Unit – 4

Hybrid storage Solutions Virtualization

Unit 04/Lecture - 01

Storage virtualization – [Rgpv/dec2013(10)]

Virtualization is the technique of masking or abstracting physical resources, which simplifies the infrastructure and accommodates the increasing pace of business and technological changes. It increases the utilization and capability of it resources, such as servers, networks, or storage devices, beyond their physical limits. Virtualization simplifies resource management by pooling and sharing resources for maximum utilization and makes them appear as logical resources with enhanced capabilities.

Forms of virtualization

Virtualization has existed in the it industry for several years and in different forms, including memory virtualization, network virtualization, server virtualization, and storage virtualization.

1. Memory virtualization

Virtual memory makes an application appear as if it has its own contiguous logi- cal memory independent of the existing physical memory resources.

Since the beginning of the computer industry, memory has been and continues to be an expensive component of a host. It determines both the size and the number of applications that can run on a host.

With technological advancements, memory technology has changed and the cost of memory has decreased. Virtual memory managers have evolved, enabling multiple applications to be hosted and processed simultaneously.

In a virtual memory implementation, a memory address space is divided into contiguous blocks of fixed-size pages. A process known as paging saves inactive memory pages onto the disk and brings them back to physical memory when required. This enables efficient use of available physical memory among different processes. The space used by vmms on the disk is known as a swap file. A swap file (also known as page file or swap space) is a portion of the hard disk that functions like physical memory (ram) to the operating system. The operating system typically moves the least used data into the swap file so that ram will be available for processes that are more active. Because the space allocated to the swap file is on the hard disk (which is slower than the physical memory), access to this file is slower.

2.Network virtualization

Network virtualization creates virtual networks whereby each application sees its own logical network independent of the physical network. A virtual lan (vlan) is an example of network virtualization that provides an easy, flexible, and less expensive way to manage networks. Vlans make large networks more manageable by enabling a centralized configuration of devices located in physically diverse locations.

3. server virtualization

Server virtualization enables multiple operating systems and applications to run simultaneously on different virtual machines created on the same physical server (or group of servers). Virtual machines provide a layer of abstraction between the operating

System and the underlying hardware. Within a physical server, any number of virtual servers can be established; depending on hardware capabilities (see figure 10-1). Each virtual server seems like a physical machine to the operating system, although all virtual servers share the same underlying physical hardware in an isolated manner. For example, the physical memory is shared between virtual servers but the address space is not. Individual virtual servers can be restarted, upgraded, or even crashed, without affecting the other virtual servers on the same physical machine.

S.no	Rgpv question	Year	Marks
Q.1	What are various forms of vertulization? Explain each in brief?	Rgpv Dec 2011	10

Unit-04/Lecture -02

Virtual LAN(VLANs) - [Rgpv/dec 2011(10)]

In simple terms, a VLAN is a set of workstations within a LAN that can communicate with each other as though they were on a single, isolated LAN. What does it mean to say that they "communicate with each other as though they were on a single, isolated LAN"?
Among other things, it means that:

- Broadcast packets sent by one of the workstations will reach **all** the others in the VLAN
- .Broadcasts sent by one of the workstations in the VLAN will not reach any workstations that are not in the VLAN
- Broadcasts sent by workstations that are not in the VLAN will never reach workstations that are in the VLAN

- The workstations can all communicate with each other without needing to go through a gateway. For example, IP connections would be established by ARPing for the destination IP and sending packets directly to the destination workstation—there would be no need to send packets to the IP gateway to be forwarded on.

- The workstations can communicate with each other using non-routable protocols.

A Local Area Network (LAN) was originally defined as a network of computers located within the same area. Today, Local Area Networks are defined as a single broadcast domain. This means that if a user broadcasts information on his/her LAN, the broadcast will be received by every other user on the LAN. Broadcasts are prevented from leaving a LAN by using a router. The disadvantage of this method is routers usually take more time to process incoming data compared to a bridge or a switch. More importantly, the formation of broadcast domains depends on the physical connection of the devices in the network. Virtual Local Area Networks (VLAN's) were developed as an alternative solution to using routers to contain broadcast traffic.

In a traditional LAN, workstations are connected to each other by means of a hub or a repeater. These devices propagate any incoming data throughout the network. However, if two people attempt to send information at the same time, a collision will occur and all the transmitted data will be lost. Once the collision has occurred, it will continue to be propagated throughout the network by hubs and repeaters. The original information will therefore need to be resent after waiting for the collision to be resolved, thereby incurring a significant wastage of time and resources. To prevent collisions from traveling through all the workstations in the network, a bridge or a switch can be used. These devices will not forward collisions, but will allow broadcasts (to every user in the network) and multicasts (to a pre-specified group of users) to pass through. A router may be used to prevent broadcasts and multicasts from traveling through the network.

