

SANRAD Application Note:

Getting the Best of Both Worlds

FC SAN and iSCSI SAN

Important Notice

The configurations described or tested in this application note are not the only available solution. This application note is not intended nor may it be construed as an endorsement of any product(s) tested.

This application note provides no warranty of any kind. SANRAD limited warranties for SANRAD products are stated in separate documentation accompanying each product. No damages or remedy of any kind shall be recoverable by any party for any claim or loss in any way arising or alleged to arise from or as a result of this application note, whether considered alone or in addition to any other claim.

Copyright SANRAD 2007

All rights reserved. The copyright and all intellectual property rights in this article belong to SANRAD. It is strictly forbidden to copy, duplicate or otherwise use this article or any part thereof in any way shape or form without the prior written consent of SANRAD.

Table of Contents

FC Reliability and iSCSI Economics Can you have the best of both worlds?.....	1
iSCSI connectivity is very low cost when compared to FC	1
Using iSCSI to put Additional Servers on the FC SAN	2
Can the V-Switch Simplify Volume Management?	3
Designed in High Availability	6
Centralized Data Protection and Backup.....	7
Conclusion	10

FC Reliability and iSCSI Economics Can you have the best of both worlds?

Some industry experts and vendors are pitting FC against iSCSI in a battle to rule the data center. Each side has strong arguments - FC proponents using performance and reliability arguments while iSCSI proponents claim better flexibility and lower costs.

The move away from FC toward iSCSI has significant implications. Do you start investing in iSCSI storage systems today because IP is much cheaper and easier than using FC? Or do you continue to invest in FC, committing the organization to a solution that may be trusted but continues to be relatively rigid and expensive?

This whitepaper reviews a third option which is to combine an iSCSI / IP network with your existing FC SAN, enabling you to deliver SAN benefits to any server including those once considered too low end for the FC SAN. It will review the SANRAD V-Switch that can extend your existing FC SAN with an IP SAN to consolidate and centrally manage the data from 100's of stranded servers.

iSCSI connectivity is very low cost when compared to FC

iSCSI connectivity is very low cost when compared to the present high cost of FC connectivity. The following chart shows the average cost of attaching an application servers using FC or iSCSI in a dual path configuration.

		Fibre Channel Cost			iSCSI Cost		
		Qty Per Server	Cost Per	Total Cost	Qty Per Server	Cost Per	Total Cost
Cost per Server	HBAs	2	\$1,000	\$2,000	N/A	N/A	\$0
	Storage Network Ports	2	\$1,000	\$2,000	2	\$50	\$100
				\$4,000	\$100		

It's difficult to justify attaching all your servers to a FC SAN using traditional FC HBAs and FC switches. As the chart indicates, the cost of attaching lower performing and less critical application servers is simply too expensive. Moreover, using traditional FC

SAN tools when creating, managing and protecting all the individual LUNs for tens to hundreds of low cost application servers is complex and time consuming.

Using iSCSI to put Additional Servers on the FC SAN

The benefits of having all the data cost-effectively stored on the same SAN are very compelling. They include:

- All storage capacity is shared and therefore more cost efficient
- Data is centralized for simpler backup and recovery
- Server maintenance and replacement is simplified
- Data is easier to secure and protect
- Existing SAN investments are leveraged
- New SAN investments are easier to justify

SANRAD's V-Switch is the simplest, most cost effective way to achieve these benefits.

V-Switches install at the edge of the FC SAN, between the FC SAN and Ethernet IP SAN network. V-Switches provide iSCSI to FC bridging with embedded storage services, enabling hundreds of servers to be connected to the SAN.

