

New IP SAN Era in Cluster Mirroring

Mr. Rajendra Purohit

Asso. Prof. (CSE/IT), JIET Jodhpur, India
E-Mail: rajendra.pur@gmail.com

Mr. Amit Mishra

Associate Professor (CSE/IT),
JIET Jodhpur, India

Abstract – Storage replication is one of the back bones for network environments. While many forms of Network Attached Storage (NAS), Storage Area Networks (SAN) and other forms of network storage exist, there is a need for a reliable storage replication technique between distant sites (> 1 Km). Such technology allows setting new standards and removes demerits of network failover and failback systems for virtual servers; specifically, the growing storage need for effective disaster recovery (DR) planning. The purpose of this manuscript is to identify growing technologies such as IP-SAN that allow with remote storage cluster replication for virtual servers. This study (Cluster Replication) provides an analysis of 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.

Keywords – NAS, SAN, Cluster Mirroring, Disaster Recovery.

I. INTRODUCTION

With much advancement in the sector of local area networks (LAN), there is still a growing need for efficient storage networks. Data repositories grow every second while each type of storage system may be more or less efficient depending on applications. IDC estimates 2007 revenue for the worldwide storage services market to be \$31.7 billion and , forecast to grow to \$39.9 billion by 2012 [5]. The extremely expansion of storage area networks created new technologies for SAN management and administration. Traditionally, scalability of storage networks has been solved “on-the-fly”, i.e. adding more storage components to LANs, which creates gaps in hardware utilization, I/O processing times and manageability of storage networks. Today, SAN administrators more than ever face the challenges of the fast growing rates of specific data repositories that require reliable data replication with failover and failback network systems. Server virtualization has been rapidly changing in the past 5 years up to SCVM, whereas application and network virtualization is not expected to have many changes in the next 10 years [4].

II. STORAGE AREA NETWORK

SAN is a storage network specifically designed to interconnect high-speed storage devices and servers, as show in Fig.1. This specific network can span a large number of systems and geographic locations. Unlike in NAS, SAN [8] 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 Fig.1.




Fig.1. Interconnect storage devices

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.

EMC and VMware work together on many solutions that address the world of virtualization of storage area networks. In this ever-changing environment, the need for storage replication is even more important. Downtime directly results in financial loss and for many types of businesses it is unacceptable. Recovery Time Objective (RTO) becomes golden for many enterprise systems with a need for high availability. Various techniques of storage replication strive to address concerns for fast RTOs. This manuscript presents the benefits of cluster replication of remote storage as one of the advanced approaches in storage replication.

III. MIRRORING

As businesses have become more and more dependent on 7x24 accesses to business critical information, data replication has become an important technique to ensure that a user can rapidly return to normal operations in the event of disasters, operator errors or technology failures. Replication or mirroring, a component of the Intelligent Network Platform, provides a number of ways IT Managers can protect their critical data on a day-by-day basis. Mirroring or replication technology provides several methods for quickly recovering from a disaster. In addition to increased uptime and disaster recovery applications, mirroring can reduce backup time and increase reliability by using mirror images that are detached from the main mirror volume to perform zero-window, server-free backups. Mirroring is a virtualization technique that creates multiple copies of a data volume in order to achieve fault tolerance against the failure of an individual drive or drive array. For reads from a mirror volume, the data can come from any of the mirror images or copies. For writes to a mirror, the data must be written to each of the mirror images in order to ensure that valid I/O data exists on each and so they can be used as a disaster recovery mechanism.

In the Storage Concentrator implementation, I/O reads from a mirrored data volume are always read from the mirror image that was the original Logical Volume that was used to create the Mirror Volume. If this image is detached, then the read are made from the next mirror image, and so on.

Replication works, how it can be used in an IT environment to protect corporate data against loss and improve availability, and what considerations must be taken into account before implementing a various mirroring solution. There are following replication [9] designed for environments where maximizing application uptime and business continuity are paramount.

- Synchronous Mirroring
- Asynchronous Mirroring
- Cluster Mirroring

IV. CLUSTER MIRRORING

The cluster mirroring [9, 11 and 13] implementation also makes use of the ability to have multiple logging implementations. The cluster log implementation is a fault-tolerant client/server logging system in figure 2. 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 [9, 10 and 13]. 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.