The workstations, hubs, and repeaters together form a LAN segment. A LAN segment is also known as a collision domain since collisions remain within the segment. The area within which broadcasts and multicasts are confined is called a broadcast domain or LAN. Thus a LAN can consist of one or more LAN segments. Defining broadcast and collision domains in a LAN depends on how the workstations, hubs, switches, and routers are physically connected together. This means that everyone on a LAN must be located in the same area.

Types of Connections

Devices on a VLAN can be connected in three ways based on whether the connected devices are VLAN-aware or VLAN-unaware. Recall that a VLAN-aware device is one which understands VLAN memberships (i.e. which users belong to a VLAN) and VLAN formats.

1) Trunk Link

All the devices connected to a trunk link, including workstations, must be VLAN-aware.

All frames on a trunk link must have a special header attached. These special frames are called tagged frames .

2) Access Link

An access link connects a VLAN-unaware device to the port of a VLAN-aware bridge. All frames on access links must be implicitly tagged (untagged) (see Figure8). The VLAN-unaware device can be a LAN segment with VLAN-unaware workstations or it can be a number of LAN segments containing VLAN-unaware devices (legacy LAN).

3) Hybrid Link

This is a combination of the previous two links. This is a link where both VLAN-aware and VLAN-unaware devices are attached (see Figure9). A hybrid link can have both tagged and untagged frames, but all the frames for a specific VLAN must be either tagged or untagged.

Advantages

- **Performance**. As mentioned above, routers that forward data in software become a bottleneck as LAN data rates increase. Doing away with the routers removes this bottleneck.

- **Formation of virtual workgroups**. Because workstations can be moved from one VLAN to another just by changing the configuration on switches, it is relatively easy to put all the people working together on a particular project all into a single VLAN. They can then more easily share files and resources with each other.

To be honest, though, virtual workgroups sound like a good idea in theory, but often do not work well in practice. It turns out that users are usually more interested in accessing company-wide resources (file servers, printers, etc.) than files on each others' PCs.

- **Greater flexibility**. If users move their desks, or just move around the place with their laptops, then, if the VLANs are set up the right way, they can plug their PC in at the new location, and still be within the same VLAN. This is much harder when a network is physically divided up by routers.

- **Ease of partitioning off resources**. If there are servers or other equipment to which the network administrator wishes to limit access, then they can be put off into their own VLAN. Then users in other VLANs can be given access selectively.

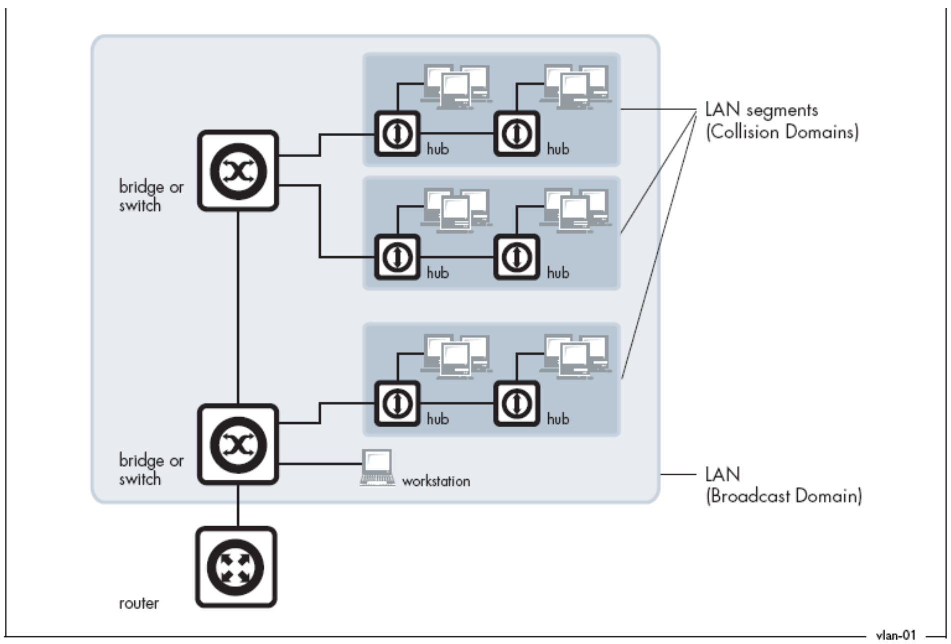

S.no	Rgpv question	Year	Marks
Q.1	What do you mean by VLANs?	Rgpv Dec 2011	10

Unit-04/Lecture -03

Management matrix – [Rgpv/dec 2012(10),rgpv dec2011(10)]

Definition: A style of management where an individual has two reporting superiors (bosses) - one functional and one operational.

matrix management is the practice of managing individuals with more than one reporting line (in a matrix organization structure), but it is also commonly used to describe managing cross functional, cross business group and other forms of working that cross the traditional vertical business units. It is a type of organizational management in which people with similar skills are pooled for work assignments, resulting in more than one manager (sometimes referred to as solid line and dotted line reports, in reference to traditional business organization charts).

