The Future of IP SAN using SCVM

[1] Rajendra Purohit, [2]Raviraj Choudhary

JNU Jodhpur / Faculty of Engg. & Tech., Jodhpur, India
Techno India NJR Inst. Tech. Udaipur / (Deptt. of CS / IT), Udaipur, India
E-mail:-[1]rajendra.pur@gmail.com

Abstract-- The SCVM, which stands for Storage Concentrator Virtual Machine, creates IP SAN storage as virtual machine, providing the user with the ability to consolidate their virtual data center. The SCVM can fully manage and has features for the storage over local as well as networked VMs. This can be attained through virtual switch or physical network connection. It provides an advanced, fully featured iSCSI SAN / Storage within an SAN Server. The user does not need another box for the storage. One just has to create an iSCSI Virtual SAN Appliance along with the Server Virtual Machines within the same hardware platform. This is not only Increases the productivity, but also simplifies the management, and reduces the power and rack space by simply loading an SCVM in a Virtual Machine. By creating an iSCSI Target within a Virtual Server, SCVM's customers may reallocate existing hardware resources to create business continuity and disaster recovery solution. This is achieved by any of two ways, be it SCVM's synchronous mirroring within the datacenter and distributed campus, or be it the asynchronous mirroring (replication) between remote facilities.

I. INTRODUCTION

The SCVM [1], which stands for Storage Concentrator Virtual Machine, creates IP SAN storage [2] as virtual machine, providing the user with the ability to consolidate their virtual data center. The SCVM can fully manage and has features for the storage over local as well as networked VMs. This can be attained through virtual switch or physical network connection. SCVM is totally hardware independent and portable. User can use, move, backup, restore, duplicate, scale and expand in local and remote data center anytime as part of physical and virtual environment. SCVM [1] creates complete Storage freedom for technology and data center manger who has embraced the virtual environment and is not tied down to vendors with traditional storage offering.

.

There is no need to purchase another box for the storage. Create an iSCSI Virtual SAN appliance along with server virtual machines within the same hardware platform. Increase productivity, maximize hardware utilization, simplify management, reduce power and rack space, and save money by simply loading an SCVM.

By creating a SCVM (Figure 1.) within a virtual environment, customers may reallocate existing hardware resources to create business continuity and disaster recovery solutions. Using award-winning Hybrid Replication which includes synchronous mirroring within the data center, distributed campus, asynchronous mirroring (remote replication) between remote facilities, and enables customers to have an instantaneous recovery and eliminate loss of access to storage during disasters. This tool enables our customers to lower operational data center costs plus provide a quick RTO (Return to Operation) in the event of a failure. In essence the SCVM is giving our customers the ability to create their own "DR (Disaster Recovery Box) [1],[2],[3]Site in a box".

.

If you are looking for all the functionality of an IP SAN, flexibility of deployment of virtual environment, ease of scalability, and future expansion, but short on budget and space, SCVM is the product you want. SCVM can be used.

•To create an IP SAN within a virtual environment that eliminates the need for a separate IPSAN hardware appliance.

•As a target for local replication from an IP SAN to a SCVM.

•As a target for remote replication from an IP SAN to a SCVM

•As a primary or secondary IP SAN Storage.

•As data flexible and portable data migration virtual appliance

"IT infrastructure is changing with explosion of virtualization and cloud environments. We are adapting to this natural change by introducing SCVM that can easily create IP SANs within a virtual environment. Other than the SCVM,

