ISSN 2278-3091

International Journal of Advanced Trends in Computer Science and Engineering, Vol.2 , No.1, Pages : 437 - 441 (2013)
Special Issue of ICACSE 2013 - Held on 7-8 January, 2013 in Lords Institute of Engineering and Technology, Hyderabad

Concept of Storage Virtualization used in Networks

1 Mr Majeed Faaraan , 2 Maimoona Begum 3 Dr Giridhar
1 M Tech Student, LIET, India, majeedfaaraan@gmail.com
2 M Tech Student, LIET, India, maimunabegum@gmail.com
3. Principal, LIET, India

Abstract

Storage Virtualization was the first tutorial of the SNIA(Storage Networking Industrial Association) Education Committee. Development began in summer 2001 and continues to this day as we incorporate new material and address new developments in the storage virtualization segment. At the time of writing this booklet, the tutorial team consisted of about 35 members from different vendor companies as well as end users interested in this topic. These team members either monitor the activities and review the con- tent, or actively develop the material to reflect new technologies and changes to existing technologies.

This paper explains the concepts of virtualization and its uses in different areas of networking.

Keywords:

Introduction:

Storage virtualization as a technology is only useful if it offers new ways to solve problems. The most severe problem that virtualization can help solve today is the overall management of storage. Today, storage infrastructure represents one of the most heterogeneous environments found in modern IT departments, with a multitude of different systems at all levels of the stack—file systems, operating systems, servers, storage systems, management consoles, management software, etc.

Virtual Storage:

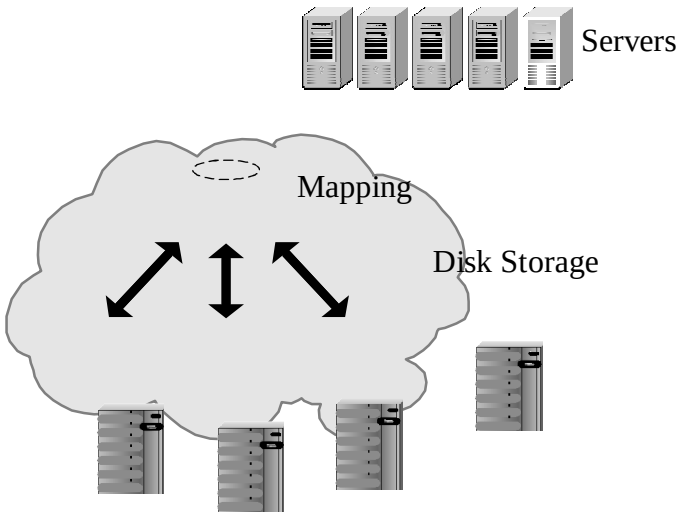

ISSN 2278-3091

International Journal of Advanced Trends in Computer Science and Engineering, Vol.2 , No.1, Pages : 437 - 441 (2013)
Special Issue of ICACSE 2013 - Held on 7-8 January, 2013 in Lords Institute of Engineering and Technology, Hyderabad

A basic model of virtual storage may be depicted by simply changing the physical devices (layer I) into virtual devices (layer II). The servers accessing these virtual devices do not "know" the difference—the interface to the virtual devices is identical to the interface to a physical device. Operations to the layer I devices such as adding, changing, or replacing components may be performed without disrupting the layers above.

Types of Virtualization:

Stack based Virtualization can be categorized as follows.

1. Host-based virtualization
2. Network-based virtualization
3. Storage-based virtualization.

Host-Based Virtualization

This type of virtualization is generally associated with logical volume managers, which, with varying levels of sophistication, are commonly found on just about every computer from the desktop to the data center. As is the case with storage- based virtualization, logical volume managers are not necessarily associated with SANs. Yet they are still the most popular method of virtualization because of their history and the fact that direct attached storage (DAS) is still very widespread. Increasingly, the logical volume manager (LVM) is a standard part of the operating system, but more advanced third-party implementations are also quite common. The most common uses of host-based LVMs are:

1. Aggregating physical storage from multiple LUNs to form a single "super- LUN" that the host OS sees as a single disk drive
2. Implementing software RAID and other more advanced functions, including snapshots and remote replication
3. Managing the health of disk resources that are under the control of the operating system.

