. In iSCSI, one or more connection makes up a session. A connection helps in carrying control messages, SCSI commands, parameters, and data within iSCSI PDUs.

Answer option B is incorrect. Initiator is the system component that originates an I/O command over an I/O bus or network. It helps in building and issuing a Command Descriptor Block (CDB), forwarding the CDB to a SCSI Target device, handling any interaction associated with that particular command (such as receiving and acknowledging data for a READ command), and checking/accepting a response (Status) from the SCSI Target device.

I/O adapters, network interface cards, and intelligent controller device I/O bus control ASICs are typical initiators.

Answer option A is incorrect. A target is the system component that receives a SCSI I/O command. The iSCSI specification refers to a storage resource located on an iSCSI server as a target. An iSCSI target is often a dedicated network-connected hard disk storage device, but may also be a general-purpose computer. Since with initiators, software for providing an iSCSI target is available for most mainstream operating systems.

QUESTION 61
Which of the following technologies is used to address the legal issue of the retaining records?

A. EDI
B. Blu-ray disk
C. BIOS
D. WORM

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
WORM is used to address the legal issue of the retaining records. WORM (Write Once, Read Many) refers to computer data storage systems, data storage devices, and data storage media that can be written to once, but read from multiple times. WORM does not include the original Compact Disc or CD-ROM, because they are pressed from master disks, and not writable by a computer.

Answer option A is incorrect. Electronic data interchange (EDI) is the structured transmission of data between organizations by electronic means.

Answer option B is incorrect. Blu-ray disks are used for storing data.

Answer option C is incorrect. BIOS (basic input/output system) is the program a personal computer's microprocessor uses to get the computer system started after you turn it on.

QUESTION 62
Choose the ST connector from the images below.

A.

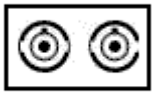

B.

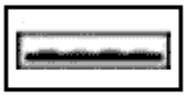

C.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
ST: A straight tip (ST) connector is a fiber-optic connector used with multimode fiber. An ST connector has a 2.5mm shaft and bayonet locking ring, and allows quick connect and disconnect of 125 micron multi-mode fiber.

ST

USB: A Universal Serial Bus (USB) connector is used with the USB cable for connecting various electronic devices to a computer. USB supports a data speed of up to 12 megabits per second. Two types of connectors are used with USB, namely USB-A Type and USB-B Type.

USB-A Type

USB-B Type

QUESTION 63
Which of the following SCSI standards has a data transfer rate of 5Mbps and a maximum cable length of 6 meters?

A. SCSI-2
B. Ultra2 SCSI
C. Fast SCSI
D. Fast Wide

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
The SCSI-2 standard has a data transfer rate of 5Mbps and a maximum cable length of 6 meters.

Answer option C is incorrect. The Fast SCSI standard has a transfer rate of 10Mbps and a maximum cable length of 3 meters.

Answer option B is incorrect. The Ultra2 SCSI standard has a data transfer rate of 40Mbps with a maximum

cable length of 12 meters.

Answer option D is incorrect. The Fast Wide SCSI standard has a transfer rate of 20Mbps and a maximum cable length of 1.5 meters.

Reference: "http://en.wikipedia.org/wiki/SCSI"

QUESTION 64
Which of the following is defined as the maximum acceptable time period needed to bring one or more applications and associated data back from an outage to a correct operational state?

A. LVM
B. RPO
C. RTO
D. MPIO

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Recovery Time Objective (RTO) is defined as the maximum acceptable time period needed to bring one or more applications and associated data back from an outage to a correct operational state.

The Recovery Time Objective is the duration of time and a service level within which a business process must be restored after a disaster in order to avoid unacceptable consequences associated with a break in business continuity. It includes the time for trying to fix the problem without a recovery, the recovery itself, tests, and the communication to the users. Decision time for users' representative is not included.

The business continuity timeline usually runs parallel with an incident management timeline and may start at the same or different points.

Answer option B is incorrect. Recovery Point Objective (RPO) is defined as the maximum acceptable time period prior to a failure or disaster during which changes to data may be lost as a result of recovery.

Data modified prior to the failure or disaster by at least this time period is preserved by recovery. Zero is a valid value and is correspondent to a "zero data loss" requirement.

Organizations determine RPO according to "acceptable loss" in a disaster situation. If the RPO of a company is 5 hours, it means that more than 5 hours of production data cannot be lost. In other words, it can be said that the data must be restored within 5 hours of the disaster.

Answer option D is incorrect. In computer storage, multipath I/O is a fault-tolerance and performance enhancement technique whereby there is more than one physical path between the CPU in a computer system and its mass storage devices through the buses, controllers, switches, and bridge devices connecting them.

Answer option A is incorrect. Logical volume management is the method used for providing a higher level-view of the disk storage on a computer than the traditional view of disks and partitions. Due to its flexibility, it is used by the system administrators for allocating storage to applications and users.

Although storage volumes need some upgrading of file system tools, they can be resized and moved around at the user's will under the control of the logical volume manager. The logical volume manager also facilitates in the management of storage volumes in user-defined groups. This helps the system administrators to deal with sensibly named volume groups rather than physical disk names.

The following are the features of the LVM (Logical Volume Manager):
- It resizes volume groups online by absorbing new physical volumes (PV) or by ejecting existing ones.
- It resizes logical volumes (LV) online by concatenating extents onto them or by truncating extents from them.

- It creates read-only snapshots of logical volumes (LVM1).
- It creates read-write snapshots of logical volumes (LVM2).
- It stripes whole or parts of logical volumes across multiple PVs, in a fashion similar to RAID 0.
- It mirrors whole or parts of logical volumes, in a fashion similar to RAID 1.
- It moves online logical volumes between PVs.
- It splits or merges volume groups in situ (as long as no logical volumes span the split). This can be useful when migrating whole logical volumes to or from offline storage.

Reference: http://www.snia.org/education/dictionary/r

QUESTION 65
Which of the following allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed?

A. MPIO
B. iSCSI
C. CIFS
D. NFS

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Network File System (NFS) is a network file system protocol that allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed. NFS, like many other protocols, builds on the Open Network Computing Remote Procedure Call (ONC RPC) system. The Network File System is an open standard defined in RFCs, allowing anyone to implement the protocol.

Answer option C is incorrect. Common Internet File System (CIFS) is a protocol that permits programs and makes requests for files and services on remote computers on the Internet. The client/server programming model is used by CIFS. A client program makes a request of a server program (usually in another computer) for accessing a file or passing a message to a program that runs in the server computer. The server takes the requested action and returns a response.

Answer option B is incorrect. iSCSI (Internet Small Computer System Interface) is an IP-based storage networking standard for linking data storage facilities. iSCSI easily transfers data over intranets and manages storage over long distances by using SCSI commands over IP networks. iSCSI can transmit data over local area networks, wide area networks, or the Internet and can enable location-independent data storage and retrieval. This protocol allows initiators to send SCSI commands to SCSI storage devices on remote servers. It is a popular storage area network protocol that allows organizations to consolidate storage into data center storage arrays while providing hosts with the illusion of locally-attached disks. Unlike traditional Fibre Channel, which requires special-purpose cabling, iSCSI can be run over long distances using existing network infrastructure.

Answer option A is incorrect. In computer storage, multipath I/O is a fault-tolerance and performance enhancement technique whereby there is more than one physical path between the CPU in a computer system and its mass storage devices through the buses, controllers, switches, and bridge devices connecting them.

QUESTION 66
Which of the following is a powerful troubleshooting tool used for generating a log or multiple logs that can be graphically analyzed to identify problem areas in system performance?

A. iostat
B. inputstat
C. perfmon
D. sysmon

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
PerfMon is a powerful troubleshooting tool used for generating a log or multiple logs that can be graphically analyzed to identify problem areas in system performance.

Answer option A is incorrect. Iostat (input/output statistics) is a computer system monitor tool used for collecting and showing operating system storage input and output statistics. It is often used for identifying performance issues with local disks or networked file systems such as NFS.

Answer option D is incorrect. Sysmon is a network monitoring tool that is designed to give high performance and accurate network monitoring. Currently, supported protocols contain SMTP, IMAP, HTTP, TCP, UDP, NNTP, and PING tests.

Answer option B is incorrect. This is an invalid option.

QUESTION 67
Which of the following statements are true regarding Direct-Attached Storage?

Each correct answer represents a complete solution. Choose three.

A. DAS uses the block I/O architecture.
B. Storage devices can only be external disk subsystems.
C. Storage is managed by a single host.
D. Storage devices are directly attached to a server's I/O bus.

Correct Answer: ACD
Section: (none)
Explanation

Explanation/Reference:
DAS stands for Direct-Attached Storage. It is a traditional way of implementing storage managed by a single host. Other host can access the storage only through that single host, over the local network. Storage devices are connected to a server's I/O bus. It uses the block I/O architecture.

Answer option B is incorrect. Storage devices can be internal or external disks attached to a SCSI bus usually.

QUESTION 68
Which of the following allows the use of advanced features of SATA, such as hot-plugging and native command queuing (NCQ)?

A. IDE
B. AHCI
C. ISA
D. PCI

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Advanced Host Controller Interface (AHCI) is an open host controller interface that allows the use of advanced features of SATA, such as hot-plugging and native command queuing (NCQ).

Answer option A is incorrect. IDE is a connection interface for hard disks.

Answer options D and C are incorrect. ISA and PCI are expansion slots.

Reference: http://en.wikipedia.org/wiki/Serial_ATA#Hotplug

QUESTION 69
Which of the following automates the migration of data from one storage to another based on the user-defined policy?

A. ADM
B. SRM
C. ILM
D. HSM

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Automated Data Migration (ADM) tools automate the migration of data from one storage to another based on the user-defined policy. These are combination of storage resource management (SRM) and hierarchical storage management (HSM). ADM extends HSM capabilities with policy engine.

Answer option D is incorrect. HSM tools are used to automate the immigration of data from one storage to another based on how old the data is and how frequently users access it.

Answer option C is incorrect. Information Life Management (ILM) refers to the creation and management of a storage infrastructure. It includes processes and technologies required to manage data for its lifetime. ILM determines the priority and value of the data, including considerations such as data placement, deletion, repurposing, and regulatory archiving.

A successful information lifecycle management strategy must be business-centric by tying closely with key processes, applications, and initiatives of the business. It should be centrally managed, providing an integrated view into all information assets of the business, both structured and unstructured.

Answer option B is incorrect. SRM applications help a SAN administrator track and control use of SAN resources.

QUESTION 70
Snapshot is a point in time copy of a defined collection of data. What are the disadvantages of using snapshots?

Each correct answer represents a complete solution. Choose three.

A. Snapshot does not provide protection against disk failure.
B. Snapshot relies on primary storage still being available.
C. Snapshot does not protect against corrupted data.
D. Snapshot data set is never 100% up to date.