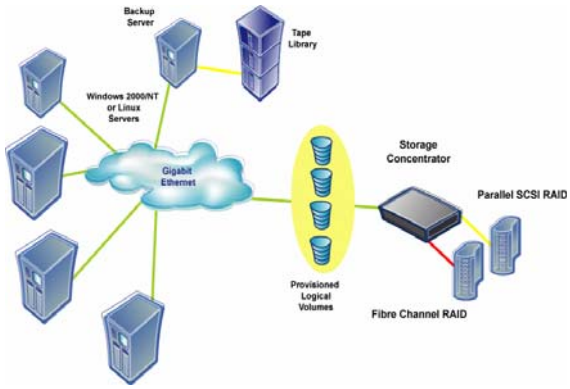

FIGURE 1: (SCVM) Protocols

Storage Concentrators provides a storage-mapping layer between the physical storage volumes and server requests, effectively moving block-level storage intelligence to the network core by consolidating disk management for full utilization of physical disk resources. Enterprise customers can, therefore, benefit from efficient storage utilization through storage consolidation, feature-rich managed storage capabilities, [4] rapid deployment of new storage, and SAN-based backup functionality. In addition, Storage Concentrators employ an easy-to-use user interface that allows administrators to easily manage storage operations from anywhere on the Internet.
When storage becomes a network resource, it can assume a range of capabilities. These include:

• Aggregation
• Logical-volume management
• Data protection services
• Support for bi-directional data transfers
• Storage-packet routing

The Storage Concentrator allows multiple servers/workstations and storage devices to store, fetch, and transfer data in a more efficient and economical way. With Storage Concentrators, enterprise storage is freed from certain physical restrictions and becomes an elastic resource that can be quickly and easily shaped to satisfy changing business and user requirements.
All Storage Concentrators are standards-compliant, offer a simple software upgrade path to new releases, and offer a wide range of performance options and features. Storage Concentrator product family includes both integrated single-chassis configurations, and highly scalable multi-chassis modular architectures.
The Storage Concentrator is designed for consolidation of physical disk resources and network storage provisioning. It makes logical disks appear as locally attached hard drives for improved storage utilization and administrative efficiency through "storage intelligence".

SCSI
SCSI is an intelligent, parallel Input/Output (I/O) bus on which various peripheral devices and controllers can

exchange information. For over two decades, SCSI has been the dominant protocol for transporting block-level data among servers. In fact, the SCSI standard was one

of the key factors in the development of open systems, providing a low-cost, high-performance storage interconnects.

Fibre Channel

Fibre Channel is another method, using a different/unique protocol, to transfer SCSI commands over a network. A Fibre Channel SAN [7] uses a switched infrastructure for storage. This infrastructure achieves high performance by relegating most of the protocol processing to hardware. However, Fibre Channel is costly and complex to install and is therefore used mostly in large data centers where the benefits outweigh the costs.

iSCSI

Storage Concentrators implement the Internet Small Computer Systems Interface (iSCSI) protocol. iSCSI [4] , [7]is a protocol that packetizes SCSI commands over industry-standard TCP/IP and Ethernet, much more ubiquitous network architecture with considerably lower up front and management costs. Since most companies already have a large network built around Ethernet, it is a natural extension to connect storage to that network rather than having to learn new ways of dealing with storage.
Almost any system with an Ethernet-based 10/100/1000 Network Interface Controller [4] (NIC) can be connected to an IP SAN through the simple addition of a software driver for iSCSI, often called an Initiator. These software drivers sits in the data path between the host file system and the network, packetizing and depacketizing standard SCSI commands so that they can traverse the IP network and communicate with other iSCSI enabled devices in the IP SAN.

A. SCVM'S ARCHITECTURE

Storage Concentrators architecture (Figure 2) that intelligently optimizes storage assets, offering the functionality traditionally associated with expensive

midrange and high-end storage systems, and host-based volume management VM.

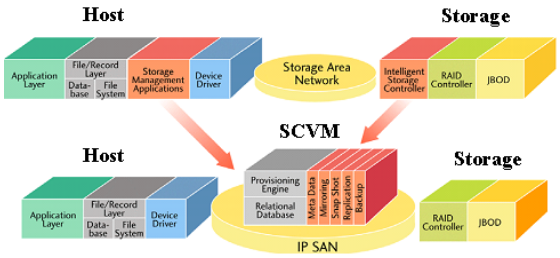

FIGURE 2: SCVM's Architecture

The SCVM architecture includes:
• An extensible in-band metadata storage-mapping layer that presents physical storage devices.
• Intelligent iSCSI storage packet routing software, providing aggregation and bi-directional data transfer for increased throughput
• A relational database that tracks physical data locations to ensure data integrity
• Online storage management to easily consolidate free storage space maximizing storage resources
• Extensibility of the SCVM architecture, allowing additional storage management applications to simply be plugged in without additional host agents

