

## Synchronous Replication of Remote Storage

Timur Mirzoev

Georgia Southern University

[tmirzoev@georgiasouthern.edu](mailto:tmirzoev@georgiasouthern.edu)

### Abstract

Storage replication is one of the essential requirements 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 synchronous storage replication technique between distant sites (greater than 1 mile). Such technology allows setting new standards for network failover and failback systems for virtual servers, specifically addressing the growing need for effective disaster recovery (DR) planning. The purpose of this manuscript is to identify the newest technologies, such as SAN/iQ and Storage VMotion, which allow for remote storage synchronous replication for virtual servers. This study provides an analysis and a comparison of various SANs that create solutions for an enterprise's needs. Additionally, the inter-operability of these technologies with the industry's leading product, VMware ESX Server, will be discussed.

### Introduction

With increased 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, forecasted to grow to \$39.9 billion by 2012 [2]. The tremendous expansion of SANs created new technologies for SAN management and administration. Traditionally, scalability of storage networks has been solved on-the-fly, by adding more storage components to LANs, which creates gaps in hardware utilization, I/O processing times, and manageability of storage networks. Today, more than ever, SAN administrators face the challenges of the growing rates of specific data repositories that require reliable data replication with failover and failback network systems. Server virtualization has been rapidly changing during the past five years, whereas application and network virtualization is not expected to have many changes in the next 10 years [12]. Companies, such as LefthandNetworks, utilize TCP/IP protocol for extending storage systems over IP and responding to the needs of virtual servers. For example, LefthandNetworks provides an excellent approach called thin provisioning. According to IDC, "thin provisioning, also known as oversubscription, allows administrators to maintain a single free space buffer pool to service the data growth requirements for all applications in a shared storage configuration" [13]. EMC and VMware work together on many solutions that address the world of virtualization of SANs. 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 synchronous replication of remote storage as one of the advanced approaches in storage replication.

### **Synchronous vs. Asynchronous**

Disaster occurrence may be caused by systems failures, human errors, or infrastructure failures; only a few to consider for DR planning. Each one of them may create a need to recover systems in a short period of time. If all remote storage targets are kept synchronized at all times, RTO can be minimized and systems may achieve 100 percent availability. Appropriate RTO and recovery point objective (RPO is the point in time to which a system needs to be recovered) change significantly with the replication approach. For example, RPO is zero for synchronous replication; for asynchronous replication, it may be small, depending on the network performance and storage systems availability. According to ITI Inc., large RTO translates to business' high tolerance to a "great loss of data" [7]. Table 1 presents the benefits and drawbacks of using synchronous or asynchronous replication of storage.

Table 1. Synchronous replication vs. asynchronous.

<b>Replication Type</b>	<b>Distance</b>	<b>Bandwidth</b>	<b>RPO</b>	<b>Availability</b>
<i>Synchronous</i>	Up to 150 miles	High bandwidth, available 100%	Zero RPO (replica is identical at all times)	100 %
<i>Asynchronous</i>	Global availability	Average	Small RPO	Depends on RTO

Each approach has its benefits and drawbacks. In general, synchronous replication requires higher investment rates but provides zero recovery time (sometimes, RTO is as long as the restart time). Distance limitations practically do not exist for asynchronous replication, which provides safer data repositories in case of disastrous events.

SAN technology, deployed at Colorado's health care institution, Denver Health, could serve as a vivid example of synchronous replication of storage for disaster recovery. Denver Health paired two technologies: 1) virtualization of servers, using VMware's ESX High Availability, 2) LeftHand SAN/iQ Multi-Site/Campus SAN [4]. According to David Boone, the Operations and Planning Manager of Denver Health, the institution saved approximately \$260,000 (running approximately 100 servers on six ESX machines), and "users don't notice when we're testing" [4]. Denver Health implemented synchronous replication of storage via low latency Gigabit Ethernet connection between two remote sites that are several blocks away with failover and failback systems, along with physical separation of storage clusters [1]. Figure 1 presents a diagram of the storage replication deployment at Denver Health.