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
Following are the disadvantages of snapshot:
- Snapshot data set is never 100% up-to-date.
- Snapshot does not provide protection against disk failure.
- Snapshot relies on primary storage still being available.

Answer option C is incorrect. It is a point-in-time copy that protects against corrupted data.

QUESTION 71
Which communication method is used in SCSI interfaces?

A. Simple
B. Duplex
C. Half-duplex
D. Full-duplex

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Small Computer System Interface (SCSI) is the second most popular drive interface in use today after the Integrated Drive Electronics (IDE) interface. The communication method used in the SCSI interface is half-duplex.

QUESTION 72
Which of the following statements are true regarding SNMP?

Each correct answer represents a complete solution. Choose three.

A. It is defined as a protocol that is used to manage the network devices.
B. It is an application-level protocol that contains three elements.
C. It is a transport-level protocol that contains two elements.
D. It is based on the client server model.

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
SNMP is an application-level protocol that is used to manage network devices. It is based on the client server model in which a manager (client) gets information from the agent (server). It contains three elements:
1. SNMP client (manager): It is defined as a device that is used to execute management applications for monitoring and controlling the network elements.
2. SNMP server (agent): It is used to collect and store management information.
3. Management Information Base (MIB): It is defined as a collection of managed objects that are contained in a virtual information store.

QUESTION 73
Which of the following Fibre Channel topologies was introduced as a replacement for parallel SCSI?

A. Redundant fabric
B. Single fabric
C. Arbitrated loop (FC_AL)
D. Point-to-point

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Arbitrated loop, also known as FC-AL, was introduced as a replacement for parallel SCSI. It is a Fibre Channel topology in which devices are connected in a one-way loop fashion in a ring topology. It is a serial architecture that is compatible with SCSI, handling up to 127 ports (devices). One port may optionally connect a loop to a fabric switch port. The bandwidth on the loop is shared among all ports. Only two ports may communicate at a time on the loop. One port wins arbitration and may open one other port in either half or full duplex mode.

QUESTION 74
What is the disadvantage of implementing iSCSI as a software iSCSI client?

A. An iSCSI HBA is needed.
B. It is an expensive solution.
C. Host CPU overhead increases.
D. FCIP is automatically enabled.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
iSCSI (Internet Small Computer System Interface) is an IP-based storage networking standard for linking data storage facilities. iSCSI easily transfers data over intranets and manages storage over long distances by using SCSI commands over IP networks. iSCSI can transmit data over local area networks, wide area networks, or the Internet and can enable location-independent data storage and retrieval. This protocol allows initiators to send SCSI commands to SCSI storage devices on remote servers. It is a popular storage area network protocol that allows organizations to consolidate storage into data center storage arrays while providing hosts with the illusion of locally-attached disks. Unlike traditional Fibre Channel, which requires special-purpose cabling, iSCSI can be run over long distances using existing network infrastructure.

iSCSI can be implemented as a software iSCSI client or a hardware-based solution using an iSCSI HBA. When you implement it as a software iSCSI client, the host CPU overhead will increase. You can run it in software (i.e., low cost solution) if you are not processor bound. Otherwise, you can put an iSCSI HBA in to implement it.

Reference: http://searchstorage.techtarget.com/news/934219/Quick-Guide-iSCSI-implementation

QUESTION 75
Which of the following ports are examples of loop-capable ports?

Each correct answer represents a complete solution. Choose two.

A. F_Port
B. G_Port
C. FL_Port
D. NL_Port

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:
A "Loop" port is capable of performing arbitrated loop functions and protocols. NL_Ports and FL_Ports are examples of loop-capable ports.

Answer option B is incorrect. A "Generic" Port can operate as either an E_Port or an F_Port. A G_Port can determine operating mode at switch port initialization, F_Port when an N_Port attachment is determined, E_Port when an E_Port attachment is determined.

Answer option A is incorrect. The "Fabric" port within a Fibre Channel fabric switch provides a point-to-point link attachment to a single N_Port. F_Ports are intermediate ports in virtual point-to-point links between end ports,

for example N_Port to F_Port to F_Port to N_Port using a single Fibre Channel fabric switch.

QUESTION 76
Which of the following is the process in which production data is copied to a device at a remote location for data protection or disaster recovery purposes?

A. Snapshot
B. Remote replication
C. Local replication
D. Clone

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Remote replication is the process in which production data is copied to a device at a remote location for data protection or disaster recovery purposes. Remote replication can be either synchronous or asynchronous. Replication can take place in the storage array, at the host (server) or in the network.

Remote replication helps in protecting data, providing protection in case of primary site failure. Remote replication is used to provide a continuous, non-disruptive, host-independent solution for disaster recovery, data backup or migration over a long distance.

Remote replication leverages the power of storage systems for copying data between storage pools or systems in IP or FC SAN at various sites. Users can leverage the disk-based remote copy to restart service in minutes when source data fails due to system malfunction or accidents. If users need to transfer responsibility back to the source, it can be quickly synced with the remote copy for only differentials.

Answer option C is incorrect. Local replication is one of the replication methods used for data migration, which can be performed for various reasons, such as migrating from a small LUN to a larger LUN. Local replication can be used to test critical business data or applications. For instance, when an application upgrade is planned, it is tested using the local replication and if the test is successful, it can be restored to the source volumes.

Answer option A is incorrect. Snapshot is a virtual representation of a volume at a given point and is ideal for short-term backup and recovery purposes. Snapshot uses less space since only modified data is copied from the primary.

Snapshot is a point-in-time copy of a defined collection of data. It is space-efficient, any-point-in-time copies of changed bits or volumes, achieving time, cost and capacity efficient. Snapshot can be of files, LUNs, file systems, or any other type of container supported by the system it can be read-only or read-write copies of files, LUNs, file systems or any other type of container.

Answer option D is incorrect. Clones closely resemble snapshots, as they are bit-identical copies of full volumes, and thus require the same storage capacity as the primary volume.

Reference: http://searchdatabackup.techtarget.com/definition/remote-replication

QUESTION 77
What is the approximate maximum network bandwidth for Fibre Channel network in 4Gb/s SAN?

A. 400MB/S Full Duplex
B. 200MB/S Full Duplex
C. 1600MB/S Full Duplex
D. 800MB/S Full Duplex

Correct Answer: D
Section: (none)

Explanation

Explanation/Reference:
The approximate maximum network bandwidth for Fibre Channel network in 4Gb/s SAN is 800MB/s Full Duplex.

Answer option B is incorrect. The approximate maximum network bandwidth for Fibre Channel network in 1Gb/s SAN is 200MB/s Full Duplex.

Answer option A is incorrect. The approximate maximum network bandwidth for Fibre Channel network in 2Gb/s SAN is 400MB/s Full Duplex.

Answer option C is incorrect. The approximate maximum network bandwidth for Fibre Channel network in 8Gb/s SAN is 1600MB/s Full Duplex.

QUESTION 78
Which of the following are the most useful tools for troubleshooting TCP/IP networking problems?

Each correct answer represents a complete solution. Choose two.

A. fcping
B. rescan
C. ping
D. traceroute

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:
Ping and traceroute are the most useful tools for troubleshooting TCP/IP networking problems.

PING
The ping command-line utility is used to test connectivity with a host on a TCP/IP-based network. This is achieved by sending out a series of packets to a specified destination host. On receiving the packets, the destination host responds with a series of replies. These replies can be used to determine whether or not the network is working properly.

Traceroute
Traceroute is a route-tracing utility that displays the path an IP packet takes to reach its destination. It uses Internet Control Message Protocol (ICMP) echo packets to display the Fully Qualified Domain Name (FQDN) and the IP address of each gateway along the route to the remote host. Traceroute sends out a packet to the destination computer with the TTL field value of 1. When the first router in the path receives the packet, it decrements the TTL value by 1. If the TTL value is zero, it discards the packet and sends a message back to the originating host to inform it that the packet has been discarded. Traceroute records the IP address and DNS name of that router, and sends another packet with a TTL value of 2. This packet goes through the first router, and then times out at the next router in the path. The second router also sends an error message back to the originating host. Now, the process starts once again and traceroute continues to send data packets with incremented TTL values until a packet finally reaches the target host, or until it decides that the host is unreachable. In the whole process, traceroute also records the time taken for a round trip for each packet at each router.

Answer options A and B are incorrect. Fcping and rescan are the useful tools for troubleshooting Fibre channel network problems.

QUESTION 79
Which of the following components are defined by the iSNS standard?

Each correct answer represents a complete solution. Choose all that apply.

A. iSNS Databases
B. iSNS Servers
C. iSNS Clients
D. Initiator
E. iSNS Protocol
F. Target

Correct Answer: ABCE
Section: (none)
Explanation

Explanation/Reference:
The iSNS standard defines four components:
- iSNS Protocol: iSNSP is a protocol that specifies how iSNS clients and servers communicate. It is intended to be used by various platforms, including switches and targets as well as server hosts.
- iSNS Clients: iSNS clients are part of iSNSP aware storage devices. iSNS clients initiate transactions with iSNS servers using the iSNSP, register device attribute information in a common Discovery Domain (DD), download information about other registered clients, and receive asynchronous notification of events that occur in their DD.
- iSNS Servers: iSNS servers respond to iSNS protocol queries and requests made by iSNS clients using the iSNSP. iSNS servers initiate iSNSP State Change Notifications and store properly authenticated information submitted by a registration request in an iSNS database.
- iSNS Databases: iSNS databases are the information repositories for iSNS server. They maintain information about iSNS client attributes.

Answer options D and F are incorrect. Initiator and Target are the basic components of Serial Attached SCSI (SAS).

Reference: http://en.wikipedia.org/wiki/Internet_Storage_Name_Service

QUESTION 80
____ refers to a set of enhancements to Ethernet local area networks for use in data center environments.

Correct Answer: DCB
Section: (none)
Explanation

Explanation/Reference:
Data center bridging (DCB) refers to a set of enhancements to Ethernet local area networks for use in data center environments. Specifically, DCB goals are, for selected traffic, to eliminate loss due to queue overflow and to be able to allocate bandwidth on links. Essentially, DCB enables, to some extent, the treatment of different priorities as if they were different pipes. The primary motivation was the sensitivity of Fibre Channel over Ethernet to frame loss. The higher level goal is to use a single set of Ethernet physical devices or adapters for computers to talk to a Storage Area Network, Local Area network, and InfiniBand fabric.

Different terms have been used to market products based on the underlying Data Center Bridging standards:
- Data Center Ethernet (DCE) was a term originally coined and trademarked by Cisco Systems. DCE referred to Ethernet enhancements for the Data Center Bridging standards, and also including a Layer 2 Multipathing implementation based on the IETF's Transparent Interconnection of Lots of Links (TRILL) standard.
- Convergence Enhanced Ethernet or Converged Enhanced Ethernet (CEE) was defined from 2008 through January 2009 by group of including Broadcom, Brocade Communications Systems, Cisco, Emulex, HP, IBM, Juniper Networks, QLogic. The ad-hoc group formed to create version 0 proposals for enhancements that enable networking protocol convergence over Ethernet, especially Fibre Channel.