Management advantages and disadvantages

Key advantages that organizations seek when introducing a matrix include:

- To break business information silos - to increase cooperation and communication across the traditional silos and unlock resources and talent that are currently inaccessible to the rest of the organization.
- To deliver work across the business more effectively – to serve global customers, manage supply chains that extend outside the organization, and run integrated business regions, functions and processes.
- To be able to respond more flexibly – to reflect the importance of both the global and the local, the business and the function in the structure, and to respond quickly to changes in markets and priorities.
- To develop broader people capabilities – a matrix helps develop individuals with broader perspectives and skills who can deliver value across the business and manage in a more complex and interconnected environment.

Key disadvantages of matrix organizations include:

- Mid-level management having multiple supervisors can be confusing, in that competing agendas and emphases can pull employees in different directions, which can lower productivity.
- Mid-level management can become frustrated with what appears to be a lack of clarity with priorities.
- Mid-level management can become over-burdened with the diffusion of priorities.
- Supervisory management can find it more difficult to achieve results within their area of expertise with subordinate staff being pulled in different directions.

Application

The advantages of a matrix for project management can include:

- Individuals can be chosen according to the needs of the project.
- The use of a project team that is dynamic and able to view problems in a different way as specialists have been brought together in a new environment.
- Project managers are directly responsible for completing the project within a specific deadline and budget.

The disadvantages for project management can include:

- A conflict of loyalty between line managers and project managers over the allocation of resources.
- Projects can be difficult to monitor if teams have a lot of independence.
- Costs can be increased if more managers (i.e. project managers) are created through the use of project teams.

S.no	Rgpv question	Year	Marks
Q.1	What do you understand by managent metrix? Explain.	Rgpv Dec 2012 Rgpv Dec 2012	10 10

Unit-04/Lecture -04

Data center infrastructure - [Rgpv/dec2013(7), Rgpv/dec2012(10), Rgpv/dec2012(10)]

Organizations maintain data centers to provide centralized data processing capabilities across the enterprise. Data centers store and manage large amounts of mission-critical data. The data center infrastructure includes computers, storage systems, network devices, dedicated power backups, and environmental controls (such as air conditioning and fire suppression).
Large organizations often maintain more than one data center to distribute data processing workloads and provide backups in the event of a disaster. The storage requirements of a data center are met by a combination of various stor- age architectures.

Core elements

Five core elements are essential for the basic functionality of a data center:

- Application: an application is a computer program that provides the logic for computing operations. Applications, such as an order processing system.
- Database: more commonly, a database management system (dbms) provides a structured way to store data in logically organized tables that are interrelated. A dbms optimizes the storage and retrieval of data.
- Server and operating system: a computing platform that runs applica- tions and databases.
- Network: a data path that facilitates communication between clients and servers or between servers and storage.
- Storage array: a device that stores data persistently for subsequent use.

These core elements are typically viewed and managed as separate entities, but all the elements must work together to address data processing requirements.

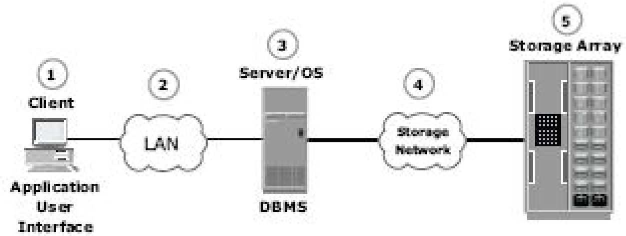

key requirements for datacenter elements

Uninterrupted operation of data centers is critical to the survival and success of a business. It is necessary to have a reliable infrastructure that ensures data is accessible at all times. While the requirements, shown in figure 1-6, are appli- cable to all elements of the data center infrastructure, our focus here is on storage systems.

Key characteristics of data center elements

- Availability: all data center elements should be designed to ensure accessibility. The inability of users to access data can have a significant negative impact on a business.
- Security: polices, procedures, and proper integration of the data center core elements that will prevent unauthorized access to information must be established. In addition to the security measures for client access, specific mechanisms must enable servers to access only their allocated resources on storage arrays.
- Scalability: data center operations should be able to allocate additional processing capabilities or storage on demand, without interrupting business operations. Business growth often requires deploying more servers, new applications, and additional

databases. The storage solution should be able to grow with the business.

- Performance: all the core elements of the data center should be able to provide optimal performance and service all processing requests at high speed. The infrastructure should be able to support performance requirements.
- Data integrity: data integrity refers to mechanisms such as error correction codes or parity bits which ensure that data is written to disk exactly as it was received. Any variation in data during its retrieval implies corruption, which may affect the operations of the organization.
- Capacity: data center operations require adequate resources to store and process large amounts of data efficiently. When capacity requirements increase, the data center must be able to provide additional capacity with- out interrupting availability, or, at the very least, with minimal disruption. Capacity may be managed by reallocation of existing resources, rather than by adding new resources.
- Manageability: a data center should perform all operations and activities in the most efficient manner. Manageability can be achieved through automation and the reduction of human (manual) intervention in com- mon tasks.

