

**SANRAD White Paper:**  
**Continuous Data Availability Solutions using**  
**SANRAD's iSCSI Virtualization Switch™**  
**WP 003-03**

SANRAD  
US Tel: +1-800-471-2616  
International Tel: +972-3-767-4800

**Copyright SANRAD 2003**

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

INTRODUCTION.....	3
LOCAL AND REMOTE SYNCHRONOUS DATA MIRRORING AND FAILOVER .....	4
Local Synchronous Data Mirroring within a LAN.....	4
Remote Synchronous Data Mirroring within a MAN or WAN.....	5
SYNCHRONOUS REMOTE REPLICATION AND HIGH AVAILABILITY .....	7
THE V-SWITCH 3000 IP TAKE-OVER FOR ENHANCED HIGH-AVAILABILITY .....	8
SYNCHRONOUS REMOTE REPLICATION OVER IP .....	9
BACKUP AND RECOVERY TO AND FROM A REMOTE SITE .....	10
ASYNCHRONOUS REMOTE REPLICATION OVER IP .....	12
CONCLUSION .....	12

## Table of Figures

FIGURE 1.	SYNCHRONOUS REPLICATION USING SANRAD'S V-SWITCH 3000 .....	5
FIGURE 2.	SYNCHRONOUS REMOTE REPLICATION USING SANRAD'S V-SWITCH 3000 .....	6
FIGURE 3.	SYNCHRONOUS REMOTE REPLICATION USING HIGH AVAILABILITY .....	7
FIGURE 4.	V-SWITCH 3000 IP TAKE-OVER .....	8
FIGURE 5.	SYNCHRONOUS REMOTE REPLICATION OVER IP .....	9
FIGURE 6.	REMOTELY ATTACHED VOLUME USING SANRAD'S V-SWITCH 3000.....	10
FIGURE 7.	OFF-SITE REMOTE BACKUP AND RECOVERY USING SANRAD'S V- SWITCH 3000.....	11

---

## Introduction

The majority of today's businesses cannot survive a catastrophic loss of corporate data. Data is growing at exponential rates, forcing businesses to continually examine and enhance the availability, security and efficiency of their data environment. As one banking executive said, "The information we have on money is more valuable than the money itself."

Businesses recognize that remote data replication, multi-path failover along with faster backup and recovery times are essential to their survival and ability to service the new 24 x 7 global economy. Several mechanisms are implemented by most businesses to insure business continuity. Most legacy methods use RAID within a storage system and local attached tape for backup. Although these methods are useful, they are slowly proving to be inadequate. Business requirements for continuity plans and fault recovery demand greater levels of data protection and business continuance requiring off-site data replication, automatic storage system failover, smaller back-up windows and quicker file recovery.

These new high availability business requirements can be delivered by two essential techniques.

The first technique is the use of iSCSI and its ability to move data efficiently and securely whenever TCP/IP and Ethernet are deployed. iSCSI is a new data communication protocol that resides between the a host's OS file system and the storage infrastructure. Hosts can now reliably and cost effectively read and write data to and from storage systems over greater distances using an IP network. iSCSI is similar to any standard device driver located on the lower layers of the host operational stack and therefore will not interfere with standard applications, operating systems and file systems. Moreover, iSCSI is an IETF standard, which eliminates interoperability issues.

The second technique is in-band virtualization, which delivers volume management, data replication, failover control and data routing from the network layer. The introduction of this control at the network layer augments or replaces traditional volume management mechanisms that reside on the periphery, within the hosts or storage systems. By placing control in the center, network administrators have greater flexibility and control over data flow and can build more reliable and secure high performance storage networks.

This white-paper will review how SANRAD's V-Switch 3000 can be deployed to provide Disaster Recovery solutions in the following areas.

- Local and Remote Synchronous Data Mirroring and Failover
- Backup and Recovery to and from a Remote Site
- Synchronous Remote Replication and High Availability