Other benefits of using iSCSI and V-Switches include:

- Remote data can be replicated to the SAN using iSCSI over IP
- Built-in multipathing and high availability
- Servers can be distributed and managed remotely across any LAN, MAN or WAN

With V-Switches, increasing the number of application servers storing data within the FC SAN is simple and straight-forward. Servers read and write data across an IP network to and from the V-Switches. V-Switches convert the traffic to FC and route it to the designated storage system/LUN within the FC SAN. V-Switches are equipped with FC ports (for connecting to the FC SAN) and iSCSI Gb Ethernet ports (for connecting to the Ethernet network to create the IP SAN).

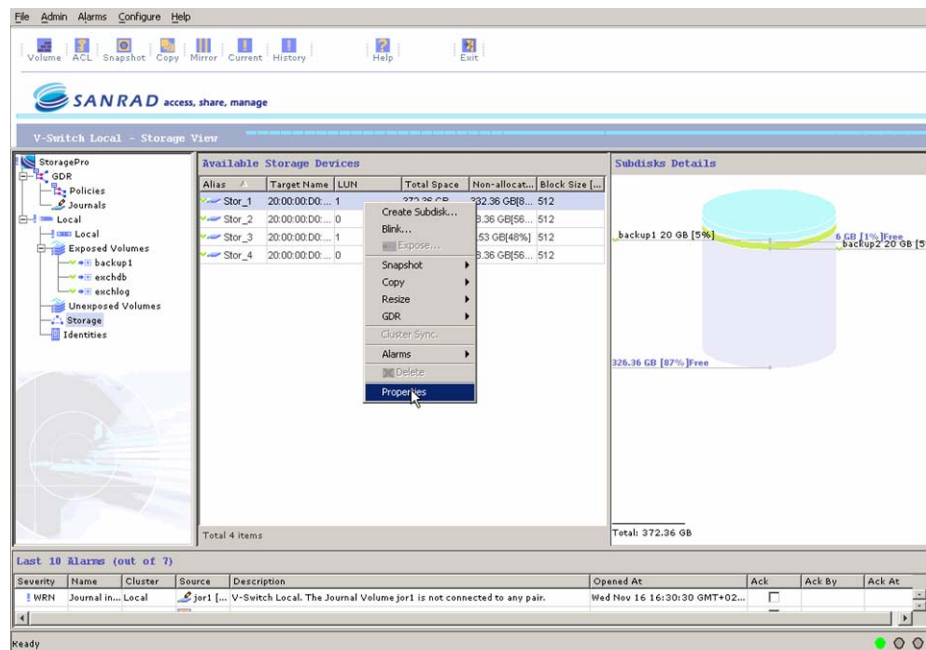
Can the V-Switch Simplify Volume Management?

One of the major issues in adding additional servers to the SAN is the task of managing storage resources for either a few to hundreds of new lower end servers. Servers connected over iSCSI can become numerous, be distributed over a large geographic area including remote offices. Traditional methods of FC SAN volume management are far too expensive, complex and maintenance intensive to support this type of infrastructure.

Unlike a simple iSCSI to FC bridge, the V-Switch has embedded volume management and storage virtualization capabilities designed to simplify the task of delivering and managing storage for hundreds of iSCSI connected servers. This allows servers to be connected, when needed, without requiring any changes to the FC SAN. iSCSI connected servers require no agents and no additional software. The only requirement is a generic iSCSI driver. This makes attaching, replacing and upgrading individual iSCSI servers a simple non-disruptive task providing vault isolation between the FC SAN and IP SAN.


The V-Switch performs volume creation and management using storage virtualization technology. The V-Switch receives a LUN from the FC SAN, just like any other FC connected server on the SAN. However, this LUN is not used “as is”. The LUN is placed into the storage pool of the V-Switch and is managed as a shareable resource. Additional LUNs from any FC array can be used to increase the capacity of the storage pool. From this central storage pool, unique virtual volumes are created for the iSCSI connected servers. By using V-Switch storage virtualization, a couple of LUNs and a few FC switch ports is all that is needed to create and deliver storage resources to hundreds of iSCSI connected servers.

