www.ijecs.in
International Journal Of Engineering And Computer Science ISSN:2319-7242
Volume 3 Issue 3 March,2014 Page No. 5088-5092

Performance Evaluation of iSCSI based storage network with cluster mirroring.

Rahul Jangid[1], Mr. Rajendra Purohit[2]

[1]Jodhpur Institute of Engineering and Technology, Jodhpur, India.
Jangid.rahul61@gmail.com

[2]Asso. Prof. (C.S.E/I.T) Jodhpur Institute of Engineering and Technology, Jodhpur, India.

Abstract: With an increase in the amount of volume of data in networks, the traditional storage architecture is greatly challenged. Today, SAN (storage area network) emerges as the main technology of network storage. The main inter networking technology of SAN is the fibre channel. But now with increasing demand for mirroring solutions for small & medium sized businesses, the high implementation costs, lack of trained staff and poor interoperability of fibre channel have proposed iSCSI to be a cost effective and viable alternative of fibre channel. Mirroring provides excellent solution for disaster-recovery, high data availability and business continuity. Synchronous mirroring or asynchronous mirroring enables customers to have an instantaneous recovery and eliminate loss of access to storage during disasters. But the cluster mirroring can transform your existing server storage into high availability storage system. The cluster mirroring tries to reduce the response-time and provide high availability storage. In this paper, we explore the details of iSCSI and attempt to evaluate the performance of iSCSI SAN and show that it can be performed satisfactorily and economically, without requiring the costlier fibre channel option.

Keywords: **SAN, iSCSI, SCVM, IP SAN.**

I. Introduction

With the enormous amount of data being exchanged daily there is a definite need for storage repositories and the ability to retrieve data with a minimum amount of time delay. In real life application, more organizations are looking at the Storage area network (SANs) to address the mid to long term storage need of businesses from small and medium business (SMBs) to the large Enterprise. The SCVM (Storage Concentrator Virtual Machine) creates IP SAN storage as virtual machine using iSCSI [1, 2], providing the user with the ability to consolidate their virtual data center. The storage concentrators consolidate physical storage resources, and provision those resources into logical volumes. With the Virtual machine, the separation of physical storage from logical volumes offers increased flexibility in storage management and fault tolerance. The Storage Resources can be located internally, externally or on a network. SCVM s can be used to create an IP SAN [1, 2] within a virtual environment that eliminates the need for a separate IP SAN hardware appliance. It can also be used as a target for local replication and remote replication from IP SAN to SCVMs, as a primary or secondary IP SAN storage,

Backend storage for an achieve appliance, security appliance, as data flexible virtual appliance. Remote mirroring enables creating exact copies of the primary data at a site far away from the primary site, but connected to it via some kind of an underlying network. The mirroring process is invisible to the application servers and so does not affect the allocated resources for the application in any way.

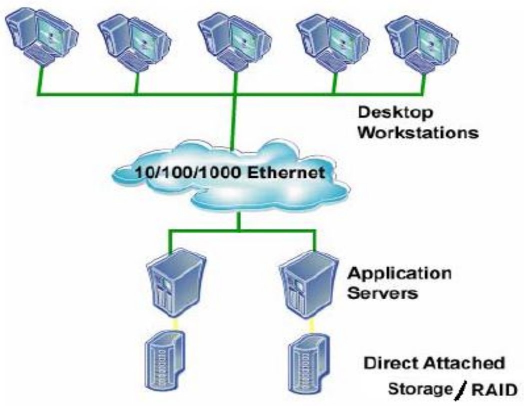

Figure 1: IP SAN using DAS / RAID

II. Storage Area Networking (SAN)

Storage area network is a storage network specifically designed to interconnect high speed storage devices and servers, as show in Figure 2. This specific network can span a large number of systems and geographic locations. Unlike in NAS, SAN [7] is connected behind the servers where storage devices and clients/servers are on same local segment. SAN provides a block-level interface to clients as opposed to the file-level interface in NAS, which provides a faster mechanism to transfer data over the network, i.e., reading and writing of data

occurs at block level or the transfer of raw data at physical level. An example of how data is transferred in actual SAN can be seen in Figure 2.
When a client on a network requests information from a storage device, the request is received by server, which in turn obtains the data from the storage device and sends it to the appropriate client over the network. Hence, the server has full control over the storage device and transfers the appropriate data requested by the client in the minimum possible amount of time. This particular block-level transfer provides a faster level of data transmission than to file-level transfers in NAS systems, since there is no overhead involved. Also, backups in the SAN network do not affect the rest of the LAN, since the back-up data passes only over the SAN, providing a LAN free back-up environment and less chances of network congestion occurred in frequent backups of data over the LAN. Hence for providing faster transmission of data over SAN, devices like SCSI are used to transfer block of data unlike other devices that are used for storage.