---

## Local and Remote Synchronous Data Mirroring and Failover

Mirroring is a RAID technique for creating and maintaining identical data sets on different physical disks within an array. Mirroring protects data and keeps applications operational in the event of disk failure. Without mirroring, a damaged disk drive would halt an application and potentially cause data loss. With RAID, if a disk drive fails within a mirror, it will have an identical partner or partners that will seamlessly take over the I/O requests previously serviced by the down disk drive. Once the failed drive is replaced, the partner rebuilds the new replaced disk drive with the current data until its data is 100% resynchronized to match the partner. At this time the replaced disk drive rejoins its partner or partners and can again support I/O requests. Read commands are generally divided among all mirrored partners. Any new data write commands must be replicated and sent to all the mirrored partners to maintain data continuity across all mirrors. All partners must acknowledge and confirm the write request before the RAID controller returns a write acknowledgment to the file system of the host(s). This insures that all mirrors remain synchronized and the file system of the host(s) remains synchronized with the data on the mirrors. If a mirrored partner does not return a write acknowledgement, then the RAID controller will retry several times to confirm that one of the partners has failed. If there is still no write acknowledgement, the RAID controller will send an alert to the system administrator but it will continue to service file system I/O requests with the surviving partner or partners. Some RAID controllers provide “self healing” where un-used disk drives within the RAID system will automatically be selected as the new “spare” drive and be automatically synchronized and enabled as the new replacement disk drive.

### Local Synchronous Data Mirroring within a LAN

The V-Switch 3000 can perform mirroring with block-level virtualization capabilities. Mirroring is performed by the V-Switch 3000 in a similar fashion as a RAID controller. The major difference is that RAID controllers mirror storage devices within a single enclosure. However, because the V-Switch 3000 is in the network layer, it can create and maintain mirrored partners/ volumes anywhere within the network, indifferent to traditional physical limitation such as enclosures and distance. Local synchronous mirroring can now be performed between two or more enclosures. For example an V-Switch 3000 in Building A with a FC attached JBOD can keep the data files on the FC JBOD synchronized with a FC attached JBOD in building B, and another FC attached JBOD in building C. The V-Switch 3000 can maintain all three as partners within a mirror.

Like a RAID controller, if one of the mirrored partners goes off-line or experiences a failure, the V-Switch 3000 will automatically remove the failed partner from operation but will continue to service the application I/O requests with the remaining two mirrored partners. See Figure 1.

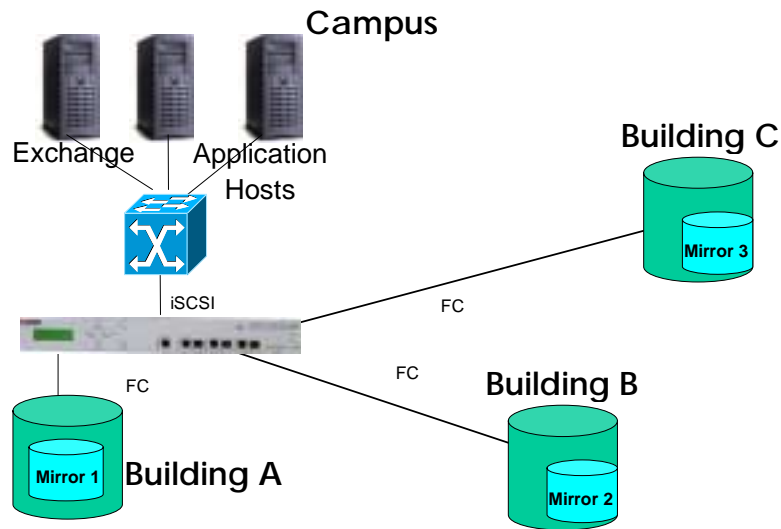


Figure 1. Synchronous replication using SANRAD's V-Switch 3000

## Remote Synchronous Data Mirroring within a MAN or WAN

High-end storage solution vendors use proprietary remote replication software to implement disaster recovery solutions for large enterprise applications. Such solutions require data to be remotely mirrored to a duplicate high-end storage system. The cost of purchasing a second storage system and the replication software can run into the hundreds of thousands of dollars, nearly 150% more than the initial storage system investment.

The V-Switch 3000 provides an alternative solution that enables companies to use any system from any vendor to create the redundant mirrored partner. An enterprise customer can use a high-end storage system for their primary storage and, with the V-Switch 3000's mirroring capability, implement a low cost remote replication solution using lower cost simple disk systems without degradation of performance levels.