This screen shows how a typical FC SAN LUN is virtualized into new smaller volumes. Complex virtual volumes can also be created to support mirroring, spanning, snapshots and/or striping.

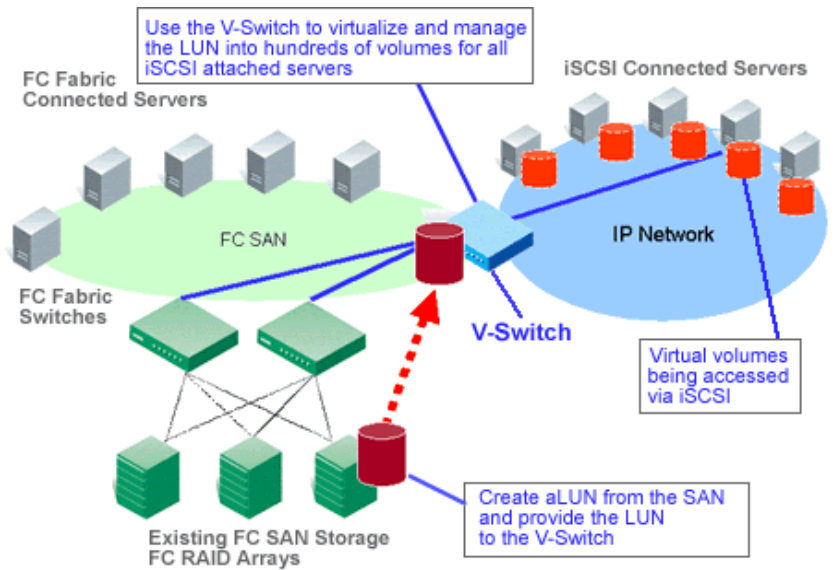


When attaching the V-Switch, the SAN infrastructure does not need to be changed. The V-Switch is attached via its standard FC initiator ports which resemble standard server HBAs with standard WWNs. The FC SAN administrator zones the storage LUN to the V-Switches in the same way as he would zone a LUN to a standalone server. The V-Switch will discover and log on to the LUN and read and write to the LUN in the same manner as any FC SAN-connected server. The entire process is as simple as adding one standalone server or an FC server to the SAN. The administration need only manage a single FC SAN LUN to deliver storage to many lower end servers.

This diagram shows a typical installation of a V-Switch at the edge of an FC SAN



This diagram shows a typical installation of a V-Switch at the edge of an FC SAN using embedded storage virtualization to create multiple virtual volumes from a single LUN.



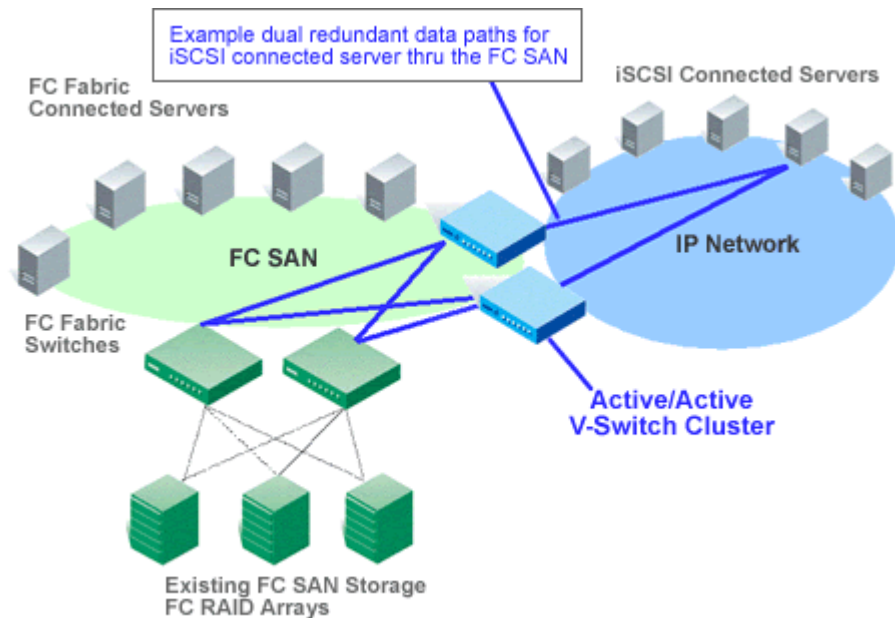
Designed in High Availability

Driven by global economy, 24 hr data availability is the norm. Using multiple data paths and V-Switch clustering ensures continuous data access for your staff and customers.