QUESTION 81
Which of the following DB-style connectors are used by the SCSI-2 drive interface standard?

Each correct answer represents a complete solution. Choose three.

A. 40-pin female
B. 68-pin female
C. 50-pin female
D. 25-pin female

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:
The following DB-style connectors are used by the SCSI-2 drive interface standard:
- 68-pin female
- 50-pin female
- 25-pin female

Small Computer System Interface (SCSI) is the second most popular drive interface in use today after the Integrated Drive Electronics (IDE) interface.

Answer option A is incorrect. There is no such DB-style connector as 40-pin female.

Reference: http://en.wikipedia.org/wiki/SCSI

QUESTION 82
Which of the following helps in providing communication between a target and an initiator in different VSANs?

A. SDH
B. ISL
C. SAS
D. IVR

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Inter-VSAN routing (IVR) helps in providing communication between a target and initiator in different VSANs, using IVR zone sets and IVR zones. This inter-VSAN routing (IVR) technology allows more scalability within a data center. It also can be used across long distances using either metropolitan-area network (MAN) technology or WAN technology with Fiber Channel over IP (FCIP).

Answer option C is incorrect. Serial Attached SCSI (SAS) is a computer bus used to move data to and from computer storage devices such as hard drives and tape drives. SAS depends on a point-to-point serial protocol that replaces the parallel SCSI bus technology. SAS offers backwards-compatibility with second-generation SATA drives.

Answer option B is incorrect. Inter-Switch Link (ISL) is a trunking method developed by Cisco to use for Ethernet and Token Ring trunk connections. Most of the Cisco switches and routers that support trunking also support ISL except some older switches such as Catalyst 4000 switches. ISL encapsulates the original frame by adding a 26-byte header and a 4-byte trailer.

Answer option A is incorrect. Synchronous Digital Hierarchy (SDH) is a standard technology for synchronous data transmission on optical media. It provides faster and less expensive network interconnection than traditional PDH (Plesiochronous Digital Hierarchy) equipment.

QUESTION 83
Which of the following is the smallest storage access unit on an IDE hard drive?

A. Spindle
B. Track
C. Sector
D. Cluster

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
A sector is the smallest storage access unit on an IDE hard drive. A hard disk drive sector is a subdivision of a track on a magnetic disk or optical disc. Each sector stores a fixed amount of data. The typical formatting of these media provides space for 512 bytes (for magnetic disks) or 2048 bytes (for optical discs) of user-accessible data per sector.

Answer option B is incorrect. A disk drive track is a circular path on the surface of a disk or diskette on which information is magnetically recorded and from which recorded information is read.

Answer option A is incorrect. Spindle provides axis for the rotation of platters.

Answer option D is incorrect. A cluster is a set of contiguous sectors and is treated as a single unit.

QUESTION 84
How many active partitions can be made on the hard disk of a computer?

A. 3
B. 4
C. 1
D. 2

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
An active partition is one that is read at startup and is expected to contain the necessary system files to boot the computer. If no partition is set as Active, a computer does not boot from the hard disk drive. There can be only one active partition in a computer. Only the primary partition can be marked as the active partition.

QUESTION 85
Which of the following RAID levels is also known as "Block striping with distributed parity"?

A. RAID 4
B. RAID 0
C. RAID 1
D. RAID 5

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
RAID 5 is also known as "Block striping with distributed parity", as the data is striped at the block level and the parity information is distributed across all disks.

Answer option A is incorrect. RAID 4 is also termed as "Block striping with dedicated parity disk", as the disks are striped at the byte level.

Answer option B is incorrect. RAID 0 is known as "Block striping without parity".

Answer option C is incorrect. RAID 1 is called "Mirroring of disks".

QUESTION 86
Which of the following terms is defined as a part of a system that, if it fails, will stop the entire system from working?

A. High availability
B. Vaulting
C. Component redundancy
D. Single point of failure

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
A single point of failure (SPOF) is a part of a system which, if it fails, will stop the entire system from working. They are undesirable in any system whose goal is high availability, be it a network, software application or other industrial system. Systems are made robust by adding redundancy in all potential SPOF and is generally achieved in computing through high-availability clusters. Redundancy can be achieved at the internal component level, at the system level (multiple machines), or site level (replication).

Answer option A is incorrect. The term High availability refers to the ability of a system to perform its function continuously for a significantly longer period of time than the reliabilities of its individual components would suggest. High availability is most often achieved through failure tolerance.

Answer option C is incorrect. Component redundancy helps in describing a component of a computer or network system that is used for guarding the primary system from failure by acting as a backup system. Redundant components can contain both hardware elements of a system, such as disk drives, peripherals, servers, switches, routers and software elements, such as operating systems, applications and databases.

Component redundancy is needed for protecting against types of physical hardware failures. If uptime is critical, some sort of high-availability system is implemented that can allow to quickly switch operations over to another device or datacenter in an emergency.

Answer option B is incorrect. Vaulting is the strategy of sending critical data out of the main location (off the main site) as part of a disaster recovery plan. The storage of off-site data is also known as vaulting, as backups are stored in purpose built vaults. Using removable storage media such as magnetic tape or optical storage, data is usually transported off-site.

QUESTION 87
How many devices can be connected together from a standard SCSI controller?

A. Eight devices, including one host bus adapter card
B. Four devices, including one host bus adapter card
C. Ten devices, excluding one host bus adapter card
D. Six devices, including one host bus adapter card

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

A standard SCSI controller allows up to eight devices to be connected to the SCSI bus, including one host bus adapter and up to seven drive controllers. Every SCSI device, including the host bus adapter (HBA), must be assigned a unique ID number. SCSI IDs begin at 0 and count upward. With 8-bit SCSI implementations, SCSI ID 7 has the highest priority and 0 the lowest.

QUESTION 88
Your company has decided to purchase a new rack mount server. An existing equipment rack in the server room has only 11.5 inches space available. Which of the following is the maximum size of the server that can be installed into the rack?

A. 7U
B. 6U
C. 9U
D. 4U

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
A rack unit is a unit of measure used to describe the height of equipment intended for mounting in a 19-inch rack or a 23-inch rack. One rack unit is 1.75 inches (44.45 mm) high. A rack unit is also known as U. Hence, a 1U product has a vertical measurement of 1.75"; 2U is 3.5"; 3U is 5.25" and so on. The rack unit size is based on a standard rack specification as defined in EIA-310.

According to the scenario, 11.5 inches space is available in the existing rack. Hence, the maximum size of the server that can be installed in the rack is 6U (i.e. 6 x1.75 = 10.5).

Answer options A and C are incorrect. The servers of size 7U and 9U will be bigger for the available space in the existing rack.

Reference: "http://en.wikipedia.org/wiki/Rack_unit"

QUESTION 89
Which of the following is a process of offloading backup procedures from a server so that the time ordinarily allotted to backup functions can be used to carry out other server tasks?

A. Server-based backup
B. Serverless backup
C. Online backup
D. LAN-free backup

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Server-less backup is a process of offloading backup procedures from a server. This is done so that the time ordinarily allotted to backup functions can be used to carry out other server tasks. Generally, the amount of time that a server can devote to process requests from applications is limited by the backup window - the amount of time that should be reserved for data backup. Server-less backup is a storage area network (SAN) solution that is designed for leading lower hardware costs and improving time-effectiveness, scalability, and fault tolerance.

Server-less backup is used to enable disk-to-tape or disk-to-disk backup without depending on server resources or network bandwidth.

Answer option D is incorrect. A LAN-free backup is a process of backing up of server data to a shared, central storage device without sending the data over the local area network (LAN). It is usually achieved by using a

storage area network (SAN). The goal of LAN-free backup is to reduce the load on LAN and reduce the time it takes to complete the backup. It helps in offering an alternative way of backup than a simple data copy to network-attached storage (NAS) over LAN.

It comes in different flavors:

- With backup server: In addition to a shared storage device (usually a traditional tape library), there exists a central server arbitrating access to device (for all the other SAN servers). The central server, however, does not handle data stream itself.
- Without backup server: The storage facility (usually a virtual tape library, or VTL) is smart enough to handle multiple data accesses without intermediate component.

The advantages of LAN-free backup contain shorter backup and recovery times and less disruption to other systems and applications.

Answer option A is incorrect. Server-based backup uses a server with large disk capacity for backing up data on other servers and workstations in the network. The data on this server is then copied to portable disks or tape for storage offsite. This solution helps in permitting for server-to-server replication. Using a second server placed at an offsite location and specialized software for replicating the data from one server to another over the Internet permits a user to copy the data from the backup server to the offsite server.

The benefit of using server-based backup accrues to those organizations that have huge amount of data. The speed of the backup increases because server-based disk drives is much faster than tape or portable disks. Also, the disk sizes on the server are much larger than what is currently present on portable drives. Server-based backup also permits for the copy process to take place during the day when other servers are being used for productive work.

Answer option C is incorrect. Online backup uses the Internet and a secure (or encrypted) connection for backing up data on the network to a third-party provider's server storage system.

Reference: http://searchstorage.techtarget.com/definition/serverless-backup

QUESTION 90
Which of the following is a data storage technique that automatically moves data between high-cost and low-cost storage media?

A. Hierarchical storage management
B. Queue depth
C. Cache
D. RAID

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Hierarchical storage management (HSM) is a data storage technique which automatically moves data between high-cost and low-cost storage media. HSM systems exist because high-speed storage devices, such as hard disk drive arrays, are more expensive (per byte stored) than slower devices, such as optical discs and magnetic tape drives. Instead, HSM systems store the bulk of the enterprise's data on slower devices, and then copy data to faster disk drives when needed. In effect, HSM turns the fast disk drives into caches for the slower mass storage devices. The HSM system monitors the way data is used and makes best guesses as to which data can safely be moved to slower devices and which data should stay on the fast devices.

Answer option B is incorrect. Queue depth is something that most storage administrators need to overlook as this is the reason for most performance issues. Queue depth is defined as the number of commands that the hba can send / receive in a single chunk - per LUN. From a host-hba point (initiator), it is defined as the number of commands that can be queued (or stored) and then sent to storage. From a storage point (target), it is defined as the number of commands it can accept in one-shot, again, per lun.

Performance degradation starts quickly once user reaches max queue depth on target and it starts backing off the queues and it will hit hard on user's response times. Ideally, if users have few hosts and luns, the default is good. If users have the luxury to dedicate ports for high performance applications, ensure that the queue depth is configured appropriately so that users get the best of performance.