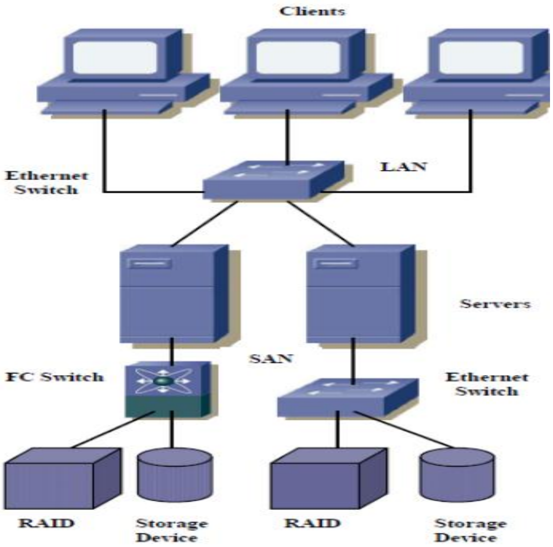

Figure 2: Storage Area Network (SAN)

III. Internet Protocol Storage Area Network (IP SAN)

A common side effect of server consolidation projects or implementing a centralized Fibre Channel SAN is scores of unused disk arrays. Many organizations are reluctant to perform what they call forklift upgrades. This study explains way to migrate captured server disks (direct-attached external or RAID) to an IP SAN (Internet Protocol Storage Area Network) as [2, 12, 14]. To access the data on a given DAS/RAID, the desktop workstations must communicate with the application server attached to the storage volume using iSCSI command sets. The storage on an IP SAN with many application servers, place a dedicated 10/100/1000 Ethernet switch [10] behind the application servers. A storage volume sits on this dedicated 10/100/1000 Ethernet IP SAN. Each application server can use generic 10/100/1000 NIC cards with iSCSI software and/or dedicated SNIC hardware, as shown in the diagram below. In this environment, users can access stored data on any of the DAS/RAID via any of the other servers, including application servers. You can use logical volumes to segregate data.

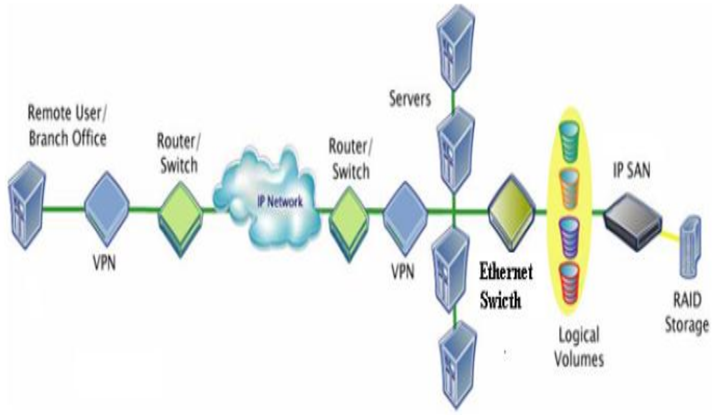

Figure 3: Internet Protocol Storage Area Networking (IP SAN)

IV. Cluster Mirroring

The cluster mirroring [8, 12] implementation also makes use of the ability to have multiple logging implementations. The cluster log implementation is a fault-tolerant client/server logging system. It requires shared storage (which should be available, given the fact that the mirror volume is meant to be clustered). The mirroring system makes calls to the log, which are passed on to the server. The server then responds to the client, which passes the results back to the system. This may cause network traffic (which is inevitable), but it reduces disk traffic. The server is smart enough to know that if one machine has marked a region dirty, it need not mark it again for another machine. Other optimizations have been made as well. In this mirroring, master IP address ensures server access to storage volumes no matter which unit is active, always contains a group of volumes. Cluster mirroring technology to eliminate all points of failure across the management, storage and hardware engines [8, 9]. Cluster-Mirrors can be used for improving the uptime and availability of SAN.