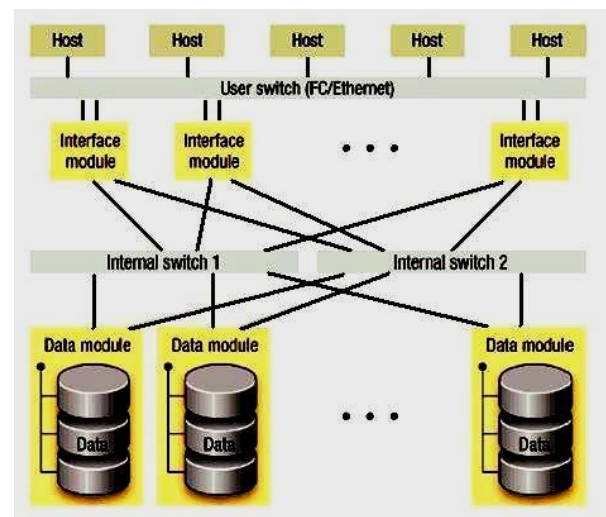


Fig.2. Cluster Mirroring Operation

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.

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.

Fail-Over

A FailOver [7, 8] 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 FailOver 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, FailOver automatically reassign the operations of the failed system to the surviving system. FailOver automatically redirects user requests from the failed system to good system.

Active-Active Clusters

A cluster consists of two Storage Concentrators work as Active - Active cluster [11, 13]. 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.

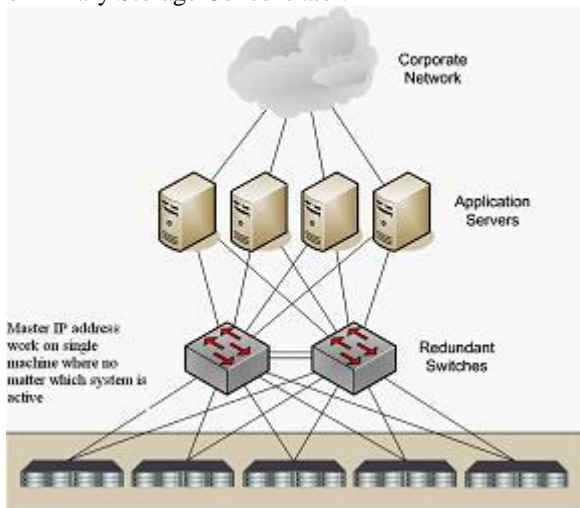


Figure 3 Cluster Mirroring Approach

Storage volume FailOver 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 and both units can initiate a FailOver event.
- An administrator can set up a FailOver 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.

How does this benefit?

The benefits of using Cluster Mirroring are in figure 4:

- Cluster Mirroring is designed to be a simple-to-administer solution that extends failover 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

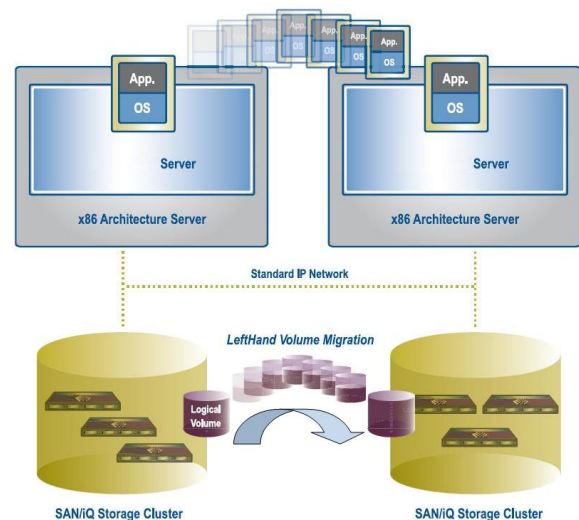


Fig.4. SAN cluster provisioning

How does it work?

Automatic failover and failback is also handled in a similar fashion - once the Primary Site is down, it automatically fails over a low-latency Ethernet connection to the Recovery Site. After the operations of the Primary site are restored, SAN acknowledges the response from the Primary site and fails the network storage back from the Recovery Site [4]. The advantage of using SAN is based in the way network storage is viewed. All storage resources may be viewed as one resource pool that allows for simplification of SAN management, remote replication, high scalability and reliability. According to IP SAN goes as far as entirely handling I/O load-balancing, which is critical for cluster scalability” [3]. The need for reliable SANs provisioning for virtual machines requires storage to be accessible at all times; moreover live migration of Logical Unit Numbers (LUNs) without interruption of

services is essential. Figure 4 illustrates how SAN technology provisions SAN clusters for servers.