Answer option C is incorrect. Cache is a component that transparently stores data so that future requests for that data can be served faster. The data that is stored within a cache (values) that have been computed earlier or duplicates of original values that are stored elsewhere. If requested data is contained in the cache (cache hit), this request can be served by simply reading the cache, which is comparatively faster. Otherwise (cache miss), the data has to be recomputed or fetched from its original storage location, which is comparatively slower. Hence the greater the number of requests that can be served from the cache, the faster the overall system performance becomes. To be cost efficient and to enable an efficient use of data, caches are relatively small.

Answer option D is incorrect. RAID is the term used to describe a storage system consisting of multiple disks acting as a single logical disk with differing levels of data redundancy and differing techniques of data distribution.

A redundant array of inexpensive disks (RAID), also called redundant array of independent disks, uses two or more derives in combination to create a fault-tolerant system that protects against physical hard drive failure and increases hard drive performance.

It is a technology that allows computer users to achieve high levels of storage reliability from low-cost and less reliable PC-class disk-drive components, via the technique of arranging the devices into arrays for redundancy.

QUESTION 91
Which of the following can cause a broadcast storm on the network?

Each correct answer represents a part of the solution. Choose two.

A. Fluctuation in voltage
B. A down network server
C. Faulty NIC on a networked computer
D. Slow workstation systems

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
A broadcast storm is a situation in which one or more network devices send jabber packets constantly and create traffic. In case of a faulty NIC that sends a jabber packet, you can detect the faulty NIC through the network monitor software or by the elimination process.

In case if a server is down, another computer on the network may attempt to communicate with the server by sending frequent requests, thus generating traffic on the network, which may result in a broadcast storm.

Answer options A and D are incorrect. Slow workstation systems or voltage fluctuation cannot cause a broadcast storm on the network.

QUESTION 92
What is the approximate maximum network bandwidth in 10 Gigabit Ethernet?

A. 2000MB/s full duplex
B. 1MB/s (half duplex only)
C. 20MB/s full duplex
D. 200MB/s full duplex

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
The approximate maximum network bandwidth in 10 Gigabit Ethernet is 2000MB/s Full Duplex.

Answer option B is incorrect. The approximate maximum network bandwidth in 10Base-T is 1MB/s (Half Duplex only).

Answer option C is incorrect. The approximate maximum network bandwidth in 100Base-T is 20MB/s Full Duplex.

Answer option D is incorrect. The approximate maximum network bandwidth in Gigabit Ethernet is 200MB/s Full Duplex.

QUESTION 93
A minimum of how many disks are needed for a RAID 0 array?

A. 2
B. 3
C. 4
D. 1

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
RAID 0 (striped disks) distributes data across multiple disks in a way that gives improved speed at any given instant. If one disk fails, however, all of the data on the array will be lost, as there is neither parity nor mirroring. In this regard, RAID 0 is non-redundant, so it should not be used for a critical system. A RAID 0 array requires a minimum of two drives. The diagrammatic representation of RAID 0 is shown below, where A, B, C, D, E, and F represent blocks:

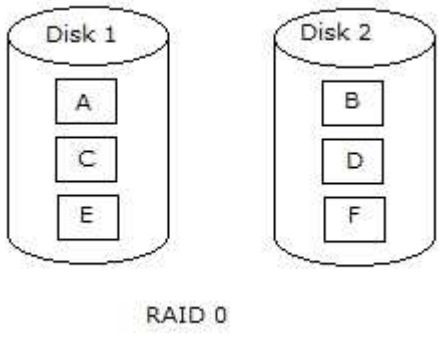

Reference: http://www.ecs.umass.edu/ece/koren/architecture/Raid/basicRAID.html

QUESTION 94
Which of the following is the storage interconnect used in mainframe enterprise computing environments?

A. HiPPI
B. VI
C. ESCON
D. FICON

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
ESCON (Enterprise System Connection) is the storage interconnects used in mainframe enterprise computing environments.

Answer option D is incorrect. FICON (Fibre Connector) allows ESCON assets to be used within a Fibre Channel SAN interface.

Answer option A is incorrect. HiPPI (High Performance Parallel Interface) connects devices at short distance and high speeds.

Answer option B is incorrect. VI (Virtual Interface) is an efficient, lightweight protocol that is being developed for message passing within clusters.

QUESTION 95
The Recovery Point Objective (RPO) of a company is 5 hours. Which of the following conclusions is correct based on this information?

A. 5 hours is the maximum limit to restart servers.
B. 5 hours is the maximum limit to resume production.
C. More than 5 hours of production data cannot be lost.
D. More than 5 hours of services have been recovered.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Recovery Point Objective (RPO) denotes the acceptable amount of data loss measured in time. Data must be recovered by the time of Recovery Point Objective (RPO) defined by an organization. Organizations determine RPO according to "acceptable loss" in a disaster situation. If the RPO of a company is 5 hours, it means that more than 5 hours of production data cannot be lost. In other words, it can be said that the data must be restored within 5 hours of the disaster.

QUESTION 96
Which architecture allows system-to-system communication and user level DMA with a very little CPU interaction?

A. FC-VI
B. FC-AE
C. FC-SB
D. FC-AV

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Fibre Channel Virtual Interface (FC-VI) is a standard for application-level distributed inter-process communication. It is based on Intel Corporation's v1.0 Virtual Interface (VI) Architecture. It allows system-to-system communication and user level DMA with as little CPU interaction. Fibre Channel Virtual Interface aims to provide a mapping between FC and VI.

Answer option B is incorrect. The Fibre Channel Avionics Environment (FC-AE) protocol has been designed for use in real time aircraft control system.

Answer option D is incorrect. The Fibre Channel Audio Video (FC-AV) protocol maps digital A/V formats onto Fibre Channel.

Answer option C is incorrect. FC-SB (Fibre Channel Single Byte) is the industry-standard command protocol for ESCON (Enterprise Systems Connection), a type of data connection, over Fibre Channel.

Reference: http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=50150

QUESTION 97
_____ encryption is a network security process that helps in applying crypto services at the network transfer layer above the data link level, but below the application level.

Correct Answer: Network
Section: (none)
Explanation

Explanation/Reference:
Network encryption (sometimes called network layer, or network level encryption) is a network security process that helps in applying crypto services at the network transfer layer - above the data link level, but below the application level.

Network encryption is implemented through Internet Protocol Security (IPSec). It is a set of open Internet Engineering Task Force (IETF) standards that is used in conjunction for creating a framework for private communication over IP networks. IPSec works through the network architecture in which end users and applications don't need to be altered in any way. Encrypted packets appear to be identical to unencrypted packets and can be easily routed through any IP network.

QUESTION 98
Which of the following classes of service helps in establishing a dedicated connection through the fabric before transferring data, much like a virtual private network?

A. Class 2
B. Class 3
C. Class 6
D. Class 1

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Class 1 service helps in establishing a dedicated connection through the fabric before transferring data, much like a virtual private network.

Answer option A is incorrect. Class 2 service does not require a dedicated channel. In Class 2 service, the receiving device sends an acknowledgment message in order to tell the sending device that the data has been received.

Answer option B is incorrect. Class 3 service also does not require a dedicated channel. It is a connection-less and unacknowledged service.

Answer option C is incorrect. Class 6 service is multi-cast (one-to-many) transmission through the fabric.

Reference: http://searchstorage.techtarget.com/tip/Service-classes-in-Fibre-Channel-SANs

QUESTION 99

Which of the following file permissions permits reading, writing, changing, and deleting of the file?

A. Modify
B. Write
C. Read
D. Full Access

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
The Full Access permission permits reading, writing, changing, and deleting of the file.

Read, Write, Modify, and Full Access are fundamental permissions that can be assigned to files and folders. These permissions may be summarized as follows:

- Read: Permits viewing and listing of files and subfolders - Permits viewing or accessing of the file's contents
- Write: Permits adding of files and subfolders - Permits writing to a file
- Modify: Permits reading and writing of files and subfolders; allows deletion of the folder - Permits reading and writing of the file; allows deletion of the file
- Full Access - Permits reading, writing, changing, and deleting of files and subfolders Permits reading, writing, changing and deleting of the file

QUESTION 100
Which of the following is a full duplex disk interface?

A. ATA
B. PATA
C. SATA
D. SAS

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Serial Attached SCSI (SAS) is a computer bus used to move data to and from computer storage devices, such as hard drives and tape drives. It is a Full-duplex interface with link aggregation (4-ports wide at 24 Gbit/s).

Answer options A, B, and C are incorrect. ATA, PATA, and SATA are not full duplex interfaces.

Exam C

QUESTION 1
A 300 GB disk can provide fault tolerance to which of the following?

Each correct answer represents a complete solution. Choose two.

A. 150 GB disk
B. 300 GB disk
C. 400 GB disk
D. 600 GB disk

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
A 300 GB hard disk can provide fault tolerance to a 150 GB disk and a 300 GB disk. A disk can provide fault tolerance to another disk, provided that the size of the disk does not exceed the size of the disk providing fault tolerance.

QUESTION 2
Which of the following is the RPM speed of the SAS drive?

Each correct answer represents a complete solution. Choose two.

A. 10000
B. 15000
C. 5400
D. 7000

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
SAS HDD's are primarily aimed for the enterprise market whose main goal is to achieve the highest performance. SAS can be found commonly at rotational speed of 10,000 RPM and 15,000 RPM.

QUESTION 3
You have 5 disks working fine in your storage network. To make a hot spare environment, how many disks should be there?

A. 7
B. 5
C. 10
D. 6

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
At least 10 disks are required for making the hot spare environment. Hot spare components are those that are online and configured as standbys into the system and are used as automatic replacements for failed components. Hot sparing does not require a user intervention.

Hot pluggable components are those that can be removed or replaced without interrupting server operations or data access. Disk in a fault-tolerance enclosure should be hot-pluggable.

QUESTION 4
What is the advantage of using LTO-4 tape drives?

Each correct answer represents a complete solution. Choose three.

A. Data transfer rate to 120 MB/s
B. Drive level encryption
C. Partition feature
D. Capacity up to 800 GB

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
LTO- 4 has the following advantages:
- Capacity up to 800 GB
- Increased data transfer rate to 120 MB/s (maximum)
- Introduced drive level encryption feature using 256-bit AES-GCM

Answer option C is incorrect. The partition feature was introduced with LTO-5 tapes. This feature allows tape to be "split" into 2 separately writable areas.

QUESTION 5
You work as a Storage Administrator for uCertify Inc. You need to implement a backup infrastructure for a remote office. This office has enough data to fill a single tape each night. You can remotely schedule and control your backup server backups throughout the night. The staff at this site is not technical and will not touch the equipment. You only need to store the last five days worth of data these tapes. What solution purchase and administrative costs?