Network-Based Virtualization:

A network-based approach to virtualization supports data center-wide storage management and is able to accommodate a truly heterogeneous SAN with a diverse range of host platforms and storage resources. Typically, network-based virtualization is implemented using "black-box" appliances in the SAN fabric and potentially some agent software installed on the host (the need for and the extent of those host-based agents depends on how virtualization is implemented in the appliance). The appliance itself can be anything from an off-the-shelf server plat- form to a dedicated, proprietary hardware design. Typical functions offered by network-based virtualization include:

- Combining several LUNs from one or more arrays into a single LUN before presenting it to a host
- Taking a single LUN from an array and slicing it into smaller virtual LUNs to present to the hosts
- Synchronous and asynchronous replication within the SAN as well as over WAN links

ISSN 2278-3091

International Journal of Advanced Trends in Computer Science and Engineering, Vol.2 , No.1, Pages : 437 - 441 (2013)
Special Issue of ICACSE 2013 - Held on 7-8 January, 2013 in Lords Institute of Engineering and Technology, Hyderabad

• Device security to ensure that access to a LUN is restricted to specified hosts Other functions include caching, advanced volume management, storage on demand, and QoS functions, but the availability of these more advanced features varies from one vendor to another.

Storage-Based Virtualization:

Storage-based virtualization is typically not dependent on a specific type of host, allowing the array to support heterogeneous hosts without worrying about variance of host operating systems or applications. Also, storage-based RAID systems deliver optimum performance in relation to their hardware because features like caching can be tuned to the specific hardware. The downside to this approach is that the storage virtualization functions are typically confined to a single array; for example, the source volume used for a snapshot and the snapshot itself are maintained on the same array, making the snapshot useless in case of hardware failure. In some cases, virtualization functions extend across multiple arrays, or a cluster of arrays or controllers; however, these solutions are typically restricted to a single-vendor implementation.

Implementation:

Implementation of virtualization concepts can be described on the following factors.

1 Reusability
2. Backup/restore of data
3.Implementing a virtualized SAN
4.Migrating data

Reusing :

SAN infrastructure. If the customer has existing infrastructure in the form of switches, hubs, HBAs, etc., are there any restrictions on its use in a virtualized SAN? The answer is that the solutions on the market today should be able to work with existing SAN infrastructure.

Direct attached storage (DAS). With many of the in-band solutions available today, DAS resources such as SCSI-attached arrays can be migrated into the SAN simply by attaching them directly to the virtualization engine using a suitable HBA. The virtualization engine handles the FC-SCSI (or FC-SSA, etc.) bridging within the virtualization engine. For an out-of-band virtualization system, it will be necessary to use stand-alone FC-SCSI bridges to per- form this function.

ISSN 2278-3091

International Journal of Advanced Trends in Computer Science and Engineering, Vol.2 , No.1, Pages : 437 - 441 (2013)
Special Issue of ICACSE 2013 - Held on 7-8 January, 2013 in Lords Institute of Engineering and Technology, Hyderabad

Back Up:

A backup is always recommended when performing a migration to a virtualized SAN. It may not be necessary to perform a restore of the data. For out-of-band virtualizers, the data on the disks should not be touched at all by the virtualization process. For in-band virtualizers, the answer is implementation dependent; most of the solutions on the market today offer the ability to take an existing LUN and simply layer the virtualization services onto the LUN without impacting the partition table or file systems on the LUN. However, using in-band virtualization with existing LUNs will prevent the virtualizer from offering a full set of advanced functions for that LUN, so it's worth discussing this issue with the vendor(s) in question.

Another option to consider is the use of host-based mirroring to mirror an existing LUN to a fully virtualized LUN as part of the implementation process.

Virtualized SAN:

The following points are focused more in the implementation of virtualized SAN.

. 1The object of the exercise is to move existing hosts to a Fibre Channel-based, virtualized SAN.

2. The hosts all have direct attached storage in which the data must be migrated to the SAN with a minimum of downtime and disruption.

3. All the hosts have some form of host-based mirroring available, either built in to the operating system or from a third-party application installed on the host

Migrating Data

Once the application servers are connected to the SAN and the virtualization engine is in place, the next step is to get a copy of the data on the direct attached storage into the SAN. There are several steps in this process:

1. A LUN equal in size to the locally attached storage is served up to the application server.
2. Using the host-based mirroring functions on the host, a mirror is created between the local storage and the SAN storage. At this point, there will be some performance degradation while the mirror is created, but scheduling the operation during off-peak hours can minimize the impact

ISSN 2278-3091

International Journal of Advanced Trends in Computer Science and Engineering, Vol.2 , No.1, Pages : 437 - 441 (2013)
Special Issue of ICACSE 2013 - Held on 7-8 January, 2013 in Lords Institute of Engineering and Technology, Hyderabad

Conclusion

Main aim of this paper is to focus on the basic concepts of virtualization and its implementation. With the discussions made in this paper it is concluded that, storage virtualization will be used to obtain the below mentioned advantages.

- Improved Storage utilization
- Reduced downtime
- Relief for administrators
- Maintaining back up of data
- Data migaration

References:

[1] www.snia.org/join
[2] www.snia.org/members/tutorials/virtualization
[3] www.snia.org/tech_activities/shared_storage_model
[4] www.snia.org/apps/org/workgroup/snia-edcomm/tut-virtualization