B. STORAGE CONCENTRATOR TOPOLOGY

Storage Concentrators are deployed using an organization's parallel SCSI or Fibre Channel storage devices, isolating servers, server operating systems, and applications from the physical limitations of storage devices (see Figure 1). (In environments lacking existing storage, Integrated Storage Concentrators can also be equipped with internal disks.)
Shielding servers from the complexities of physical devices simplifies media management and the allocation of storage space to each server or host computer. This "off-host" approach also avoids the introduction of additional latency caused by an extra device being present in the data path.

II. HOW THE STORAGE CONCENTRATORS WORK

Storage Concentrators operate in the network core and provide logical-volume management of physical storage volumes. The network administrator uses the Storage Concentrators to provide hosts with controlled access to these logical volumes. To ensure data integrity, Storage Concentrators use metadata to track physical data locations. This approach maximizes network intelligence by supplying non-disruptive online storage management to maximize physical storage resources. The results are improved storage utilization and administrative efficiency through in-network "storage intelligence."

IP SANs offer key features is including synchronous [5] and asynchronous [6] mirroring in SCVM.

SYNCHRONOUS MIRRORING

Synchronous Mirroring [5] option offers the ability to define a synchronous mirror for any disk managed (virtualized or service-enabled). Synchronous disk mirroring (See in Figure 3.) can protect data from storage device and channel failures. This solution is ideal for satisfying the demand for non-stop operations. Synchronous mirroring sends write transactions to both storage devices at the same time in parallel, updating both copies of the data before the write operation is confirmed.

Ideal for batch-oriented applications, synchronous mirroring ensures both data sets are kept in synch at all times. With synchronous mirroring, user can also migrate data from older disks to newer disks with no downtime, and failure protection of RAID can be enhanced with the ability to mirror across multiple cabinets.
Synchronous hard drive mirroring traditionally affects server performance, but with a solution the storage server can take care of the mirroring operations, which reduces I/O requests and CPU demands on processing dual write commands. User can have primary and synchronously mirrored disk from different vendors, even if one is Fibre and the other is iSCSI.
The mirror can be defined on disks that are not necessarily identical to the primary disk in terms of interface (SCSI, FC, iSCSI, Infiniband…). The process of creating the mirror does not cause any down time to the application. A mirror can protect against the consequences of device/cabinet/frame level failure. It is also a good tool to migrate data from older disks to newer disks without any downtime.

How does it work?

When a mirror is first created, the primary and mirrored disks (virtual or service-enabled) are synchronized to match data on both sides. After the synchronization is complete, all write-requests from the associated application server are delivered simultaneously to both sides of the mirror.

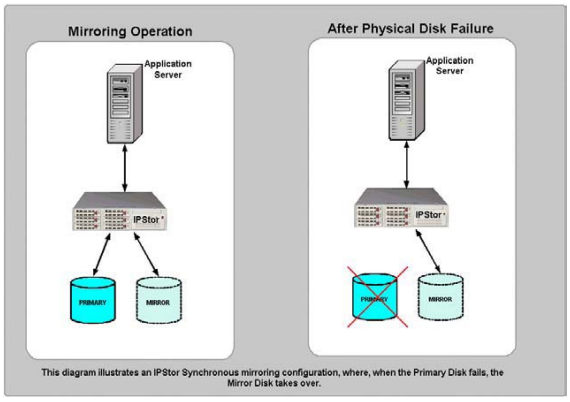

FIGURE 3: Synchronous Mirroring

It is important to realize that the dual-write process is controlled. With this design, only a single I/O request needs to traverse the front-end of the storage network, thereby eliminating the extra storage traffic created by host-based mirroring as well as relieving the host CPU from having to process dual write commands. The primary and its mirror can be swapped any time. Mirror-and-swap is a technique to effectively migrate data from old disks to new disks without any down time for the application server. After the swap, the mirror (which was the primary/old disk) can be removed and retired. All mirroring is managed from a single console, using the same creation process. Storage administrators no longer have to contend with application and/or OS-

specific host-based mirroring schemes, thereby greatly reducing management cost and complexity.