A. Deploy a tape library with multiple tape drives to ensure quick backup and enough tape slots to support a new tape each night.
B. Deploy a single direct attached tape drive to the backup server and instruct onsite personnel how to swap tapes to ensure a new blank tape is available each night.
C. Deploy a tape autoloader with enough tape slots to support a new tape for each night
D. Deploy a set of five individual tape drives to the backup server each with a tape preloaded then rotate the tape drive you use each night.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
As you need to store only five days worth of data these tapes, you should deploy a tape autoloader with enough tape slots to support a new tape for each night. The tape autoloader contains one tape drive with many tapes and robotics that can insert and change the tapes. This system is ideal when one drive is sufficient for the read/ write performance requirement.

Answer option D is incorrect. You cannot insert and change tapes in a tape drive automatically. So deploying a set of five individual tapes will not help in this scenario.

Answer option A is incorrect. The tape library consists of multiple tape drives with lots of tapes and robotics that can insert and change tapes. This system is ideal when multiple drives are required to maintain read/write performance. Deploying a tape library will not help in this scenario.

Answer option B is incorrect. Instructing onsite personnel about swapping tapes to ensure a new blank tape is available each night is not a recommended action.

QUESTION 6
What is a difference between a tape autoloader and a tape library?

A. An autoloader uses robotics to insert and change tapes.
B. An autoloader contains one tape drive, whereas a tape library can contain multiple tape drives.
C. An autoloader provides automatic tape change, but a tape library is operated manually.
D. A tape library uses robotics to insert and change tapes.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
The tape autoloader contains one tape drive with many tapes and robotics that can insert and change the tapes. This system is ideal when one drive is sufficient for the read/write performance requirement.

The tape library consists of multiple tape drives with lots of tapes and robotics that can insert and change tapes. This system is ideal when multiple drives are required to maintain read/write performance.

QUESTION 7
Which of the following terms refers to the condition when the data transfer rate falls below the minimum threshold at which the tape drive heads were designed to transfer data?

A. Multiplexing
B. Shoe-shining
C. Thin provisioning
D. Multistreaming

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Shoe-shining is a disadvantage with read/write tape drives. During shoe-shining, the data transfer rate falls below the minimum threshold at which the tape drive heads were designed to transfer data to or from a continuously running tape. In this situation, the modern fast-running tape drive is unable to stop the tape instantly.

Answer options A and D are incorrect. Multiplexing is the process of alternately writing or interleaving multiple clients-data to a single tape. Multistreaming allows multiple backups (streams) of a single client to occur simultaneously to multiple tape drives.

Answer option C is incorrect. Thin provisioning is the act of using virtualization technology to give the appearance of more physical resource than is actually available.

QUESTION 8
You are transferring some files actively to another place through a network. The data is in encrypted form, but data on storage device is not encrypted. What type of encryption is this?

A. Data-at-rest encryption
B. Host-based encryption
C. Network encryption

D. Data-in-flight encryption

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Data in flight is a data security term that refers to data that is actively moving through a network. Data in flight is one aspect of a Data Loss Prevention (DLP) system. In case of data-in-flight, data is moving through a communications media such as a copper wire, optical cable, or even air (in the case of wireless communications). Data-at-rest refers to data that has been written to a media such as a disk drive, tape cartridge, or CD.

Answer option B is incorrect. Host encryption is a host-based encryption solution that helps in providing compliance and maintaining data confidentiality for data from the host to disk on storage. The main advantage of host encryption is that it helps in protecting information in storage environment in the event it becomes compromised through unauthorized access or disk removal.

Answer option A is incorrect. Data at rest is a term that is sometimes used to refer to all data in computer storage while excluding data that is traversing a network or temporarily residing in computer memory to be read or updated. In order to keep data at rest from being accessed, stolen, or altered by unauthorized people, security measures such as data encryption and hierarchical password protection are commonly used.

Answer option C is incorrect. Network encryption is a network security process that helps in applying crypto services at the network transfer layer.

QUESTION 9
You are a network administrator for uCertify.com. You are designing the network of the company. You are required to accomplish the following task:

- The cabling must support data transfer speed up to 50 Mbit/sec.
- The cost of cabling must be as low as possible.

Which of the following cables will you use?

A. CAT 2
B. CAT 5e
C. CAT 6
D. CAT 4

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
According to the question, you are required to use the cable that can support data speed up to 50 Mbit/sec and it must be the cheapest of the available choices. In order to fulfill both the requirements, you will have to use the CAT 5e cable.

Answer option C is incorrect. Although CAT 6 cables can also handle data speed up to 1000 Mbit/sec, they are much more expensive than CAT 5e cables.

Answer options A and D are incorrect. CAT 2 and CAT 4 cables are used in 4 Mbit/sec and 16 Mbit/sec token ring networks respectively.

QUESTION 10
You are connecting two buildings with a high-speed network. Distance between buildings is 1650 feet. Which of the following cables will you use?

A. 100BaseT
B. CAT5
C. CAT6
D. 1000Base-SX

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
1000BASE-SX uses short wavelength laser over multimode fibers. Its maximum segment length is 550 meters (1800 ft).

Answer options B, C, and A are incorrect. The maximum allowed length of a 100 Base T, CAT 5, or CAT 6 cable is 100 meters (330 ft).

QUESTION 11
Which of the following cables will provide maximum bandwidth?

A. CAT 6
B. CAT 2
C. CAT 4
D. CAT 5

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
CAT 6 provides the maximum speed of the above cable types. Cat 6 is a cable standard for Gigabit Ethernet and other network physical layers that is backward compatible with the Category 5/5e and Category 3 cable standards.

Answer option D is incorrect. CAT 5 provides speeds up to 100 Mbps.

Answer options B and C are incorrect. CAT 2 and CAT 4 cables are used in 4 Mbit/sec and 16 Mbit/sec token ring networks respectively.

QUESTION 12
You work as a Storage Administrator for uCertify Inc. After installing, cabling, zoning, and device masking for a new server, you completed HBA configuration successfully. You observe that hosts cannot see the storage after executing a device find command. After the investigation at the switch, you found that the host is not "seen" at the switch port. Light is visible at either end of the fiber, but no link light is seen on the switch. What is the most likely cause of the problem?

A. The zone has not been activated.
B. Transmit and receive fibers are reversed.
C. The switch ports are disabled.

D. The HBA on the subsystem is not online

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
The optical fiber always transmits from point A to point B while the other fiber is transmitting from B to A. Therefore, both ends of a full-duplex system have both transmitters and receivers. Because of this, you must be careful not to mix them up during installation by reversing transmit and receive fibers. Sometimes the fibers are color coded, or they may have some other identifying feature such as a ridge marking one fiber on duplex cables. In any case, be careful to not reverse the two fibers when installing connectors on the ends.

QUESTION 13
Which of the following devices are Hot Pluggable?

Each correct answer represents a complete solution. Choose two.

A. Disk
B. Power supply
C. HBA
D. Disk chassis

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Hot swapping and hot plugging are terms used to describe the functions of replacing storage system components without shutting down the system. Disks and HBAs are hot pluggable devices.

QUESTION 14
You need to connect several ISLs, what should you do?

A. Connect D ports of the switches
B. Connect switches via consoles
C. Connect E ports of the switches
D. Connect F ports of the switches

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
E port is the connection between two fibre channel switches. When E_ports between two switches form a link, that link is referred to as an inter-switch link (ISL). ISL stands for Inter Switch Link. If two switch ports are joined together, an inter switch link (ISL) is created. Links are the physical media that connect ports together in a Fibre Channel SAN.

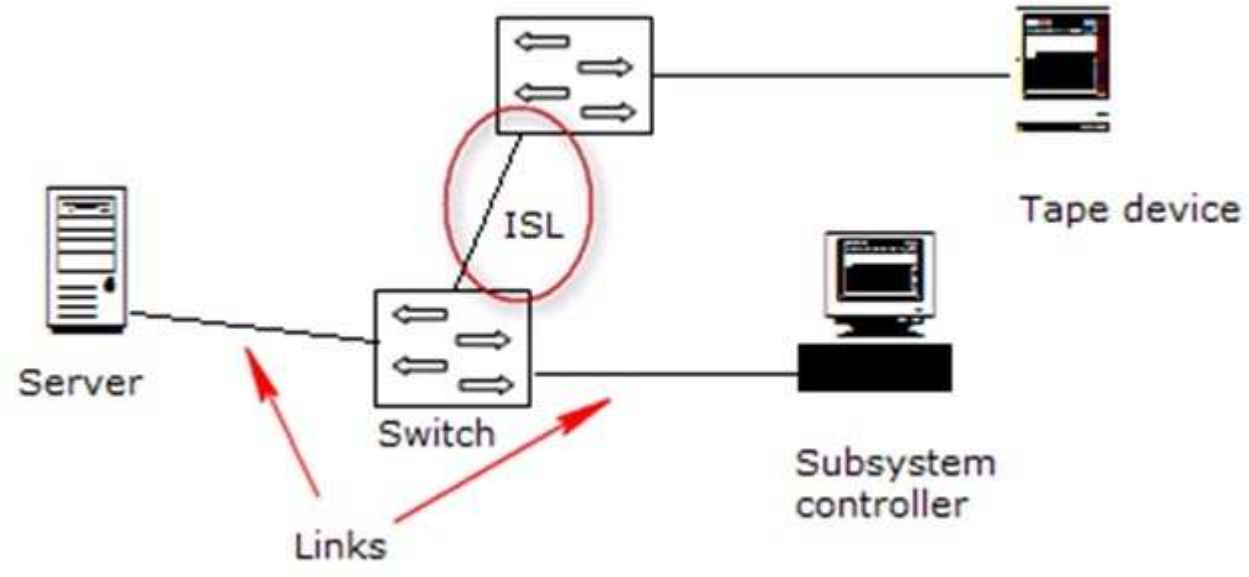

Answer option B is incorrect. Connecting switches via consoles is useless. The console port is used to connect a switch with a computer for configuration purposes.

Answer option D is incorrect. F port is a port on the switch that connects to a node point-to-point.

Answer option A is incorrect. D port is a diagnostic port, used solely for the purpose of running link-level diagnostics between two switches

QUESTION 15
You work as a storage administrator for uCertify Inc. You have set up an iSCSI network storage systems. You have configured a host with an iSCSI host bus adapter (HBA). You observe that storage system is not connecting to target LUNs (logical units). What should be your next step to resolve the issue?

A. Allow UDP port 6257
B. Allow TCP port 3260
C. Allow UDP port 636
D. Allow TCP port 389

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Port 3260 is the official port for iSCSI or Internet SCSI. This service allows clients (initiators) to send SCSI-based commands to a SCSI storage device on a remote terminal. Port 3260 is used with TCP/IP to allow 2 host computers to communicate SCSI commands using basic IP networks.

Answer option D is incorrect. Port 389 is a server port for the Lightweight Directory Access Protocol.