Multiple Data Paths: Within every server, there is usually only one iSCSI initiator. iSCSI initiators are native software drivers that are included with most popular operating systems. The initiator is used to connect to and move data to and from the V-Switches and servers. Within the initiator there is a connection layer. The connection layer manages the TCP/IP network connections between the server the V-Switches. In common non-critical applications there is only one iSCSI TCP/IP connection. But for high availability, there is a primary connection and an alternate connection. All iSCSI traffic between the server and the V-Switch travels over the primary connection. However, if the primary connection were to fail for any reason, the alternate connection would immediately become the new connection and automatically take over all iSCSI traffic. This mechanism eliminates the single connection concern. Some operating systems (e.g. Microsoft, SUN) have even more sophisticated mechanisms built into their iSCSI drivers.

V-Switch Clustering: Clustering is another key to enabling high availability and fail-over paths. In the event a V-Switch is temporarily offline or the network connection to the V-Switch is lost, a partner V-Switch attached to the same storage and the same host network can take over the IP addresses and data communication for the offline switch. All V-Switch clusters are “active/active” servicing their assigned servers, but they can also provide a “passive” failover path for other servers within the network. All V-Switches maintain the configuration information of other V-Switches within the same network topology and monitor the heartbeat of their designated cluster partner. When a V-Switch goes offline its cluster partner will sense its offline state and exposes the IP addresses from its cluster partner. The iSCSI connection layer, discussed previously, will discover the re-exposed IP addresses and create a new connection thus enabling the servers to proceed with data communication through a new V-Switch to the storage systems. The V-Switch will continue to service its own servers and the servers of the offline V-Switch until the original offline V-Switch is brought back on-line and the connection paths repaired. The entire failover and re-connect sequence occurs within a matter of seconds and has no effect on application server operation.

This diagram shows a typical high availability installation of a V-Switch cluster using multiple paths and platform clustering.



Centralized Data Protection and Backup

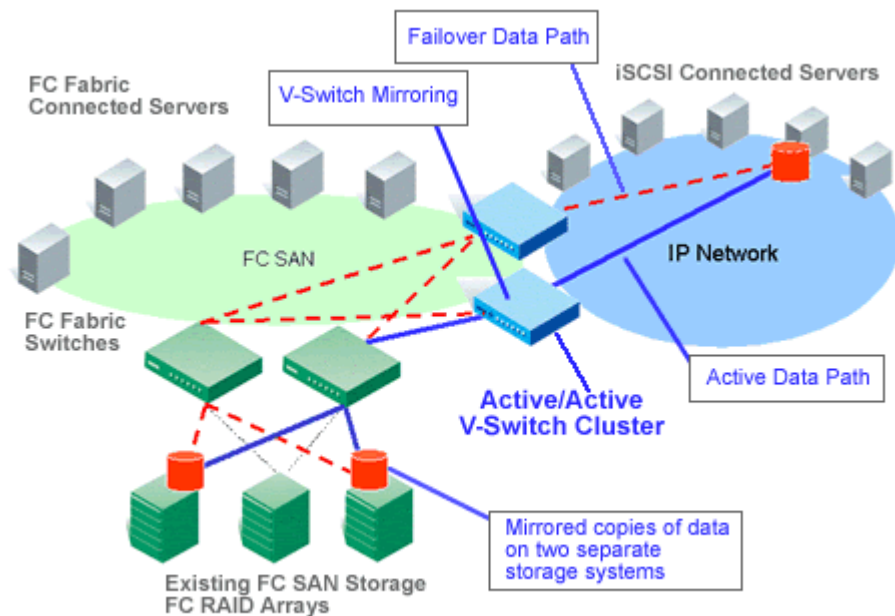
Business continuance and data protection are at the top of every “must have” list when deploying a new storage architecture, The V-Switches have three features that provide data protection for the iSCSI servers connected to the SAN:

- Data mirroring
- Local, campus-wide and remote replication
- Snapshot and rollback.

Data Mirroring: SANRAD’s V-Switch can create mirrored virtual volumes where the two volumes of the mirror are located on different storage systems within the FC SAN. Mirroring is done in real-time and also improves access performance because data requests can be serviced by any mirror member. Mirroring is performed by the V-Switches in a similar fashion as a RAID controller. The major difference is that RAID controllers mirror storage devices within a single array. Because the V-Switch operates in the network layer, it can create and maintain mirrored volumes anywhere within the FC SAN indifferent to traditional physical limitation such as enclosures and distance. Local synchronous mirroring can now be performed simultaneously to 2, 3 or 4 independent storage systems. The V-Switch can write and read directly and simultaneously to all storage systems/LUNs being mirrored. Like a RAID controller, if one of the mirrored partners goes offline or experiences a failure, the V-Switch will automatically remove the failed member from operation and continue to service iSCSI

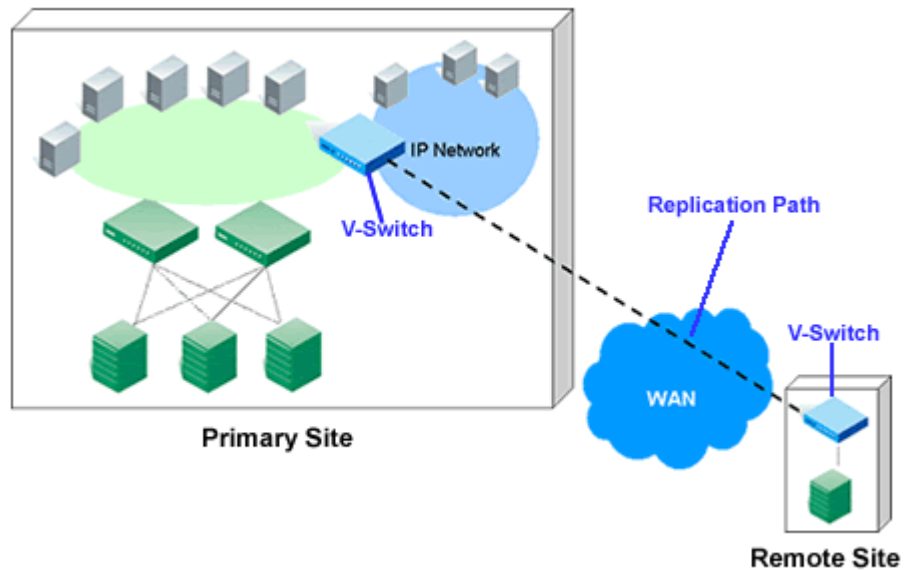
server requests with the remaining mirror member(s). The V-Switch can also take advantage of dark fibre to allow mirroring across a large campus topology. With V-Switch mirroring, server operation is not interrupted in any way if a storage system goes offline. The failover to the surviving mirror members/storage systems occurs in real-time and is completely invisible to the server.

This diagram shows a typical high availability installation of a two V-Switches in a cluster with data mirroring and full redundancy.