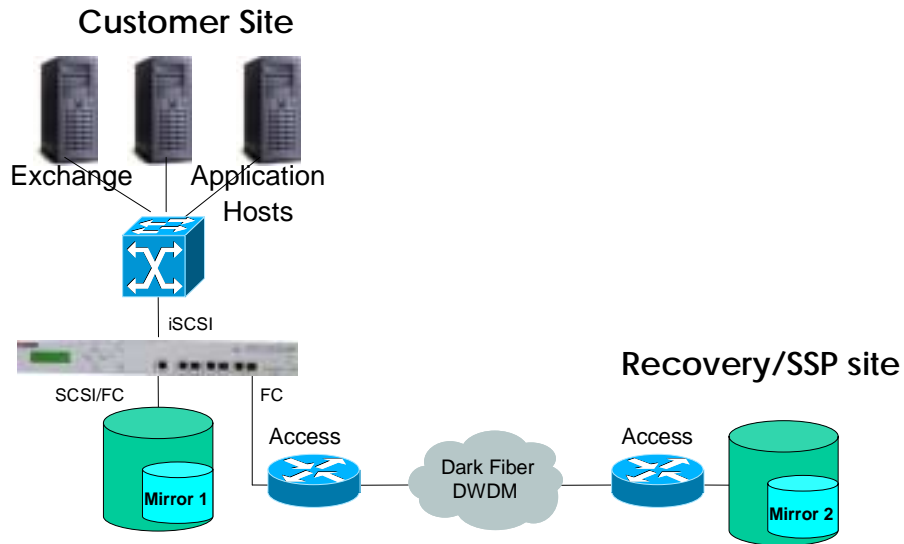
SANRAD's V-Switch 3000 provides synchronous mirroring using direct attached SCSI storage systems and direct attached or fabric attached FC storage systems. Remote Synchronous Data Mirroring within a MAN or WAN is accomplished by using standard FC tunneling techniques.

In the primary site, a local storage system is attached to the V-Switch 3000 using either a SCSI or FC connection. In the recovery site an FC disk subsystem is attached to the same V-Switch 3000 through a dark fibre/ tunneling technique like a DWDM connection. Dark fibre / DWDM service provides a high-speed low latency connection which is as fast as a direct FC or SCSI connection. Using SANRAD's iSCSI Virtualization Software, mirrored partner volumes are defined from the local and remote storage system. See Figure 2. During normal system operation, blocks that are written to the mirror are duplicated and written to the local mirror over FC or SCSI and the remote mirror partner using dark fibre/DWDM. Just like a RAID controller, a write acknowledgement is sent to the host only after the blocks are written to both local and remote mirror partners. In case of a disk subsystem failure in a primary site, the V-Switch 3000 directs

---

all read/write operations to the remote mirror partner connected over dark fibre/ DWDM. The hosts in the primary site continue working without interruption through the V-Switch 3000 to the remote storage systems even though the primary storage systems are off-line.

This method assures complete data integrity at all times. Similar to a RAID controller, after the disk subsystem in the primary site is replaced, this new local mirror partner is re-synchronized by the mirror partner at the recovery site and the system is recovered. The example below shows how an SSP (storage service provider) can use the V-Switch 3000 at a customer site to replicate and protect customer data.



**Figure 2. Synchronous remote replication using SANRAD's V-Switch 3000**

## Synchronous Remote Replication and High Availability

We can expand a simple remote synchronous mirroring installation into a complete storage disaster recovery configuration. We can use one site as a disaster recovery site for another or, as in this example, we will use two sites where each site is acting as the failover over disaster recovery site for the other. Either way, this is accomplished by having an V-Switch 3000 in both the primary and recovery sites. V-Switch 3000es are physically connected to all the hosts for both sites through the IP network. But each host is only communicating with the V-Switch 3000 within its primary site. The primary disk systems for each site are located and locally connected to the primary V-Switch 3000. For disaster recovery purposes, a mirror of the local disk systems is connected to the same primary V-Switch 3000, but located remotely through dark fiber/DWDM within the other site. To make the configuration complete, the V-Switch 3000 in each respective recovery site is given a physical connection to the mirrored disk belonging to the respective primary site. See Figure 3.

Using this topology, we form a symmetric configuration where each site is acting as a primary site and as well as a disaster recovery site for the other site. The V-Switch 3000 high availability feature monitors the availability of the other site by using the IP network to constantly send and receive heartbeat packets between the two V-Switch 3000es. For example, in case of a disaster in site 1, the V Switch in site 2 will detect the failure at site 1 and perform a failover process. Each connection between hosts, the down V-Switch 3000 and the down storage at off-line site 1 will be terminated and moved automatically to the designated recovery site - site 2. The operational V-Switch 3000 at site 2 will automatically re-establish a connection to the hosts seeking storage from down site 1. Once authentication has been confirmed, the V-Switch 3000 at site 2 will connect the hosts to the mirror of site 1 data. This mirror of site 1 data has been kept synchronized up until the time of the disaster with site 1 data by the V-Switch 3000 at site 1 performing synchronous mirroring over the dark fiber DWDM connection. This failover is done through an V-Switch 3000 IP-takeover process. See the next section.