Managing storage infrastructure

Managing a modern, complex data center involves many tasks. Key management activities include:

- Monitoring is the continuous collection of information and the review of the entire data center infrastructure. The aspects of a data center that are monitored include security, performance, accessibility, and capacity.
- Reporting is done periodically on resource performance, capacity, and utilization. Reporting tasks help to establish business justifications and chargeback of costs associated with data center operations.
- provisioning is the process of providing the hardware, software, and other resources needed to run a data center. Provisioning activities include capacity and resource

planning. Capacity planning ensures that the user's and the application's future needs will be addressed in the most cost-effective and controlled manner. Resource planning is the process of evaluating and identifying required resources, such as personnel, the facility (site), and the technology. Resource planning ensures that adequate resources are available to meet user and application requirements.

S.no	Rgpv question	Year	Marks
Q.1	What are the data centre? What are the requirement for the design of a secure data centre.	Rgpv Dec2013 Rgpv Dec2012 Rgpv Dec2012	7 10 10

Unit-04/Lecture -05

Backup & disaster recovery – [Rgpv/dec 2013(10), Rgpv/dec 2012(7),Rgpv/dec 2012(10)]

Backup is a copy of production data, created and retained for the sole purpose of recovering deleted or corrupted data.
With growing business and regulatory demands for data storage, retention, and availability, organizations are faced with the task of backing up an ever-increasing amount of data. This task becomes more challenging as demand for consistent backup and quick restore of data increases th roughout the enter prise wh ich may be spread over multiple sites. Moreover, organizations need to accomplish backup at a lower cost with minimum resources.

Backup purpose

Backups are performed to serve three purposes: disaster recovery, operational backup, and archival.

Disaster recovery

Backups can be performed to address disaster recovery needs. The backup copies are used for restoring data at an alternate site when the primary site is incapacitated due to a disaster. Based on rpo and rto requirements, organizations use different backup strategies for disaster recovery. When a tape-based backup method is used as a disaster recovery strategy, the backup tape media is shipped and stored at an offsite location. These tapes can be recalled for restoration at the disaster recovery site. Organizations with stringent rpo and rto requirements use remote replication technology to replicate data to a disaster recovery site. This allows organizations to bring up production systems online in a relatively short period of time in the event of a disaster.

Operational backup

Data in the production environment changes with every business transaction and operation. Operational backup is a backup of data at a point in time and is used to restore data in the event of data loss or logical corruptions that may occur during routine processing. The majority of restore requests in most organizations fall in this category. For example, it is common for a user to accidentally delete an important e-mail or for a file to become corrupted, which can be restored from operational backup.

Operational backups are created for the active production information by using incremental or differential backup techniques, detailed later in this chapter. An example of an operational backup is a backup performed for a production data- base just before a bulk batch update. This ensures the availability of a clean copy of the production database if the batch update corrupts the production database.

Archival

Backups are also performed to address archival requirements. Although CAS has emerged as the primary solution for archives, traditional backups are still used by small and medium enterprises for long-term preservation of transaction

Records, e-mail messages, and other business records required for regulatory compliance.

Apart from addressing disaster recovery, archival, and operational require- ments, backups serve as a protection against data loss due to physical damage of a storage device, software failures, or virus attacks. Backups can also be used to protect against accidents such as a deletion or intentional data destruction.

Backup methods

Hot backup and cold backup are the two methods deployed for backup. They are based on the state of the application when the backup is performed. In a hot backup, the application is up and running, with users accessing their data during the backup process. In a cold backup, the application is not active during the backup process.

The backup of online production data becomes more challenging because data is actively being used and changed. An open file is locked by the operating system and is not copied during the backup process until the user closes it. The backup application can back up open files by retrying the operation on files that were opened earlier in the backup process. During the backup process, it may be possible that files opened earlier will be closed and a retry will be successful. The maximum number of retries can be configured depending on the backup application. However, this method is not considered robust because in some environments certain files are always open.

In such situations, the backup application provides open file agents. These agents interact directly with the operating system and enable the creation of consistent copies of open files. In some environments, the use of open file agents is not enough. For example, a database is composed of many files of varying sizes, occupying several file systems. To ensure a consistent database backup, all files need to be backed up in the same state. That does not necessarily mean that all files need to be backed up at the same time, but they all must be syn- chronized so that the database can be restored with consistency.

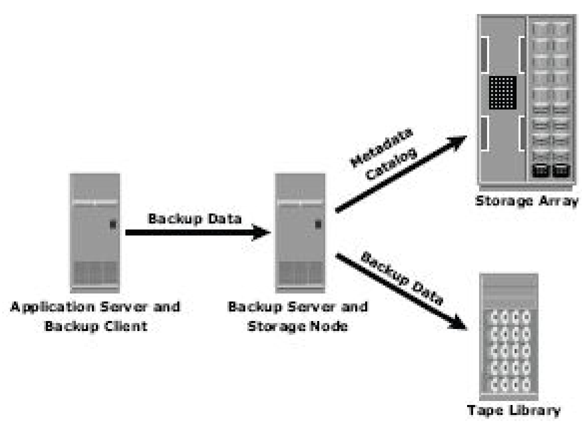

Backup architecture and process

The storage node is responsible for writing data to the backup device (in a backup environment, a storage node is a host that controls backup devices). Typically, the storage node is integrated with the backup server and both are hosted on the same physical platform. A backup device is attached directly to the storage node's host platform. Some backup architecture refers to the storage node as the media server because it connects to the storage device. Storage nodes play an important role in backup planning because they can be used to consolidate backup servers.