Cluster mirrors take advantage of reciprocal mirroring feature. In cluster, primary and secondary storage volumes are maintaining back and forth a copy of each other's data [11, 13]. It improves the uptime of the overall system. No single point of failure operation is enabled by redundant connections to the SAN for the host systems, redundant active/active Storage Volumes and the associated cabling, and data mirrored. The cluster mirroring model most commonly used replication process in the disaster recovery [7]. It is most preferable in many application such as banking database, Organizations data base and etc.

How does it benefit?

In the event of a master failure, a Read Master (replica) is automatically promoted to become the new Master [12], the Master's Virtual IP addresses (VIPs) are instantly and automatically switched to the newly promoted Read Master, and updates continue to be processed without any service interruption. If a Read Master (replica) fails, its load is automatically switched and load balanced to surviving nodes. Fail-over is automatic and completes within a few seconds with no service interruption and no data loss. Moreover, using the synchronous cluster replication reduced downtime by 85% to 95% compared to the asynchronous or semi-synchronous replication

V. ADVANTAGES OF CLUSTER MIRRORING

- A. Removes distance limit of the mirror without impacting local performance at the primary site
- B. Sensible trade-off between performance of primary site and amount of unsynchronized data
- C. Real time protection against disk / frame / channel errors
- D. Vendor-neutral, protocol independent [5] , [7] (FC, SCSI, iSCSI)
- E. Disk upgrade without downtime
- F. Disk Failover functionality improve performance.

VI. CONCLUSION

Cluster Mirroring provide an intelligent centralized storage system for increasing disk utilization and disk space efficiency, and reducing data duplication. The results are significant benefits in manageability, resilience, and scalability. Cluster Mirroring fit seamlessly into existing storage (SAN/NAS) and TCP/IP data networks and provide storage provisioning for heterogeneous servers and storage systems in a SAN.

Cluster Mirroring is ideal for environments consisting of two storage systems located at geographically distant sites wherein the data must be kept as closely synchronized as possible. It also serves as an effective tool for zero-downtime migration of data from older disks to newer disks arrays. It can be reduce the downtime up to 30 percentages and improve the efficiency up to 20 percentages.

VII. ABBREVIATIONS

SAN	Storage Area Network
NAS	Network Attached Storage
IP SAN	Internet Protocol Storage Area Network
SCVM	Storage Concentrator Virtual Machine
iSCSI	Internet Small Component System Interface
VM	Virtual Machine
RTO	Return to Operation
NIC	Network Interface Controller
TCP/IP	Transmission Control Protocol / IP
FC	Fibre Channel

REFERENCE

- [1] Stonefly, Inc. technical report on , Storage Concentrators of "A New Dimension in SAN" Oct 2008
- [2] Chandler, Doug. Worldwide Storage Services 2008–2012 Forecast. Document # 212063. Storage and Data Management Services, IDC 2008.
- [3] Fegreus, Jack. SANitize DAS to enhance virtual infrastructure. InfoStor Lab Review 2007 http://www.infostor.com/article_display.content.global.enus.articles.infostor.topnews.lab-review-sanitize-das-toenhancevirtualinfrastr ucture.1.html
- [4] LefthandNetworks.MatchingServer Virtualizat with Advanced Storage Virtualization2007 <http://www.lefthandnetworks.com/webcasts.aspx>
- [5] Schreck, Galen. Forrester TechRadar™: Infrastructure Virtualization, 2008. Forrester Research, Inc. 2008
- [6] U.Troppens, R.Erkens, W.Muller. Storage Networks Explained: Basics and Application of Fibre Channel SAN, NAS, iSCSI and InfiniBand. John Wiley, 2004.
- [7] Cluster diagnose on <http://www.publib.boulder.ibm.com/infocenter/ibmxiv/r2/topic/com.ibm.help.xiv.d oc/technologicalprimer>
- [8] 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.
- [9] EMC2 Corporation, Navisphere SAN Management Suite, "http://www.emc.com/products/storage/management/navisphere.jsp"; accessed July 15, 2007.
- [10] Tony Redmond, "Designing storage for Microsoft Exchange 2007", Syngress, USA. Page no. 320 - 340, 350 – 378 and 510 – 522.
- [11] 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 (MSS'03), 2003.
- [12] <http://www.databasejournal.com/features/mysql/mysql-synchronous-cluster-replication.html>
- [13] Prof. Rajendra Purohit " Performance Evolution of IP SAN using Mirroring " IP SAN Performance using Mirroring in Publisher: LAP LAMBERT Academic Publishing (July 27, 2012)