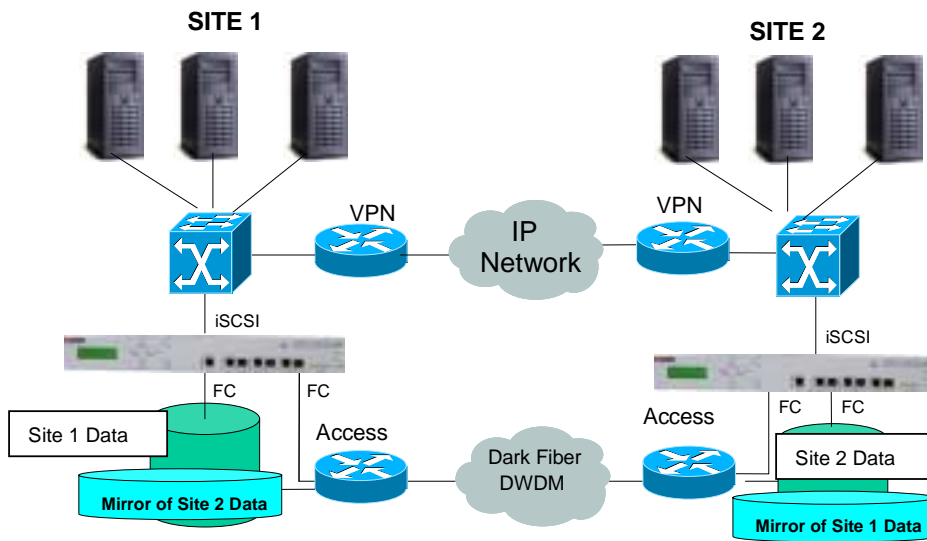


Figure 3. Synchronous remote replication using High Availability

## The V-Switch 3000 IP Take-Over for Enhanced High-Availability

IP Take-over is key to enabling high availability and fail-over paths with the V-Switch 3000. In the event an V-Switch 3000 is temporarily off-line or becomes overloaded, other V-Switch 3000es attached to the same storage and the same host network can take over the IP addresses and data communication for the off-line switch. All V-Switch 3000es are “active” servicing their assigned hosts but they can also provide a “passive” failover path for other hosts within the network. This is because all V-Switch 3000es maintain the configuration information of other V-Switch 3000es within the same network topology and monitor the heartbeat of their designated partner or partners. When a site or V Switch goes off-line, iSCSI will terminate the host connections with the problematic site but maintain the iSCSI session within the host while waiting for the IP addresses for storage to be re-exposed. Another V-Switch 3000 will now expose the IP addresses from the down site or V Switch. iSCSI will discover the re-exposed IP addresses and create a new connection thus enabling the hosts to proceed with communication through a new V Switch to the storage systems. The V-Switch 3000 will continue to service it’s own hosts and the hosts of the off-line V Switch until the original off-line V-Switch is brought back on-line and the connection paths or site are repaired. The IP addresses will be restored to the returning switch and the original connections will be re-established, thus providing continuous data availability. See Figure 4.

IP take-over can be used with a single data center to provide multiple failover paths to storage systems. It can be used across multiple data centers to automatically provide a new route to storage resources in the event of a sectional site failure. As mentioned in the previous section, when used in conjunction with V-Switch 3000 mirroring and FC or third party FC tunneling techniques, it can provide an off-site disaster recovery facility which re-connects hosts with mirrored partners within seconds of primary site failure.




Figure 4. V-Switch 3000 IP Take-over

## Synchronous Remote Replication over IP

Future additions of iSCSI initiators into the V-Switch 3000 leverage the remote replication process to work over the IP network. In this case we are not bound to the limitation of the fiber channel transport or fibre channel tunneling techniques such as dark fibre/DWDM. We can use any IP transport that provides enough performance bandwidth for the synchronous remote replication process.

In Figure 5 the volume is replicated between the primary and the remote site using SANRAD's V-Switch 3000 through a secured IP network within the LAN.