## Campus SAN – Denver Health Deployment




Figure 1. Denver Health Campus SAN deployment.

For Denver Health, simplicity of DR management, SAN reliability, and scalability with direct saving on hardware resulted from the integration of server and storage virtualization. Another example of remote replication is by Mitsubishi Electric Automation, which runs SAP ECC6.0 database using SAN/iQ Snap, SAN/iQ Advanced Provisioning, and SAN/iQ Remote Copy. Mitsubishi Electric Automation was looking for redundancy, management, and performance in its business operations [8]. In the future, according to LefthandNetworks, Mitsubishi is considering adding approximately 15 TB of data for DR, CRM, and other applications [8].

### Virtualization of Storage

The virtualization concept is not a new idea. It is based on a time-sharing concept originally developed by scientists at Massachusetts Institute of Technology (MIT) in 1961 [3]. Server virtualization has rapidly advanced to the level that network storage systems must be available 100 percent of the time and provide reliability, manageability, and scalability. Highly scalable systems save money in the long run but may require higher initial investments. Although many companies claim high scalability, SAN administrators immediately know the limitations of deployed SANs as soon as storage space becomes inadequate. The popularity of VMware's Virtual Infrastructure 3 (VI3) enterprise solution progresses the advancement of virtualization of the server to the next level—distributed resource scheduling (DRS), high availability (HA), enhanced HA, and distributed power management (DPM). These are some of the essential technologies provided by VI3. One of the major components for VI3 is not just ESX machines and Virtual Infrastructure Management; it is the networked storage behind virtualization of servers. The availability of 100 percent for virtual servers is highly dependent on the availability of SAN running in the

background of ESX servers. LANs are usually viewed as the front-end networks, whereas SANs are considered back-end. Migrating virtual machines live, without interruption of service, is only possible if the SAN is available at 100 percent of the time. The importance of network storage has been stressed many times; however, with technologies such as virtualization of servers, the rules of enterprise networking are constantly changing. Various Fibre Channel (FC) technologies for network storage exist today, but certain limitations, such as high cost, scalability, and proprietary hardware, forced further advancement of Ethernet.

IP-based systems provide great scalability and standards. Protocols such as FCIP and iFCP provide vast benefits to enterprise network storage systems but also have certain limitations. FC-based SANs are great when there is no need to extend the SAN over a distance. As soon as the need for distance is involved, IP-based data provisioning proves to be more efficient. Perhaps for those reasons, iSCSI evolved a protocol that does not involve the use of any FC equipment (i.e., it is all Ethernet-based.) iSCSI simply transports SCSI commands over TCP/IP [11]. Ethernet-based systems are cost-effective and highly scalable systems, but they are limited by the bandwidth of the Ethernet channel. Today, 10 Gb/s is standard and no longer the “golden” bandwidth. Forty Gb/s and 100 Gb/s transfer rates are soon to become new standards [5].

Founded in 1999, LefthandNetworks is the company that takes Ethernet-based iSCSI SANs virtualization to a new standard—SAN/iQ. SAN/iQ technology was developed in 2001, and it is an x86-based highly scalable IP SAN [10]. Figure 2 presents the logical connectivity of SAN/iQ systems. LefthandNetworks calls its IP SAN storage units Virtual Storage Appliances (VSA). VSAs connect various network storage clusters, while providing excellent SAN administration tools. For example, physical disks may be added or removed from SAN, and the SAN/iQ management software sees the change and distributes storage resources without any need for the administrator’s attention [9].




Figure 2. SAN/iQ structure.

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/iQ acknowledges the response from the primary site and fails the network storage back from the recovery site [9]. The advantage of using SAN/iQ 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 an InfoStor lab test, “SAN/iQ goes as far as entirely handling I/O load-balancing, which is critical for cluster scalability” [6]. 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 3 illustrates how SAN/iQ technology provisions SAN clusters for VMware ESX servers.