For higher levels of availability, mirrored images maintained in Active/Active Cluster mirroring can provide a system with No Single Points of Failure, which is designed to improve the overall uptime of the storage system for organizations with 7x24x365 requirements. The user can protect from server hardware failure and reduce the planned or unplanned downtime. Cluster mirroring provides all of the advantages of storage consolidation in addition to clustering advantages. These advantages include centralized storage management, better storage utilization and lower storage management costs.
No single point of failure operation is enabled by redundant connections to the SAN for the host systems, redundant SAN with redundant drive arrays and cabling, and data mirrored by the Intelligent Network System. In the event of a failure of any redundant component, the system will remain operational and will continue to deliver data to the attached servers. When the failed components have been repaired, the system will automatically resume fault tolerant operation and will be able to sustain another failure. If, because of the failure, the data in a mirror volume becomes out of sync, it will automatically be resynchronized when the failed component is restored to normal operation.

1. Fail-Over

A Fail-Over [8, 9] is a high redundancy storage provisioning solution comprised of two storage volumes configured so that one is actively managing storage volumes and the other is in standby mode monitoring all fault condition If the active storage volumes fails for any reasons, including loss of power, failure of a Gigabit Ethernet / FC port, loss of any storage connection, the inability to serve any volume, the inability to communicate with the host or the monitoring storage volume, then a

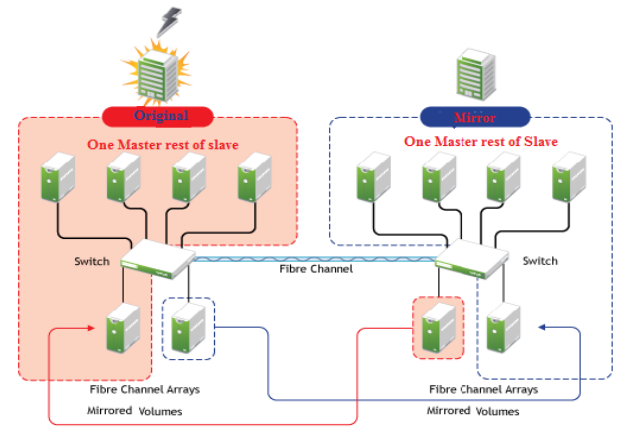

Figure 4: Cluster Mirroring Operation

Fail-Over event will be initiated. The standby unit is always ready to assume the active units duties when needed. If a component in either cluster member fails, Fail-Over automatically reassign the operations of the failed system to the surviving system. Fail-Over automatically redirects user requests from the failed system to good system.

2. Active-Active Clusters

A cluster consists of two Storage Concentrators work as Active-Active cluster [12]. In a Cluster system there is one primary and one Secondary Storage Concentrator. All host access the cluster through a single IP address. This allows all iSCSI target discovery to be accomplished with a single entry in the iSCSI initiator interface on the host. The session between the host and its discovered targets may be directed to either of the Storage Concentrators by the Administrator. The balance of IO traffic between the two Storage Concentrators may be adjusted at any time using the Administrative Interface on the Primary Storage Concentrator.

3. Storage volume Fail-Over Benefits

- Two units provide redundant CPUs, power supplies, hard drives, port connections and operating system.
- All customer configuration data from the active Storage Concentrator is replicated to the standby unit.
- Both the active and standby units can initiate a Fail Over event.
- An administrator can set up a Fail-Over event as a test or to service the active Storage Concentrator.
- Redundant Storage Concentrators ensure storage volumes are continuously available.
- A master IP address ensures server access to storage volumes no matter which Storage Concentrator is active.

4. How does this benefit?

The benefits of using Cluster Mirroring are:

- Cluster Mirroring is designed to be a simple-to-administer solution that extends fail-over capability from within a data center to a remote site.
- Cluster Mirroring is also designed to provide replication of data from the primary site to a remote site, which helps to keep the data at the remote site current.
- The combination of failover and data replication aids in the recovery from disaster helping prevent loss of data in less time than otherwise possible.
- Replicates data from the primary site to the remote site to ensure that data there is completely up-to-date and available.
- If Site A goes down, Cluster Mirroring allows us to rapidly resume operations at a remote site minutes after a disaster

IV. Performance analysis of IP SAN

The most important performance measures for the case of storage networks are average response time and network throughput. In this thesis the two parameters are employed to evaluate the iSCSI -based SAN. Average response time is the average time elapsed since the I/O request is sent by the initial device, until the I/O operations are finished by the destination device and the termination finishing message was sent back to the initial device by the destination device.
Throughput is the maximum amount of I/O request processed by the storage subsystem per unit time. Average response time and throughput can be measured from different views, for example that of the users, OS, and disk controller, etc. For simplicity, in our experiment the two parameters are measured from the OS of the server. To confirm that whether the iSCSI technology is well suitable to implement the SAN, we have carefully designed two groups of experiments, one for the FC-SAN (group 1), the other for iSCSI-based SAN (group 2).
In the experiment, the I/O meter is used to measure the throughput and average response time in different block sizes of the I/O request. The values of the experiment are shown in table, and the corresponding curves of throughput and average response time dependence on the block size are shown in diagrams.