Remote Replication: The V-Switch can also replicate data across long distances using iSCSI over IP networks (LAN, MAN, WAN). Remote replication between sites using V-Switches can be deployed in *one-to-one*, *many-to-one*, or *many-to-many* topologies. This replication can be done synchronously (real-time) but is usually done asynchronously because IP links between remote sites are generally slower than internal LAN speeds. To keep replication as efficient as possible the IT professional can specify which volumes need to be replicated. During replication, new data changes to local production volumes are copied and cached in a local disk based journal volume. These data changes are formed into files call PiTs (Point in time snapshots) Based on the scheduled time PiTs are closed and replicated to the local site, between the V-Switches. Once the PiT reaches the remote site, they are checked for accuracy and then written to the corresponding remote volumes by the V-Switch. PiT technology provides a time/data consistent application mount point. The data and volumes at the remote site can be used recover a primary site or can act as a failover site. Remote site data and volumes can also be tested while replication is progressing to validate a recovery process.

This diagram shows a typical disaster recovery topology where data is replicated between the primary site and remote site.

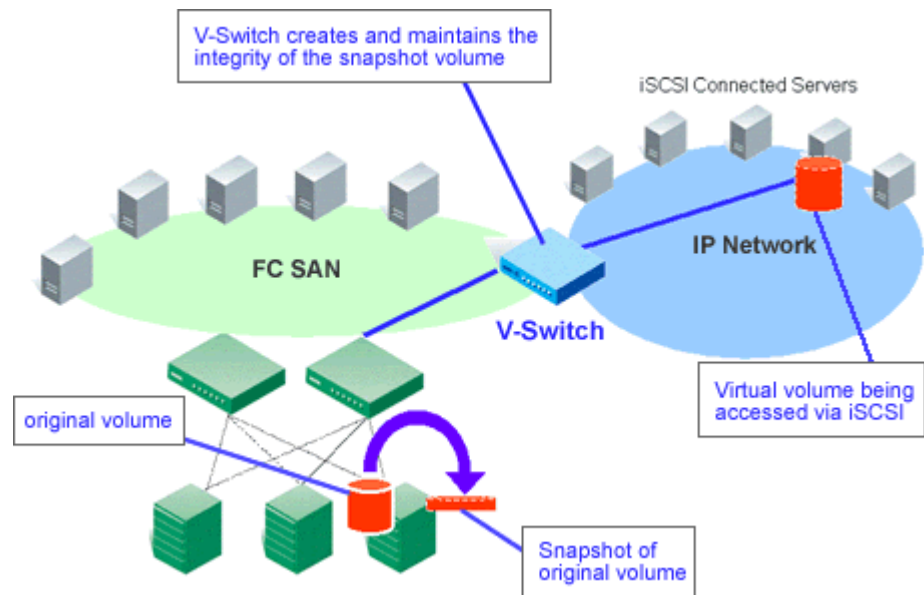


Snapshot and Rollback: A V-Switch snapshot represents a frozen image of the virtual iSCSI volume created and managed by the V-Switch for a specific iSCSI connected host. The source of a snapshot is called an "original." When the V-Switch creates a snapshot, it looks exactly like the original volume at that point in time. The snapshot remains the same even if changes are made to the original iSCSI volume.

SANRAD snapshots allow you to keep the original volume online and in production while using the snapshot for backing up to tape or disk. Using SANRAD snapshots is much more convenient than standard server based backups. SANRAD snapshots are 100% LAN-free and server-less.

SANRAD rollback allows you to use a snapshot to quickly return any original volume to the point the snapshot was taken thus eliminating the need to use tape to recover a volume. Snapshots store delta changes only and require much less storage capacity than original volumes.

This diagram shows a typical volume being protected by V-Switch snapshot



Conclusion

FC SANs represent a significant investment. Increased FC SAN and data center ROI is quickly realized by consolidating data management and protection tasks for hundreds of distributed servers into a single solution. FC SAN extension using SANRAD V-Switches leverage existing SAN investments, delivers storage consolidation with embedded storage services and takes full advantage of low cost and easy to use iSCSI.

For more information please contact your local SANRAD representative or visit www.sanrad.com