Backup software also provides extensive reporting capabilities based on the backup catalog and the log files. These reports can include information such as the amount of data backed up, the number of completed backups, the number of incomplete backups, and the types of errors that may have occurred. Reports can be customized depending on the specific backup software used.

S.no	Rgpv question	Year	Marks
Q.1	Explain briefly the following a) Disaster b) Operational backup c) Archival	Rgpv Dec 2013	7
Q.2	What do you understand by backup and disaster recovery? Explain.	Rgpv Dec 2013	10

Unit 04/Lecture - 06

Backup topologies – [Rgpv/2013(7)]

Three basic topologies are used in a backup environment: direct attached backup, lan based backup, and SAN based backup. A mixed topology is also used by combining lan based and SAN based topologies.

In a direct-attached backup, a backup device is attached directly to the client. Only the metadata is sent to the backup server through the lan. This configuration frees the lan from backup traffic.

Depicts use of a backup device that is not shared. As the environment grows, however, there will be a need for central management of all backup devices and to share the resources to optimize costs. An appropriate solution is to share the backup devices among multiple servers. In this example, the client also acts as a storage node that writes data on the backup device.

In lan-based backup, all servers are connected to the lan and all storage devices are directly attached to the storage node. The data to be backed up is transferred from the backup client (source), to the backup device (destination) over the lan, which may affect network performance. Streaming across the lan also affects network performance of all systems connected to the same segment as the backup server. Network resources are severely constrained when multiple clients access and share the same tape library unit (tlu).

This impact can be minimized by adopting a number of measures, such as configuring separate networks for backup and installing dedicated storage nodes for some application servers.

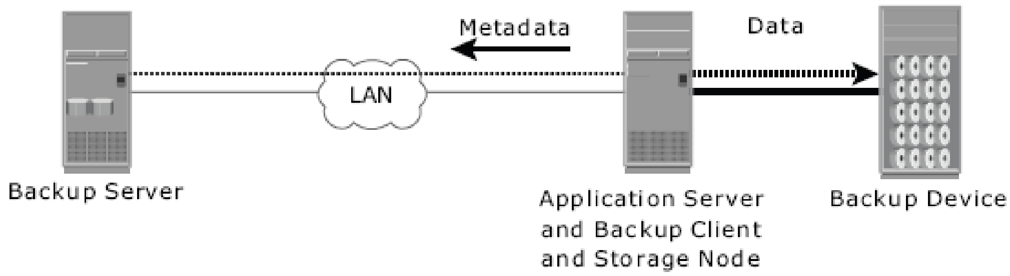

Lan-based backup topology

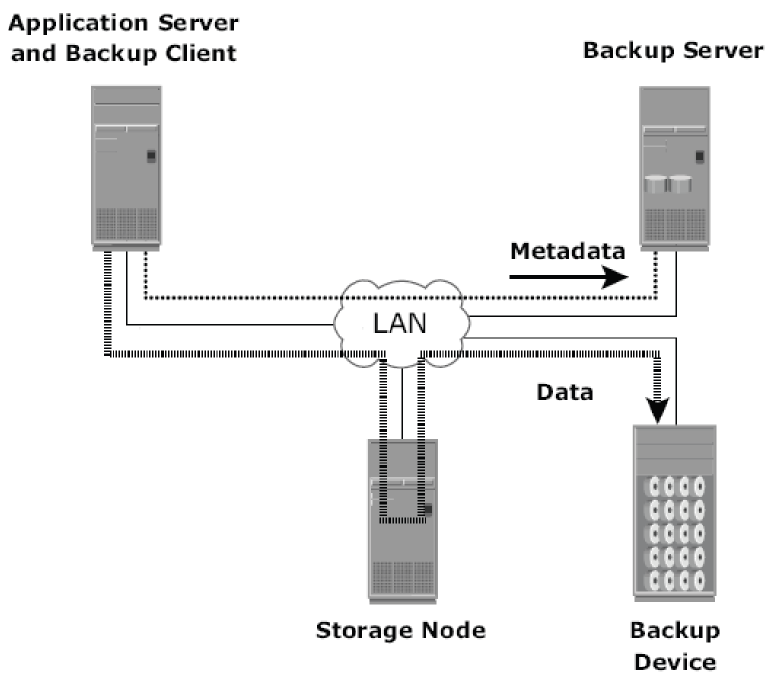

SAN-based backup topology

The mixed topology uses both the lan-based and SAN-based topologies, as shown in figure. This topology might be implemented for several reasons, including cost, server location, reduction in administrative overhead, and performance considerations.