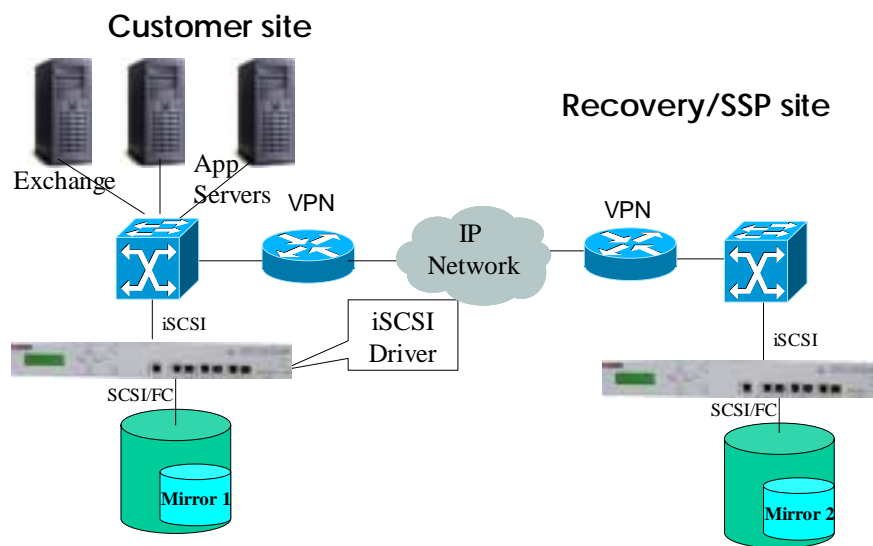
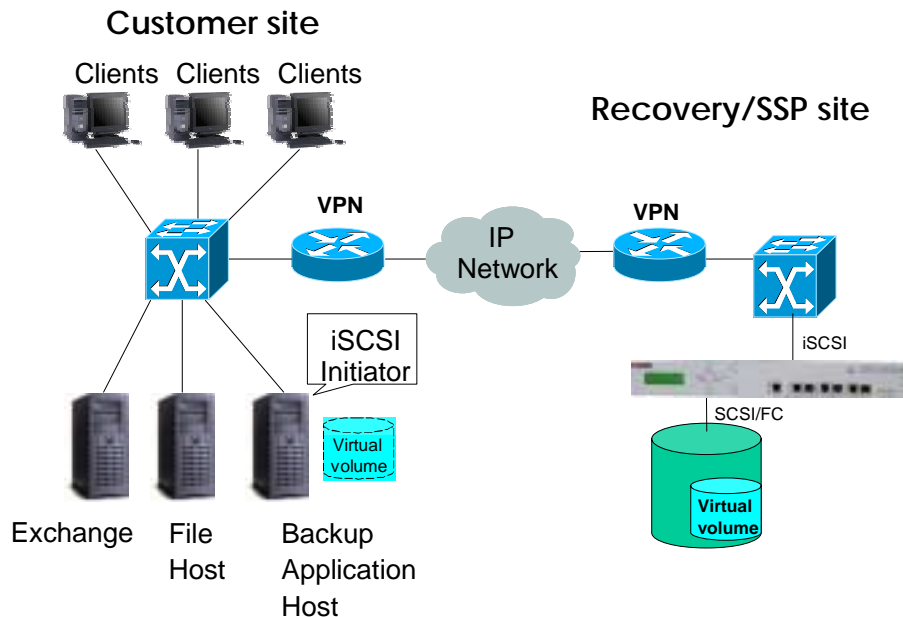


Figure 5. Synchronous remote replication over IP

## Backup and Recovery to and from a Remote Site

SANRAD's V-Switch 3000 enables secure off-site remote backup and recovery over any IP network. Using the V-Switch 3000, a disk or tape storage device can be used as a remote disk or tape recovery system. To configure this type of a solution, an iSCSI driver or iSCSI HBA (NIC) is installed in the local backup software application host. This host "sees" both the local storage systems and hosts to be backed up and, by way of iSCSI, sees the remote V-Switch 3000 and the attached tape or disk storage systems. The backup software application views and communicates with the remote iSCSI-attached backup systems as local storage devices even though they may be miles away in a different building, city or even in another country. The IT manager can use a simple file copy command, a backup software application or even an OS based mirroring feature to securely replicate files to the remote V-Switch 3000 and thus the remotely located storage systems. In this same fashion, files can also be read and recovered from the remote systems via the V-Switch 3000 and returned to the local backup host where they can be used to replace damaged or deleted files.

In Figure 6 the V-Switch 3000 presents virtual volumes to the application host. Critical files and data can be transferred securely to the remote recovery site using a simple copy command or backup software application.