Table: 1 Comparison between iSCSI based IP SAN and FC SAN

	Throughput(IO/s) in T time (ms)		Throughput(MB/s) in T time (ms)		Average Response Time (ms)	
Block size	FC	iSCSI	FC	iSCSI	FC	iSCSI
1KB	252.2162	1538.791	0.237378	1.40380	1.23856	3.84327
2KB	252.4374	1384.010	0.39518	2.60522	1.5383	3.83862
4KB	228.699	1201.280	0.797651	4.59580	2.17464	4.24649
8KB	205.5713	940.3271	1.51461	7.25428	3.26873	4.73416
16KB	178.5502	582.0236	2.707035	9.13099	4.02832	5.46148
32KB	150.2857	379.0472	4.630789	11.77945	4.96364	6.50105
64KB	123.0156	253.0086	7.651026	15.77455	6.02847	8.04707
128KB	96.72113	149.1703	12.12765	18.68377	8.33292	10.11468
256KB	75.76741	100.0032	19.18439	25.26578	12.46234	13.04359
512KB	49.12502	56.6429	25.11241	28.77153	17.247622	19.80151
1024KB	29.00699	30.7123	30.00799	31.71238	30.24288	32.21477

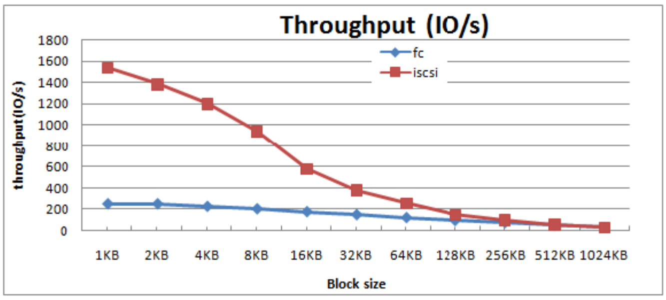

Figure 5(a): Throughput-IOPS

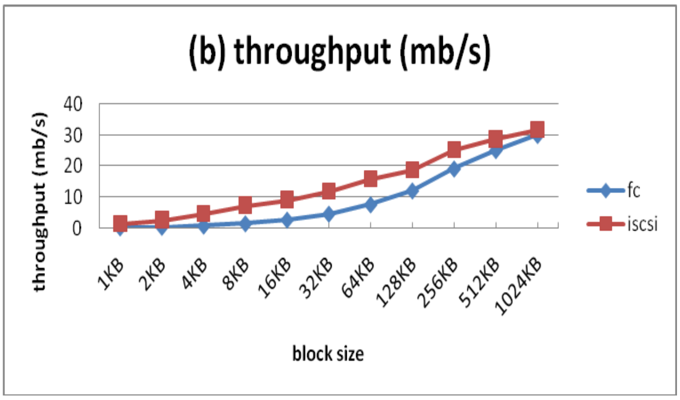

Figure 5(b): Throughput-MBPS

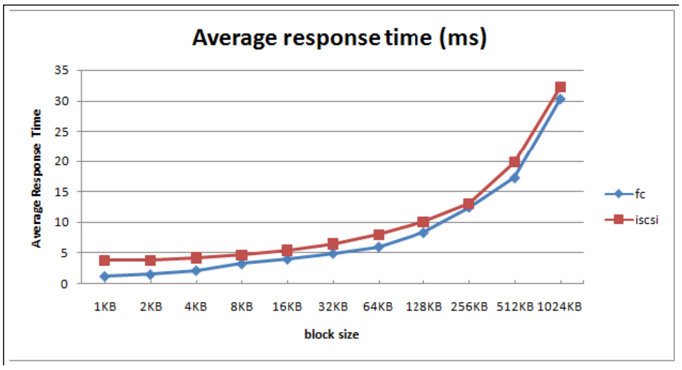

Figure 5(c): Average Response Time

Figures 5(a), 5(b) show the effect of block size on the throughput. As seen from Figure 5(b), when block size is increased, the throughput (MB/s) increases (but the IO/s decreases). Most importantly, the throughput iSCSI is bigger than that of FC. As can be seen in figure 5(a), when the block size is small (smaller than 64K), the throughput of the iSCSI is obviously larger than that of the FC. In most cases, the write/read request initiated by the application is decomposed into small I/O requests with block size of 4K-64K.