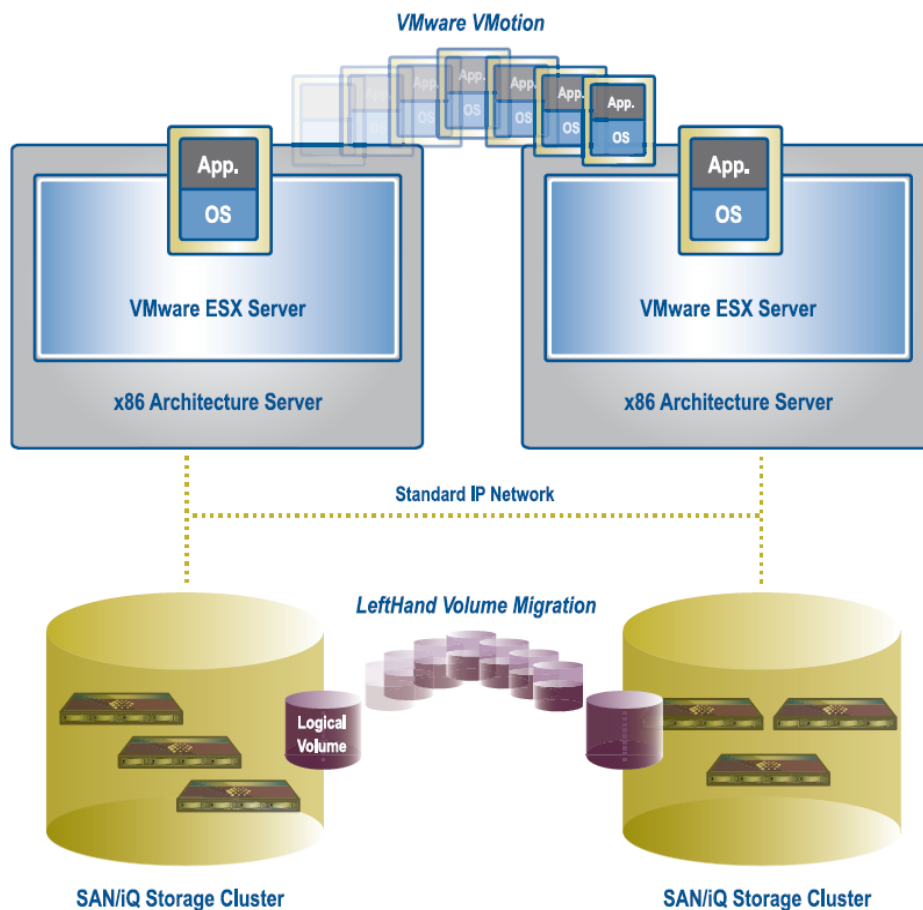


Figure 3. SAN/iQ cluster provisioning.

LefthandNetworks is not the only company that does storage virtualization. VMware is working on the Storage VMotion technology that virtualizes storage resources across IP-based networks. Figure 4 presents the Storage VMotion logical connectivity. The technology is similar to SAN/iQ clustering.




Figure 4. Storage VMotion.

With a successful deployment of virtualized IP SANs, remote replication and disaster recovery become extremely simplified and reliable. However, according to LefthandNetworks, there are certain limitations for SAN/iQ-based synchronous replication deployment: 1) low latency (2 ms), 2) high bandwidth, and 3) distance up to 20 miles [9].

## Conclusion

Synchronous replication of storage is not a novel technology, but today, it has progressed to the level where virtualized SANs increase return on investment (ROI) while providing effective solutions for storage replication, disaster recovery, and business continuity. The biggest advantages of virtualization technology of storage are the effective use of resources, availability, scalability, and manageability. Scalability becomes essential to network storage administrators when systems have to be reconfigured, upgraded, or migrated to a different physical location. IP SANs, specifically SAN/iQ, provide highly scalable solutions, while SAN management is simple with an increase in reliability. Today, server virtualization is essential, but it is time to realize the full benefits of storage virtualization technologies such as VSA (SAN/iQ) and Storage VMotion, which complement server virtualization with many benefits.