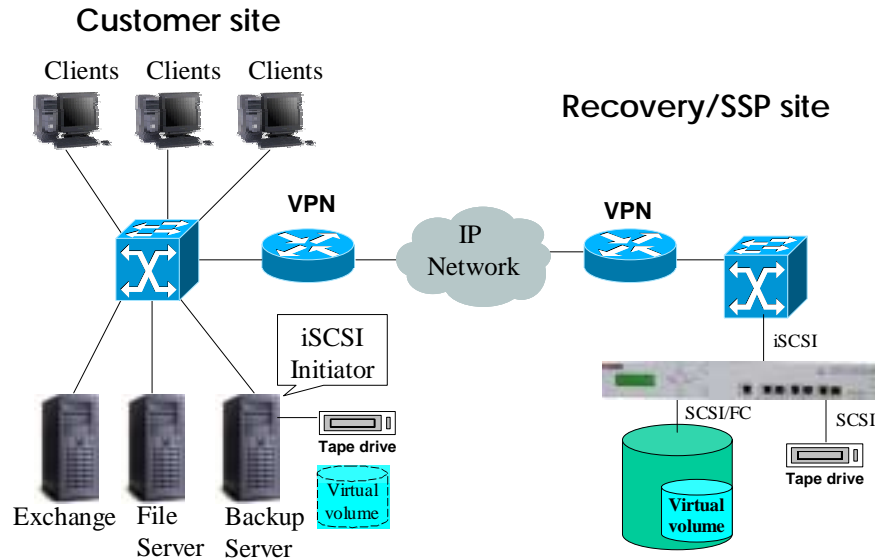
**Figure 6. Remotely attached volume using SANRAD's V-Switch 3000**

Using backup SW such as Veritas Backup Exec or BakBone NetVault we offer a highly efficient, secure and easy to implement off-site backup and recovery solution. Using incremental backup, open file backup, snapshot and file backup policies, we can configure remote backup and recovery routines to optimally utilize the existing bandwidth of the network without interfering with the normal operation of production systems and

applications. With off-site remote backup, data is replicated periodically from the backup host located in the main site into a remote disk or tape located at the recovery site.

With a backup application software and SANRAD's V-Switch 3000, remote backup can be performed over any IP-based network including wireless connections or the Internet.

In figure 7 an additional local tape backup system can be installed. Data can be backed up locally to tape in addition to remote disk and/or tape.



**Figure 7. Off-site remote backup and recovery using SANRAD's V-Switch 3000**

When adopting the remote backup option, an enterprise may decide to locate its recovery site with a data-hosting provider that specializes in hosting and facilitating the application. By locating the backup data at the service provider facilities, data is protected against a variety of natural and human disasters. By transferring data over an IP network, the geographic location of the enterprise recovery site is unrestricted.

As noted earlier, another benefit obtained through the V-Switch 3000 when combined with proper backup policies and applications, is that it allows for both tape and disk to be used as backup devices. Disk systems are generally faster than tape systems and thus can decrease the backup window as well as the time it takes to recover recently backed up and replicated files. Tape systems provide the best and most efficient media for longer-term backup and archiving. By using the V-Switch 3000 and combining new cost effective and simple disk systems with tape solutions, the backup and recovery times can vary from a few seconds for recent files still stored on disk to several minutes or hours for files migrated to tape. Network bandwidth and the amount of data that needs to be recovered will also effect recovery times.

We have presented an affordable remote backup and recovery solution for customers that would like to have a reasonable disaster recovery solution with minimal investment using their existing backup software application. SANRAD's V-Switch 3000 enables the IT manager to easily configure a remote storage device, either disk or tape, and have it used by the existing backup software to perform remote backup and recovery of critical files and data.

---

## Asynchronous Remote Replication over IP

Today, synchronous remote replication methods provide the most reliable and instantaneous disaster recovery solutions. They require a high bandwidth and low latency network to provide optimum performance. These solutions provide the best answer for implementing a disaster recovery solution today but they may be beyond the budgets of some businesses.

Future asynchronous remote data replication techniques over long distances will require a technology to address the limitations of lower speed networks. SANRAD offers a wide variety of solutions for disaster recovery and will also provide a cost efficient asynchronous remote replication option, which is totally transparent to the hosts and disk subsystems within an infrastructure. SANRAD is quickly developing an asynchronous remote replication option for its existing platform.

## Conclusion

Data is growing at exponential rates and businesses are continually examining and enhancing the security, reliability and efficiency of their data environment.

SANRAD fully understands the importance of continuous data availability, efficient and reliable backup and secure and rapid system and file recovery. The V-Switch 3000 provides businesses with local and remote mirroring/data replication and multi-path failover along with faster backup and recovery times. All these are essential to a business's ability to service the new 24 x 7 economy.