Answer option A is incorrect. Port 6257 is the default port used by WinMX P2P (Peer-to-Peer) file sharing program.

Answer option C is incorrect. The 636 port is used for connecting to directory services being run on a TCP/IP connection.

QUESTION 16

What is Data STaaS?

A. Platforms with simple TCP/IP implementations
B. Third party remote storage subscription provider
C. Third party storage application provider
D. Multi-threaded TFTP Server for uploads and download

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Storage-as-a-Service (STaaS) is a comprehensive set of enterprise-class solutions and services that deliver secure and accessible cloud storage for business data. Storage-as-a-Service makes stored data actionable, reduces compliance risk, and maximizes efficiencies.

QUESTION 17
Which of the following medium is used by FCoE?

A. Bluetooth
B. Ethernet
C. FDDI
D. Wi_Fi

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Fibre Channel over Ethernet (FCoE) is an encapsulation of Fibre Channel frames over Ethernet networks. It allows Fibre Channel to use 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol. FCoE maps Fibre Channel over selected full duplex IEEE 802.3 networks for providing I/O consolidation over Ethernet and reducing network complexity in the datacenter. The FCoE protocol specification replaces the FC0 and FC1 layers of the Fibre Channel stack with Ethernet.

QUESTION 18
Which of the following is the most important criteria for implementing Rack stabilization?

A. Humidity
B. Weight consideration
C. Antistatic devices
D. Cooling

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
It is very important that you have proper rack to store your equipment. Make sure that there is no hindrance to the stability of the rack. You should always store heavy objects on lower shelves of racks. Try to store materials inside cabinets, files, and lockers.

QUESTION 19
You work as a Storage Administrator for uCertify Inc. An HBA has failed in an application server that connects to a Fabric that uses WWPN zoning. Which of the following steps should be completed after replacing the HBA to bring the application server back online?

A. Change the application metadata server
B. Change the WWPN login of the HBA
C. Change the LUN masking configuration
D. Change the zone configuration

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
WWPN zoning uses the world wide port name to identify the HBA ports on the other ends of the cables. After replacing the HBA, the WWPN name will probably change and so the zoning must be changed. You would normally use aliases to make the WWPNs more comprehensible, so if you do this, then you just need to change the alias. WWPN zoning is essential if you want to use virtualization.

QUESTION 20
You work as a Storage Administrator for uCertify Inc. You have recently upgraded the HBA driver of the storage network. You observed that a UNIX server has been unable to mount its file systems after the HBA driver has been replaced. What is the cause of the problem?

A. The file system table was not updated to reflect the changes in device IDs
B. The disk array firmware was not updated
C. Worldwide naming has not been updated on the switch
D. The switch zoning was based on port assignments

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
The UNIX server has not been able to mount its file system because the world wide naming has not been updated on the switch. WWN zoning uses name servers in the switches to either allow or block access to particular World Wide Names (WWNs) in the fabric. WWN zoning is susceptible to unauthorized access, as the zone can be bypassed if an attacker is able to spoof the World Wide Name of an authorized HBA.

Reference: http://www.sansecurity.com/san-security-faq.shtml

QUESTION 21
Which of the following provides block level architecture?

Each correct answer represents a complete solution. Choose two.

A. CIFS
B. FCIP
C. iFCP
D. NTFS

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
iFCP and FCIP provide block level architecture. iFCP and FCIP are intended for interconnecting storage area network (SAN) devices to support data movement. iSCSI, iFCP, and FCIP are network transport protocols for moving block data over IP-based networks.

Answer options A and D are incorrect. CIFS and NTFS are the component of file level architectures.

QUESTION 22
Which default network file sharing protocol is used by Microsoft Windows clients in a NAS environment?

A. CIFS
B. IPX
C. NFS
D. DNS

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
The Common Internet File System (CIFS), also known as Server Message Block (SMB), is a network protocol whose most common use is sharing files on a Local Area Network (LAN). NAS connects the storage to the TCP/IP network and helps in enabling users to access storage resources using both CIFS (Microsoft based Common Internet File System) and NFS (Unix based Network File System) protocols over the familiar TCP/IP networks.

QUESTION 23
Which of the following is a Windows-based protocol?

A. NFS
B. TLS
C. FCIP
D. CIFS

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
The Common Internet File System (CIFS), also known as Server Message Block (SMB), is a network protocol whose most common use is sharing files on a Local Area Network (LAN). NAS connects the storage to the TCP/IP network and helps in enabling users to access storage resources using both CIFS (Microsoft based Common Internet File System) and NFS (Unix based Network File System) protocols over the familiar TCP/IP networks.

QUESTION 24
Which of the following is a neighbor discovery protocol, and can be used by any vendor?

A. LLDP
B. TFTP
C. CDP
D. LDAP

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Link Layer Discovery Protocol (LLDP), standardized by the IEEE as part of 802.1ab, enables standardized discovery of nodes in multi-vendor networks.

Answer option C is incorrect. Cisco Discovery Protocol (CDP) is a Cisco proprietary neighbor discovery protocol and cannot be used in multi-vendor environment.

Answer option D is incorrect. Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing and maintaining distributed directory information services over an IP network.

Answer option B is incorrect. Trivial File Transfer Protocol (TFTP) is a file transfer protocol used for automated transfer of configuration or boot files between machines in a local environment.

QUESTION 25
You work as s Storage Administrator for uCertify Inc. There is inconsistency in communication in the network. What should be used?

A. Class of service
B. Capacity planning
C. ACL
D. QoS

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Quality of Service (QoS) traffic management guarantees a level of service in terms of bandwidth and delivery time.

Answer option A is incorrect. Class of Service (CoS) is a way of managing traffic in a network by grouping similar types of traffic together and treating each type as a class with its own level of service priority. Unlike Quality of Service (QoS) traffic management, Class of Service technologies do not guarantee a level of service in terms of bandwidth and delivery time; they offer a "best-effort."

Answer option B is incorrect. Capacity planning is the process of determining the production capacity needed by an organization to meet changing demands for its products.

Answer option C is incorrect. Storage administrators use access control lists (ACLs) to manage storage devices and access to them. An ACL consists of one or more entries, where each entry grants permissions to a scope.

QUESTION 26
You work as a Storage Administrator for uCertify Inc. You have added new LUN in your storage network, but clients are not able discover the new LUN. What should you do next?

A. Upgrade the firmware
B. Do a rescan of the storage
C. Reboot the client machines
D. Configure zoning

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
In this case, you need to do a rescan of storage. A rescan of storage devices is needed when a storage device has been added, removed, or changed from the array. The rescan and data store detection are asynchronous processes. As a result, the detection process for new data stores may complete before the detection process for new LUNs is complete.

Answer option D is incorrect. Zoning allows an administrator to control who can see what in a SAN.

Answer option C is incorrect. If rescanning is not helpful, then you may consider this option.

Answer option A is incorrect. There is no need to upgrade the firmware, you should first rescan the storage, and then reboot the client system.

QUESTION 27
Conflicting zone configuration can create which of the following basic types of zone merge errors?

Each correct answer represents a complete solution. Choose two.

A. Configuration mismatch
B. Connectivity
C. Misconfigured Fibre Channel cable
D. Content mismatch

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Zoning helps in enabling access control between storage devices and user groups. Zones are created to increase network security and to prevent data loss or corruption.

Zoning limits access to specific storage devices, increases security, and decreases traffic over the network. Zoning is configuration for restricting which node ports can access which other node ports in the fabric. With zoning, make sure that the zones are configured correctly as when two fabrics are merged, there is no guarantee that the fabrics will merge since it will check the zone configurations between the two fabrics. Conflicting zone configuration can create two basic types of zone merge errors:

- Zoning configuration mismatch: In this, enabled zone configurations are different in each fabric.
- Zoning content mismatch: In this type of error, the definition of a zone object in one fabric is different from the definition of a zone object with the same name in the other fabric.

Answer options C and B are incorrect. These are other types of fibre channel problems.

QUESTION 28
Which of the following statements are true regarding Direct-Attached Storage (DAS)?

Each correct answer represents a complete solution. Choose two.

A. Direct access to data is available from single server.
B. Storage is managed by a single host.
C. Storage devices can only be external disk subsystems.
D. DAS uses the file I/O architecture.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
DAS stands for Direct-Attached Storage. It is a traditional way of implementing storage managed by a single host. Other host can access the storage only through that single host, over the local network. Storage devices are to a server's I/O bus.

Answer option D is incorrect. DAS uses the block I/O architecture.

Answer option C is incorrect. Storage devices can be internal or external disks attached to a SCSI bus usually.

QUESTION 29
The network-attached storage (NAS) device is a file server running on a dedicated device connected to the network. Which of the following are the advantages of the NAS device?

Each correct answer represents a complete solution. Choose three.

A. It is used when computer systems are retired and disk drives are required to be zeroed out.
B. It permits the addition of more hard disk storage space to a network without shutting them down for maintenance and upgrades.
C. It enables the Windows, UNIX/Linux, Mac OS, and Novell clients to use the same storage and to access the shared data.
D. It is easy to install, configure, and manage by using a web browser.

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:
The network-attached storage (NAS) device is a server used for file sharing. It is a file server running on a dedicated device connected to the network. These types of devices are based on Linux or UNIX derivatives.

Various advantages of the NAS device are as follows:
- It is easy to install, configure, and manage by using a web browser.
- It enables the Windows, UNIX/Linux, Mac OS, and Novell clients to use the same storage and to access the shared data because it can communicate with the network by using a variety of protocols, such as the TCP/IP, IPX/SPX, NetBEUI, or AppleTalk protocols.
- It does not required to be located within the server, but it can be placed anywhere in LAN.
- It permits the addition of more hard disk storage space to a network without shutting them down for maintenance and upgrades.

Answer option A is incorrect. Hard drive wiping is a degaussing process that applies a strong magnetic field for the purpose of initializing the media. It is used to remove data from the hard drive permanently. It is used when computer systems are retired and disk drives are required to be zeroed out.

QUESTION 30
Which of the following storage methods uses CIFS and NFS?

A. SAN
B. NAS
C. CAS
D. DAS

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
NAS connects the storage to the TCP/IP network and helps in enabling users to access storage resources using both CIFS (Microsoft based Common Internet File System) and NFS (Unix based Network File System) protocols over the familiar TCP/IP networks.

QUESTION 31
You have six disks of 250GB each in a RAID type 5. What is the maximum capacity you will get?

A. 1.25 TB
B. 1.5 TB
C. 250 GB

D. 500 GB

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
In a RAID 5 array, one hard disk is dedicatedly used for keeping the parity information. In the above question, you will be able to use only five hard disks for storage purpose as one of the hard disks will be used for parity. Each hard disk can store data up to 250 GB, so there will be only 1.25 TB space left for storing files.

QUESTION 32
Which of the following RAID types use parity for data protection?

Each correct answer represents a complete solution. Choose two.