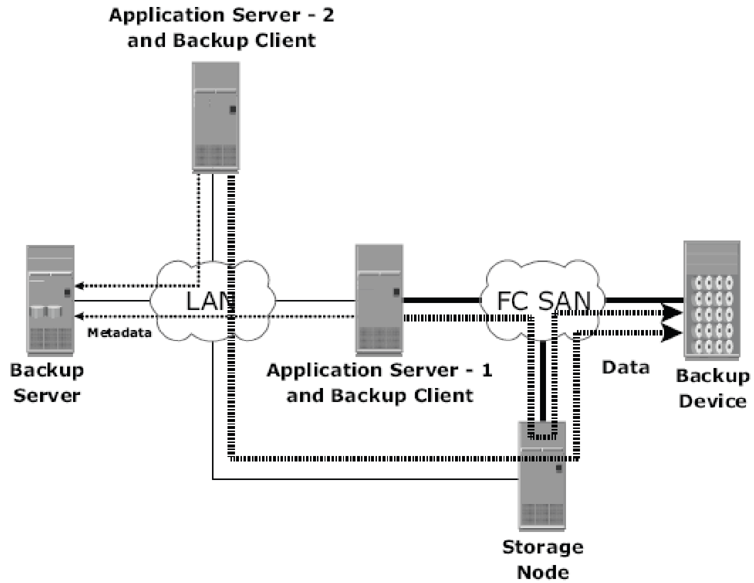

Mixed backup topology

Serverless backup

Serverless backup is a lan-free backup methodology that does not involve a backup server to copy data. The copy may be created by a network-attached controller, utilizing a scsi extended copy or an appliance within the SAN. These backups are called serverless because they use SAN resources instead of host resources to transport backup data from its source to the backup device, reducing the impact on the application server.

Another widely used method for performing serverless backup is to lever- age local and remote replication technologies. In this case, a consistent copy of the production data is replicated within the same array or the remote array, which can be moved to the backup device through the use of a storage node.

S.no	Rgpv question	Year	Marks
Q.1	List and explain different topologies for backup?	Rgpv Dec 2013	7

Unit 04/Lecture -07

SNIA storage virtualization taxonomy - [Rgpv/dec2013(7)]

The snia (storage networking industry association) storage virtualization taxon- omy provides a systematic classification of storage virtualization, with three levels defining what, where, and how storage can be virtualized.

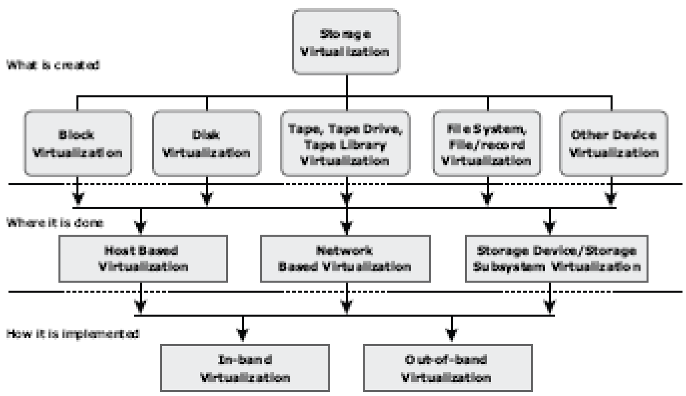

Snia storage virtualization taxonomy

The first level of the storage virtualization taxonomy addresses "what" is created. It specifies the types of virtualization: block virtualization, file virtual- ization, disk virtualization, tape virtualization, or any other device virtualiza- tion.

The second level describes "where" the virtualization can take place. This requires a multilevel approach that characterizes virtualization at all three levels of the storage environment: server, storage network, and storage, as shown in . An effective virtualization strategy distributes the intelligence across all three levels while centralizing the management and control functions. Data storage functions—such as raid, caching, checksums, and hard- ware scanning—should remain on the array. Similarly, the host should control application-focused areas, such as clustering and application failover, and vol- ume management of raw disks. However, path redirection, path failover, data access, and distribution or load-balancing capabilities should be moved to the switch or the network.

S.no	Rgpv question	Year	Marks
Q.1	Explain SNIA storage virtualization taxonomy?	Rgpv Dec 2013	7

Unit - 04/Lecture-08

Types of storage virtualization- [Rgpv/2013(7)]

Virtual storage is about providing logical storage to hosts and applications independent of physical resources. Virtualization can be implemented in both SAN and NAS storage environments. In a SAN, virtualization is applied at the block level, whereas in NAS, it is applied at the file level.

Block-level storage virtualization

Block-level storage virtualization provides a translation layer in the SAN, between the hosts and the storage arrays. Instead of being directed to the luns on the individual storage arrays, the hosts are directed to the virtualized luns on the virtualization device. The virtualization device translates between the virtual luns and the physical luns on the individual arrays. This facilitates the use of arrays from different vendors simultaneously, without any interoperability issues. For a host, all the arrays appear like a single target device and luns can be distributed or even split across multiple arrays.