III. ASYNCHRONOUS MIRRORING

Asynchronous Mirroring [6] offers the ability to define a near real-time mirror for any disk managed (virtual or service-enabled) over long distances between data centers. If two storage systems located at geographically distant locations, asynchronous mirroring (See in Figure 4.) can keep data closely synchronized without draining performance. With asynchronous mirroring, the write transaction is acknowledged as soon as the local storage device confirms and completes the request.

Distance limitations between storage devices are also eliminated, because data can be committed to the local storage device without waiting for acknowledgement from the mirrored remote site, providing a significant performance advantage over many other mirroring solutions.

Ideal for applications that demand high-interactivity, asynchronous mirroring offers a performance advantage over asynchronous mirroring.

Asynchronous 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.

How does it work?

When an asynchronous mirror (See in Figure 4.) is first created, administrators create a dedicated staging area and associate this staging area to any disk managed (virtual or service-enabled). Once the mirror is created, the primary and secondary disks (virtual or service-enabled) are synchronized to match data on both sides.

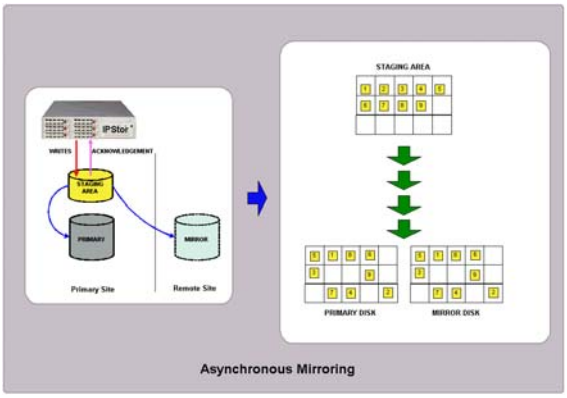

FIGURE 4: Asynchronous Mirroring

After the synchronization is complete, all write-requests from the associated application server are sequentially delivered to the dedicated staging area associated to the disks managed (virtual or service-enabled). The data written to the staging area is then committed to both the primary and its mirror as a separate process in the

background. The staging area can also be mirrored, for added protection.

IV. ADVANTAGES OF SCVM'S

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] , [6] (FC, SCSI, iSCSI)
E. Disk upgrade without downtime
F. Disk Failover functionality improve performance.

V. CONCLUSION

Storage Concentrators 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. Storage Concentrators fit seamlessly into existing storage and TCP/IP data networks and provide storage provisioning for heterogeneous servers and storage systems in a SAN. The Storage Concentrators maintain metadata that controls the mapping of real devices to the logical volumes that are presented to the individual servers. With the introduction of the Storage Concentrators, large, medium, and small organizations alike can finally leverage the benefits of storage networking using IP, including security, manageability, and quality of service, using their existing infrastructure.

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.

VI. ABREVIATIONS

SAN	Storage Area Network
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
DR Box	Disaster Recovery Box
NIC	Network Interface Controller
TCP / IP	Transmission Control Protocol / IP
RAID	Redundant Array Independent Disc
FC	Fibre Channel

REFERENCE

[1]. Stonefly, Inc. technical report on , Storage Concentrators of "A New Dimension in SAN" Oct 2008

[2]. U.Troppens, R.Erkens, W.Muller. Storage Networks Explained: Basics and Ap-

[3]. plication of Fibre Channel SAN, NAS, iSCSI and In_niBand. John Wiley, 2004.

[4]. Storage Networking Industry Association Technical Council, "SNIA Shared Stora

[5]. Model. A framework for describing storage architectures,"

[6]. 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

[7]. Synchronous Mirroring Option– At a Glance www.falconstor.com

[8]. Asynchronous Mirroring Option– At a Glance www.falconstor.com

[9]. A SNIA IP storage Forum, " iSCSI Building Blocks for IP Storage Networking" .