## References

- [1] Boone, D., Scott, B. (2007). *How Denver Health Is Using VMware and LeftHandNetworks Multi-Site SAN for Simple, Cost Effective Disaster Recovery*. VMworld 2007 BC26. Retrieved June 1, 2008, [http://inet-gw.maphis.homeip.net/training/VMWare/VMWorld-2007/Sessions/BC/PS\\_BC26\\_289343\\_166-1\\_FIN\\_v4.pdf](http://inet-gw.maphis.homeip.net/training/VMWare/VMWorld-2007/Sessions/BC/PS_BC26_289343_166-1_FIN_v4.pdf).
- [2] Chandler, D. (2008). "Worldwide Storage Services 2008–2012 Forecast." Document #212063. Storage and Data Management Services, IDC.
- [3] Corbató, F., Daggett, M., and Daley, R. (1961). *An Experimental Time-Sharing System*. Retrieved June 9, 2007, <http://larch-www.lcs.mit.edu:8001/~corbato/sjcc62>.
- [4] Denver Health (2007). *VMware Success Story: Denver Health*. Retrieved June 2, 2008, [http://www.vmware.com/files/pdf/customers/08Q2\\_ss\\_vmw\\_denver\\_health\\_english.pdf](http://www.vmware.com/files/pdf/customers/08Q2_ss_vmw_denver_health_english.pdf).
- [5] Dornan, A. (2008). "One Network to Rule Them All." *InformationWeek*, May 18, 2008.
- [6] Fegreus, J (2007). "SANitize DAS to Enhance Virtual Infrastructure." *InfoStor Lab Review*, retrieved September 5, 2008, [http://www.infostor.com/article\\_display.content.global.en-us.articles.infostor.top-news.lab-review-sanitize-das-to-enhance-virtual-infrastructure.1.html](http://www.infostor.com/article_display.content.global.en-us.articles.infostor.top-news.lab-review-sanitize-das-to-enhance-virtual-infrastructure.1.html).
- [7] Highleyman, B., Holenstein, P., and Holenstein, B. (2003). "Replicating Applications for Disaster Recovery." ITI Inc. White paper, retrieved June 10, 2008, [www.iticsc.com](http://www.iticsc.com).
- [8] LefthandNetworks (2007). "Case Study: Mitsubishi Electric Corporation, Inc." Retrieved May 25, 2008, <http://www.lefthandnetworks.com/document.aspx?oid=a0e000000000064AAA>.
- [9] LefthandNetworks (2007). "Matching Server Virtualization with Advanced Storage Virtualization." Webcast available at <http://www.lefthandnetworks.com/webcasts.aspx>.
- [10] LefthandNetworks (2007). "Matching Server Virtualization with Advanced Storage Virtualization." White paper.
- [11] Microsoft Storage (2005). *Storage Glossary: Basic Storage Terms*. Retrieved June 12, 2007, <http://www.microsoft.com/windowsserversystem/storage/storgloss.mspx>.
- [12] Schreck, G. (2008). *Forrester TechRadar: Infrastructure Virtualization, Q2 2008*. Forrester Research, Inc. 2008.

- [13] Waxman J., Yezhkova N., DuBois L. (2008). *LeftHandNetworks — Delivering Open iSCSI SANs for Midmarket and Enterprises*. IDC Filing Information, September 2007, IDC #208807, Volume 1, Tab: Vendors.

### **Biography**

TIMUR MIRZOEV is a professor of Information Technology Department at Georgia Southern University, College of Information Technology. Dr. Mirzoev has more than 10 years of experience in information technology, administration, and higher education. Some of Timur's research interests include server and network storage virtualization, information systems and security, LAN networks security, and WLAN networks security.