A. RAID 10
B. RAID 5
C. RAID 1
D. RAID 4

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
RAID type 4 and RAID type 5 use parity for data protection.

The following table depicts the differences between the RAID levels:

RAID Level of the Volume	Main Characteristics	Fault Tolerance	Performanc
RAID 0	Striping, but no parity	No	Yes
RAID 0 + 1	Striping with mirroring	Yes	Yes
RAID 1	Mirroring or duplexing, but no parity	Yes	Yes in readin
RAID 1 +	Mirroring with striping	Yes	Yes
RAID 2	Bit-level striping with parity	Yes	Parity slows d improves it.
RAID 3	Byte-level striping with parity	Yes	Parity slows d improves it.
RAID 4	Block-level striping with parity	Yes	Parity slows d improves it.
RAID 5	Both data and parity are striped.	Yes	Parity slows d improves it.

QUESTION 33
RAID 3 is configured; you have five disks, how many disks are required to provide a full tolerance environment?

A. 6
B. 5
C. 7
D. 10

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
RAID 3 is also called data striping with dedicated parity and parallel access. One drive is dedicated for parity. Data is regenerated in the event of a drive failure. So, in the above scenario, you are required to have six disks to provide fault tolerance.

QUESTION 34
How would you configure a storage system to prevent it from crashing when storage is full, without wasting too much unused storage space?

A. Disk virtualization
B. RAID 5
C. Thin provisioning

D. De-duplicatoin

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Thin provisioning is the act of using virtualization technology to give the appearance of more physical resource than is actually available. Thin provisioning is considered as a key component of storage. It permits over-allocation of storage capacity to get the enhanced storage utilization, increased application uptime, and simplified storage capacity management. It reduces downtime and enhances reliability with business continuity and disaster recovery. It can increase energy efficiency by making use of less storage.

QUESTION 35
Which of the following eliminates white space which helps avoid the poor utilization rates?

A. Zoning
B. Thin provisioning
C. Thick provisioning
D. LUN masking

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Thin provisioning is the act of using virtualization technology to give the appearance of more physical resource than is actually available. This methodology eliminates almost all white space which helps avoid the poor utilization rates. Thin Provisioning can be implemented in storage systems or in host software. The basic idea of thin provisioning is to consume storage capacity as data is written, as opposed to the traditional way of reserving capacity in advance and then filling it with written data.

Thin Provisioning works well for most applications, but there are some scenarios that you should avoid, such as applications where there is a high percentage of data being created and deleted.

Answer option C is incorrect. Thick provisioning allows you to allocate or reserve storage space while initially provisioning the virtual disk. The allocated storage space for the thick provisioned virtual disk is guaranteed. This operation ensures that there are no failures because of lack of storage space.

Answer option D is incorrect. LUN Masking is used to control device visibility by allowing specific devices (for example, tape drives or robots) to be seen only by a selected group of host initiators.

Answer option A is incorrect. Zoning is the partitioning of a Fibre Channel fabric into smaller subsets to restrict interference, add security, and to simplify management.

QUESTION 36
Which of the following automates the task of right sizing the thin provisioned storage, without affecting application uptime?

A. Thin reclamation
B. Thick provisioning
C. Zoning
D. LUN masking

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Thin reclamation automates the task of right sizing the thin provisioned storage, without affecting application uptime.

Answer option B is incorrect. Thick provisioning allows you to allocate or reserve storage space while initially provisioning the virtual disk. The allocated storage space for the thick provisioned virtual disk is guaranteed. This operation ensures that there are no failures because of lack of storage space.

Answer option C is incorrect. Zoning is the partitioning of a Fibre Channel fabric into smaller subsets to restrict interference, add security, and to simplify management.

Answer option D is incorrect. LUN Masking is used to control device visibility by allowing specific devices to be seen only by a selected group of host initiators.

QUESTION 37
Which of the following is an authorization process that makes a LUN available to some hosts and unavailable to other hosts?

A. SAN zoning
B. Port zoning
C. WWN zoning
D. LUN masking

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
LUN Masking is used to control device visibility by allowing specific devices to be seen only by a selected group of host initiators.

Answer option A is incorrect. SAN zoning is a method of arranging Fibre Channel devices into logical groups over the physical configuration of the fabric.

Answer option B is incorrect. Port zoning utilizes physical ports to define security zones.

Answer option C is incorrect. WWN zoning uses name servers in the switches to either allow or block access to particular World Wide Names (WWNs) in the fabric.

QUESTION 38
Which Security method is used to insure that a particular host only has access to a particular LUN?

A. Soft zoning
B. LUN masking
C. Port binding
D. Hard zoning

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
LUN Masking is used to control device visibility by allowing specific devices to be seen only by a selected group of host initiators.

Answer option C is incorrect. The switch-based fabric supports port binding. Port binding is implemented by

defining the WWPN that may log into a specific port. All ports in the fabric will have a WWPN, which is assigned port by port.

Answer option D is incorrect. Hard zoning restricts actual communication across a fabric

Answer option A is incorrect. Soft zoning restricts only the fabric name service, to show only an allowed subset of devices.

QUESTION 39
You are configuring a host to have access to three virtual devices. You want to strip these devices into one logical volume. Which of the following are the ways used to configure storage virtualization?

Each correct answer represents a complete solution. Choose three.

A. Layer based
B. Host based
C. Storage device based
D. Network based

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:
Storage virtualization is a storage technique used in a storage area network (SAN) for the management of storage devices. The various ways used to configure storage virtualization are as follows:
- Host-based
- Storage device-based
- Network-based

Answer option A is incorrect. This is an invalid answer option.

QUESTION 40
Which of the following protocols is required to setup automatic alerts to measure the performance of a storage area network?

A. ICMP
B. SNMP
C. RDP
D. LDAP

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Storage Management performs alert monitoring and logging. Some alerting software sends an SNMP trap if the operating system's SNMP service is installed and enabled. Every alert should consist of the following information:
- Severity level
- ID of the alert
- Category of alert
- Description
- Date and Time

QUESTION 41
Which of the following is an open management protocol?

A. CDP
B. SNMP
C. SMI-S
D. NetBIOS

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Storage Management performs alert monitoring and logging. Some alerting software sends an SNMP trap if the operating system's SNMP service is installed and enabled. Every alert should consist of the following information:

- Severity level
- ID of the alert
- Category of alert
- Description
- Date and Time

QUESTION 42
You noticed that disk write is slow. How will you confirm this?

A. Compare the size of buffer credit with virtual memory.
B. Upgrade the firmware and check the performance.
C. Isolate the disk and test it in a lab environment.
D. Compare the performance with recorded base line.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
To understand storage performance, you must have a baseline to compare against. You should establish a baseline of what is normal or typical performance for your storage environment as well as for the applications running in it. Recording baseline allows you to manage the predictable performance changes and more importantly.

QUESTION 43
Which of the following are management protocols?

Each correct answer represents a complete solution. Choose two.

A. WBEM
B. TFTP
C. TLS
D. SMI-S

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
SMI-S defines a method for the interoperable management of a heterogeneous Storage Area Network (SAN).

Web-based Enterprise Management (WBEM) is a set of management and internet standard technologies

developed to unify the management of enterprise computing environments.

Answer options B and C are incorrect. These are not the management protocols. TFTP is used for trivial file transfer, and TLS is a transport layer security protocol.

QUESTION 44
Which of the following provides data immutability?

A. NFS
B. CAS
C. NAS
D. DAS

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Mutable storage is the term for data that can be overwritten at any time. Immutable storage cannot be overwritten, but can be changed. Content-addressable storage (CAS) is a mechanism for storing information that can be retrieved based on its content, not its storage location. It is typically used for high-speed storage and retrieval of fixed content e.g., electronic documents, contracts, claims, CAD/CAM designs, presentations, medical X-rays, MRIs, security surveillance, audio voicemail, etc.

Answer options C, D, and A are incorrect. NAS, DAS, and NFS do not provide data immutability.

QUESTION 45
Which of the following is the removal of redundant data after a backup is complete?

A. Post Processing De-duplication
B. In line de-duplication
C. Data tiering
D. Zoning

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Postprocessing de-duplication (PPD) is the analysis and removal of redundant data after a backup is complete and data has been written to storage. It is also known as asynchronous de-duplication.

Answer option B is incorrect. In-line de-duplication is performed on the client side or when data is being transferred from the source to server.

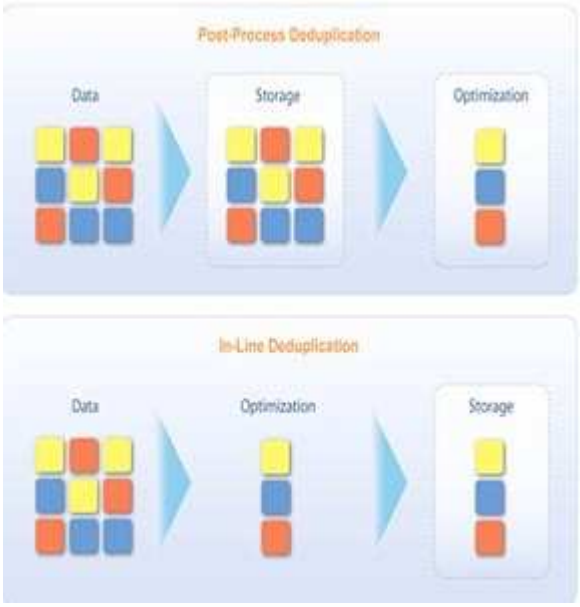

Answer option C is incorrect. Data tiering monitors data usage in order to determine where data should be stored. Data tiring is a type of software program that moves data files, volumes, or blocks between tiered storage according to company-defined policy.

Answer option D is incorrect. Zoning is the partitioning of a Fibre Channel fabric into smaller subsets to restrict interference, add security, and to simplify management.

QUESTION 46
Which of the following is most benefited from data de-duplication?

A. Zip file
B. PDF file
C. XML file
D. JPEG file

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Data de-duplication is a specialized data compression technique for eliminating coarse-grained redundant data. XML files will have the most benefit from de-duplication. JPEG, PDF, and Zip files are already compressed files.

The technique is used to improve storage utilization and can also be applied to network data transfers to reduce the number of bytes that must be sent across a link.

QUESTION 47
Which of the following is used to provide power to a system when the primary source of power is unavailable?

A. Cache battery backup
B. Virtual backup
C. Virtual tape library
D. Cache mirroring

Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:
A cache backup battery is used to provide power to a system when the primary source of power is unavailable. The battery-backup events occur when a modification in the battery-backup cache is detected.

Battery backup is a popular method that helps in protecting the cache against unexpected power loss. If the external power source (the wall power) is removed, internal or external batteries keep power applied to the cache memory chips. Since memory chips consume minimal power, small, low-capacity batteries can be used for maintaining the cache, sometimes resulting in nominal cost savings.