Block-level storage virtualization extends storage volumes online, resolves application growth requirements, consolidates heterogeneous storage arrays, and enables transparent volume access. It also provides the advantage of non- disruptive data migration.

In traditional SAN environments, lun migration from one array to another was an offline event because the hosts needed to be updated to reflect the new array configuration. In other instances, host cpu cycles were required to migrate data from one array to the other, especially in a multi vendor environment. With a block-level virtualization solution in place, the virtualization engine handles the back-end migration of data, which enables luns.

To remain online and accessible while data is being migrated. No physical changes are required because the host still points to the same virtual targets on the virtualization device. However, the mappings on the virtualization device should be changed. These changes can be executed dynamically and are transparent to the end user.

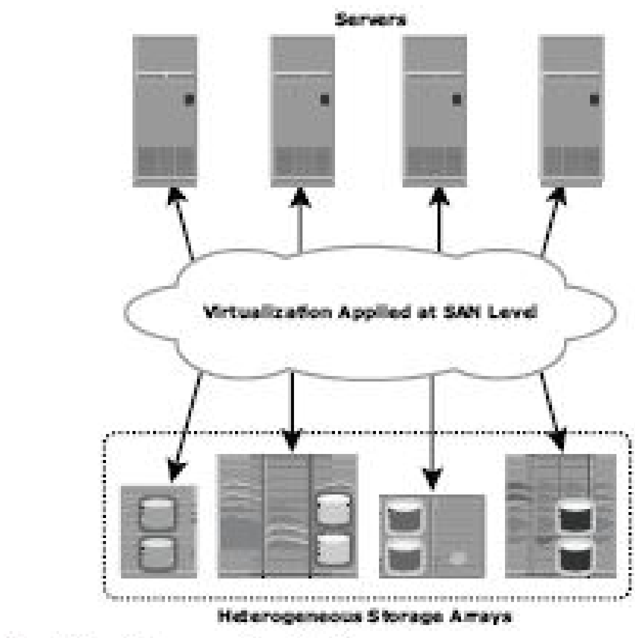

Block-level storage virtualization

Deploying heterogeneous arrays in a virtualized environment facilitates an information lifecycle management (ILM) strategy, enabling significant cost and resource optimization. Low-value data can be migrated from high- to low-per- formance arrays or disks.

file-level virtualization

File-level virtualization addresses the NAS challenges by eliminating the dependencies between the data accessed at the file level and the location where the files are physically stored. This provides opportunities to optimize storage utilization and server consolidation and to perform nondisruptive file migrations

S.no	Rgpv question	Year	Marks
Q.1	Explain block level storage vertulization?	Rgpv Dec 2013	7

Additional Topic Unit -04/Lecture -09

Managing & monitoring-

Snmp- the *snmp* protocol was the standard used to manage multi-vendor san environments. However, snmp was primarily a network management protocol and was inadequate for providing the detailed information and functionality required to manage the san environment. The unavailability of automatic discovery functions, weak modeling constructs, and lack of transactional sup- port are some inadequacies of snmp in a san environment. Even with these limitations, snmp still holds a predominant role in san management, although newer open storage san management standards have emerged to monitor and manage these environments more effectively.

Storage management initiative

The storage networking industry association (snia) has been engaged in an initiative to develop a common, open storage, and san management interface. Smi-s is based on web-based enterprise management (wbem) technology and the dmtf's common information model (cim). The initiative was formally created to enable broad interoperability among heterogeneous storage vendor systems and to enable better management solutions that span these environ- ments. This initiative is known as the storage management initiative (smi).

The smi specification, known as smi-s, offers substantial benefits to users and vendors. It forms a normalized, abstracted model to which a storage infra- structure's physical and logical components can be mapped, and which can be used by management applications such as storage resource management, device management, and data management for standardized, effective, end-to-end control of storage resources .

Using smi-s, the storage software developers have a single normalized and unified object model comprising the detailed document that contains information about managing the breadth of san components. Moreover, smi-s eliminates the need for development of vendor-proprietary management interfaces, enabling vendors to focus on added value functions and offering solutions in a way that will support new devices as long as they adhere to the standard. Using smi-s, device vendors can build new features and functions to manage storage subsystems and expose them via smi-s. The smi-s-compliant products lead to easier, faster deployment, and accelerated adoption of policy-based storage management frameworks.

The information required to perform management tasks is better organized or structured in a way that enables disparate groups of people to use it. This can be accomplished by developing a model or representation of the details required by users working within a particular domain. Such an approach is referred to as an information model. An information model requires a set of legal statements or syntax to capture the representation and expressions necessary to manage common aspects of that domain.