Thus, from the point of throughput the iSCSI is fit for implementing the storage networks. When the block size is larger than 64K, the difference of throughput decreases. Currently the performance of the storage subsystem is becoming more and more depending on the disk controller. When the block size is 1024K, the throughput of the iSCSI and the FC is similar. It can be understood easily that , in the experiment, to create similar experiment environment, the maximum throughput of the two disk subsystems is similar; and that when the block size is 1024K, the throughput of the storage subsystem has reached the maximum throughput of the disk subsystems.

In fact, after the throughput reached the maximum, increasing block size does not guarantee increased throughput but decreases the throughput instead.

Figure 5(c) shows the effect of block size on average response time. Note that average response time increases with block size. When the block size is more than 1024K, performance of the storage subsystem decreases intensively. Thus, large block size is not fit for the I/O request. Most importantly, when the block size is smaller than 1024K, the average response time of iSCSI is lower than that of the FC SAN. This indicates that the performance of iSCSI is better than that of FC as to the average response time.

V. Conclusion

In this thesis, we evaluate whether iSCSI is more suitable to implement the SAN, we analyze and compare iSCSI SAN and FC SAN with respect to several component of a network protocol which impacts the performance of the network. A Storage Concentrator and cluster mirroring enabled IP SAN delivers all the benefits of a traditional Fibre Channel SAN at a fraction of the cost, complexity and effort. By reducing management overhead, enabling clustering, strategic data protection and disaster recovery efforts, optimizing the use of existing and future storage. Cluster mirroring based IP-SAN model significantly increases the return on the overall storage applications.

This work presents how we can way out to of diagnose of problems occurred in IP SAN, Cluster mirroring with composition of different technology SCVM. This works provide the optimal solution for improving reliability, high availability, management and high performance. A Storage Concentrator and cluster mirroring enabled IP SAN delivers all the benefits of a traditional Fibre Channel SAN at a fraction of the cost, complexity and effort. By reducing management overhead, enabling clustering, strategic data protection and disaster recovery efforts, optimizing the use of existing and future storage. Cluster mirroring mixture model significantly increases the return on the overall storage application for SMBs.

References

[1] Technical report on implementing a Storage Concentrator in a Microsoft Exchange Network, StoneFly Incorporation.

[2] Technical report on DAS to SAN Migration Using a Storage Concentrator, Stone Fly Incorporation.

[3] Drew B, Network Storage-The Basics, http://www.enterprisestorageforum.com/technology/features/article.php/947551.

[4] De-Zhi Han, SNINS: A Storage Network Integrating NAS And SAN, Dept. of Computer Science, Jinan University.

[5] U.Troppens, R.Erkens, W.Muller. Storage Networks Explained: Basics and Application of Fibre Channel SAN, NAS, iSCSI and InfiniBand. John Wiley, 2004.

[6] Cluster diagnose on http://www.publib.boulder.ibm.com/infocenter/ibmxiv/r2/topic/com.ibm.help.xiv.doc/technologicalprimer

[7] Kriegel, Hans-Peter; Krger, Peer; Zimek, Arthur, Clustering High-dimensional data: A survey on subspace clustering, pattern-based clustering, and correlation clustering, ACM Transactions on Knowledge Discovery Data, Vol. 3, No. 1,Article 1, March 2009.

[8] EMC2 Corporation, Navisphere SAN Management Suite, http://www.emc.com/products/storagemanagement/navisphere.jsp; accessed July 15, 2007.

[9] Tony Redmond, Designing storage for Microsoft Exchange 2007, Syngress, USA. Page no. 320 - 340, 350 378 and 510 522.

[10] The case study of The benefits of clustered block storage http://silvertonconsulting.com/clusterblockstorage.html.

[11] W. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson. On the Self-Similar Nature of Ethernet Traffic (extended version). IEEE/ACM Transactions on Network.

[12] S. Aiken, D. Grunwald, and A. R. Pleszkun, A Performance Analysis of the iSCSI Protocol, in 20th IEEE/11th NASA Goddard Conference on Mass Storage Systems and Technologies (MSS03), 2003.

[13] Technical report on IP SAN Appliance and Virtual Appliance Users Guide from StoneFly Incorporation.

[14] Technical report on IP SAN Appliance and Virtual Appliance Users Guide from StoneFly Incorporation.