Answer option D is incorrect. Cache mirroring helps in protecting against cache data loss that occurs when a storage processor fails. As a prerequisite to cache mirroring, the storage system should have redundant storage processors with failover capability. If one controller fails, the other controller is assumed to have the ownership of the failed processor's storage and continues where it left off.

A complex operation, cache mirroring is used to transfer large amounts of rapidly changing data between controllers. Many storage systems are implemented such that the cache data path between controllers is the same as the data path (for example, SCSI bus or Fibre Channel loop) to the disk drives.

Answer option C is incorrect. A virtual tape library (VTL) is a data storage virtualization technology used typically for backup and recovery purposes. A VTL is used to present a storage component (usually hard disk storage) as tape libraries or tape drives for use with existing backup software.

Answer option B is incorrect. This is an invalid option.

QUESTION 48
You work as a storage Administrator for uCertify Inc. The company has experienced a localized disaster; the Business Continuance Plan is reviewed. The plan is revised to include a Recovery Time Objective (RTO) of four (4) hours and a Recovery Point Object (RPO) of one hour. What two processes must be implemented to support this requirement?

Each correct answer represents a complete solution. Choose two.

A. Hot or warm site
B. Operating system and application reinstallation media
C. Database snapshots of data and applications
D. Mirrored snapshots of data and applications

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
You can take advantage of a mirror database in the above scenario. To use a mirror database, you can create a database snapshot on the mirror database and direct client connection requests to the most recent snapshot. A database snapshot is a static, read-only, transaction-consistent snapshot of its source database as it existed at the moment of the snapshot's creation.

A hot site is a duplicate of the original site of the organization, with full computer systems as well as near-complete backups of user data. A warm site is a compromise between hot and cold. These sites will have hardware and connectivity already established, though on a smaller scale than the original production site or even a hot site.

Answer option C is incorrect. Unlike the mirror database itself, a database snapshot is accessible to clients. As long as the mirror server is communicating with the principal server, you can direct reporting clients to connect to a snapshot. Note that because a database snapshot is static, new data is not available.

Answer option B is incorrect. Operating system and application reinstallation media will not help in this scenario.

QUESTION 49
Which of the following allows you to make identical copies of the data on your volumes as you need them?

A. Disk mirroring
B. Synchronous replication
C. Cloning
D. Snapshot

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Cloning provides an additional layer of protection because it allows you to make identical copies of the data on your volumes as you need them. Clones closely resemble snapshots, as they are bit-identical copies of full volumes, and thus require the same storage capacity as the primary volume.

Answer option D is incorrect. Snapshots are space-efficient, any-point-in-time copies of changed bits or volumes, achieving time, cost, and capacity efficiency. Snapshots can be read-only or read-write copies of files, LUNs, file systems, or any other type of container.

Answer option A is incorrect. Disk mirroring is replication within locally connected disks, and it is the fault-tolerant mechanism used to define RAID policies.

Answer option B is incorrect. Synchronous and asynchronous replication creates a second copy of volumes and snapshots to capture continuous changes. Synchronous replication is used by mission critical environments to ensure application consistency, an RPO of no data loss and RTO of seconds, scalability, and highest performance SLA's. Asynchronous replication is used when RTOs of seconds to 30 minutes or longer are allowable.

QUESTION 50
Which of the following is considered the most important criteria when designing an OLTP implementation for a financial services customer?

A. Data integrity
B. Data striping
C. Data shredding
D. Data repository

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Data integrity is the property that data has not been altered or destroyed in an unauthorized manner.

Answer option B is incorrect. Data striping is a disk array data mapping technique in which fixed-length sequences of virtual disk data addresses are mapped to sequences of member disk addresses in a regular rotating pattern.

Answer option C is incorrect. Data shredding is a process for deleting data that is intended to make the data unrecoverable.

Answer option D is incorrect. Data repository is a central place of storage in which data is kept and maintained in an organized way.

QUESTION 51
Which of the following policy should be used to keep back up to protect against the natural disaster?

Each correct answer represents a complete solution. Choose three.

A. Dedicate and empower staff.
B. Try and test disaster recovery solution product.
C. Design a dependent plan.
D. Align disaster recovery with application development.

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:
A natural disaster refers to the results of the combination of natural hazards (physical events such as volcanic eruptions, landslides, sinkholes, blizzards, drought, hailstorms, heat waves, hurricanes, tropical storms, typhoons, ice ages, tornadoes, earthquakes, landslides, etc.) and human activities.

Disaster recovery consists of processes, policies, and procedures related to the preparation for recovery that is critical to an organization after a natural disaster. The following are some best practices for disaster recovery:
- Dedicate and empower staff.
- Divide and conquer.
- Design an independent plan.
- Provide facilities.
- Align disaster recovery with application development.
- Try and test disaster recovery solution products.

QUESTION 52
Which of the following disaster recovery setups uses a stage backup approach?

A. D2D
B. D2T
C. D2R2T
D. D2D2T

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
D2D2T (Disk to Disk to Tape) uses a stage backup approach. This would be used in conjunction with split mirrors or some other sort of remote replication.

Answer option B is incorrect. D2T (Disk to Tape) move our information from disk to tape. This is the traditional way of doing backups.

Answer option A is incorrect. D2D (Disk to Disk) provides faster backup and recovery. It is already in usable format.

Answer option C is incorrect. In D2R2T (Disk to Router to Tape), we do not store the information on the router. D2R2T is actually a third party copy where a storage router is used instead of media server.

QUESTION 53
RTO is 2 minute, and RPO is 5 minutes. Which of the following devices should be used for backup?

Each correct answer represents a complete solution. Choose two.

A. SATA
B. SSD
C. Tape library
D. Tape

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:
RPO and RTO, both are calculated when planning for disaster recovery. So, you should use the tape and tape library for the above scenario as these two are most commonly used for such scenarios.

QUESTION 54
A company has planned that the data must be restored within 5 hours of the disaster. What is the term used for this plan?

A. Downtime
B. BIA
C. RPO
D. RTO

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
The Recovery Time Objective (RTO) is the duration of time and a service level within which a business process must be restored after a disaster or disruption in order to avoid unacceptable consequences associated with a break in business continuity.

Answer option C is incorrect. Recovery Point Objective (RPO) denotes the acceptable amount of data loss measured in time. Data must be recovered by the time of Recovery Point Objective (RPO) defined by an organization. Organizations determine RPO according to "acceptable loss" in a disaster situation. If the RPO of a company is 5 hours, it means that more than 5 hours of production data cannot be lost. In other words, it can be said that the data must be restored within 5 hours of the disaster.

Answer option A is incorrect. The term downtime is commonly used to refer to periods when a system or service is unavailable.

Answer option B is incorrect. A business impact analysis (BIA) is a crisis management and business impact analysis technique that identifies those threats that can impact the business continuity of operations.

QUESTION 55
Which of the following are related to access management?

Each correct answer represents a complete solution. Choose two.

A. HBA
B. LUN
C. ACL
D. Physical security

Correct Answer: CD

Section: (none)
Explanation

Explanation/Reference:
Storage administrators use access control lists (ACLs) to manage storage devices and access to them. An ACL consists of one or more entries, where each entry grants permissions to a scope.

Physical security describes measures that are designed to deny access to unauthorized personnel from physically accessing a building, facility, resource, or stored information.

QUESTION 56
What is Exchange Completion Time?

A. The duration in which a disk drive is replaced by another disk in a fault tolerance environment
B. The duration in which a SCSI bus rescans the newly connected HBA
C. The duration in which a transaction takes from being initiated by the host, received by the storage, and acknowledged back by the host in real-time
D. The maximum amount of work that an organization is capable of completing in a given period

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Exchange Completion Time is the time duration in which a transaction is initiated by the host, received by the storage, and acknowledged back by the host in real-time. Measuring performance via Exchange Completion Times enables measurement and monitoring of storage I/O performance, hence ensuring that applications can be correlated and assigned to their most cost-effective storage tier without sacrificing SLAs. By incorporating the Exchange Completion Time metric i.e., measuring I/O conversations across the SAN infrastructure into your tiering considerations, tiering can now accurately be based on comprehensive real time performance as opposed to device specific views.

QUESTION 57
You work as a Storage Administrator for uCertify Inc. You observe that there is a high latency in data reading. What should be your first step to troubleshoot the performance?

A. Compare current utilization against established baseline
B. Isolate cause of resource overload
C. Set thresholds for resource overload
D. Check host or network application

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Latency is a highly reliable indicator of changes in performance and is often one of the first indicators of a resource overload. When latency increases, you should take following steps to troubleshoot the performance:
- Isolate the cause that might be causing a resource overload.
- Compare current utilization levels against your established baselines.
- The recommended order for drilling into resources is: CPU, disk, and network.
- Also check for other factors, such as network, host OS, or application performance, can also impact performance.

QUESTION 58
Why is it necessary to estimate the queue depth for each node while designing a configuration for a large SAN?

A. To run a bus scan
B. To avoid application failures
C. To update HBA driver
D. To configure zones

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
The queue depth is the number of I/O operations that can run in parallel on a device. If you are designing a configuration for a large SAN, you must estimate the queue depth for each node in order to avoid application failures.

QUESTION 59
Data which is accessed most often is stored on SSD, and data which is not accessed often is stored on SATA, What is this called?

A. Thin provisioning
B. Zoning
C. Data tiering
D. Capacity planning

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Data tiering monitors data usage in order to determine where data should be stored. Data tiering is a type of software program that moves data files, volumes, or blocks between tiered storage according to company-defined policy. Frequently accessed data will reside on high performance Fibre Channel or solid-state drives while infrequently accessed data gets moved to lower-cost, high-capacity drives in-house or cloud storage.

Answer option D is incorrect. Capacity planning is the process of determining the production capacity needed by an organization to meet changing demands for its products.

Answer option B is incorrect. Zoning allows an administrator to control who can see what in a SAN.

Answer option A is incorrect. Thin provisioning is the act of using virtualization technology to give the appearance of more physical resource than is actually available.

QUESTION 60
Which of the following is defined as "Queue Depth"?

A. Number of commands that a host bus adapter can send or receive
B. Number of I/O buses connected to SCSI devices
C. Number of ports connected to a fibre channel switch
D. Number of ISL links per switch

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Queue depth is defined as the number of commands that the HBA can send or receive in a single chunk per LUN. From a host-HBA point (initiator), it is defined as the number of commands that can be queued (or stored)

and then sent to storage. From a storage point (target), it is defined as the number of commands it can accept in .

Performance degradation starts quickly once user reaches the max queue depth on target, and it starts backing off the queues and it will hit hard on user's response times. Ideally, if a user has few hosts and luns, the default is good. If a user has the luxury to dedicate ports for high performance applications, ensure that the queue depth is configured appropriately so that the user gets the best of performance.