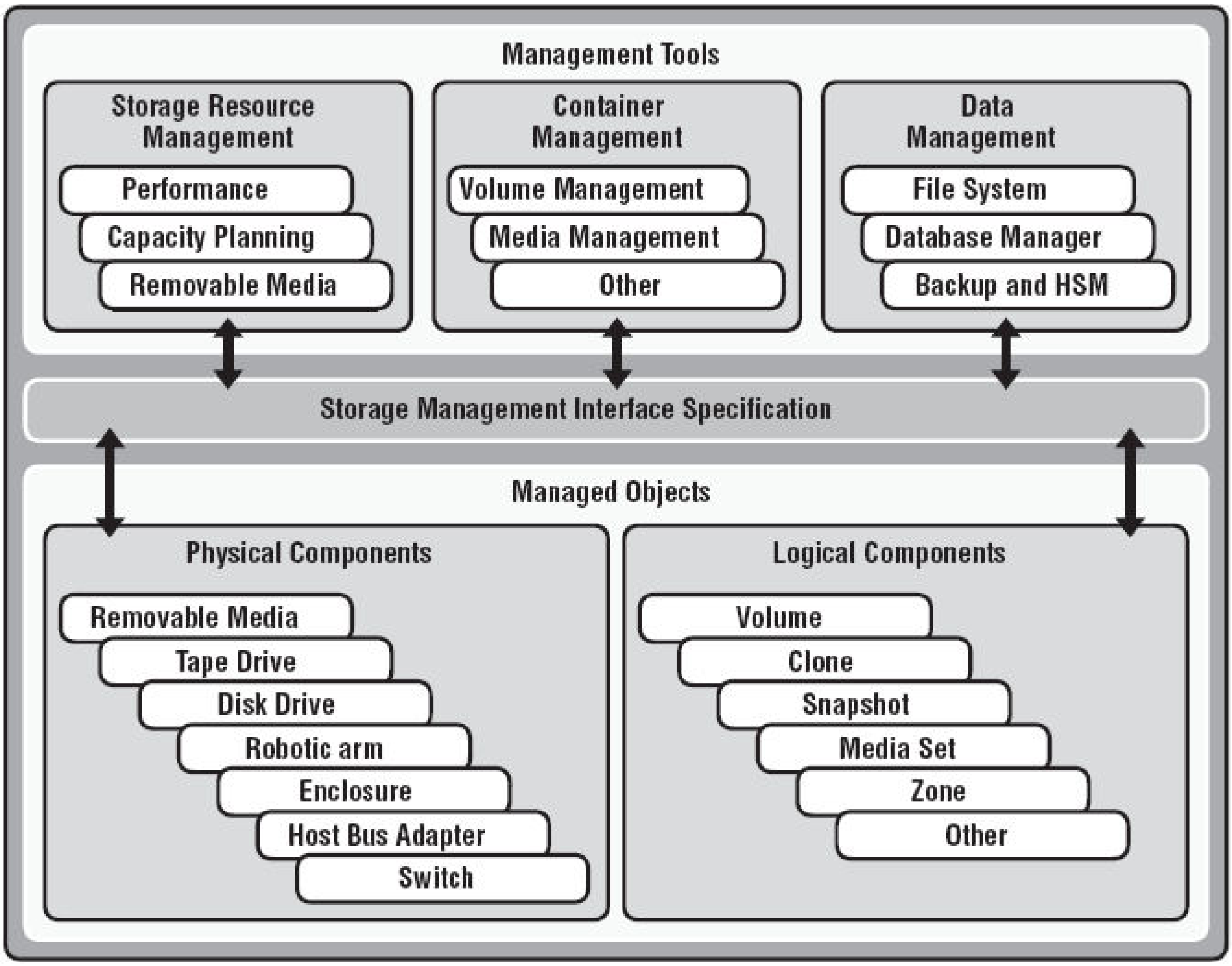

The cim is a language and methodology for describing management elements. A cim schema includes models for systems, applications, networks, and devices. This schema also enables applications from different vendors working on different platforms to describe the management data in a standard format so that it can be shared among a variety of management applications.

The following features of smi-s simplify san management:

Common data model: smi-s agents interact with an smi-s-enabled device, such as a switch, a server, or a storage array, to extract relevant management data. They can also interact at the management layer to exchange information between one management application and another. They then provide this information to the requester in a consistent syntax and format.

Interconnect independence: smi-s eliminates the need to redesign the management transport and enables components to be managed by using out-of-band communications. In addition, smi-s offers the advantages of specifying the cmi-xml over the http protocol stack and utilizing the lower layers of the tcp/ip stack, both of which are ubiquitous in today's networking world.

Multilayer management: smi-s can be used in a multilayer and cross- domain environment—for example, server-based volume managers and network storage appliances. Many storage deployment environments currently employ this combination.

legacy system accommodation: smi-s can be used to manage legacy systems by using a proxy agent or can be directly supported by the device itself. Smi-s can coexist with proprietary apis and agents as well as providing integration framework for such mechanisms.

Policy-based management: smi-s includes object models applicable across all classes of devices, enabling a san administrator to implement policy-based management for entire storage networks.

Reference

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