



## **SANRAD Global Data Replication and Disaster Recovery Solution**

**WP-009-01**

### **SANRAD**

US Tel: 800-461-2616

International Tel: 972-3-767-4800

Copyright SANRAD 2005

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.

## INTRODUCTION

Businesses recognize that remote data replication, remote backup, site failover and disaster recovery are essential to their survival and to their ability to service the 24 x 7 global economy. There are various Data Replication and Disaster Recover (DR) solutions on the market – each trying to answer, with varying degrees of success, some of the many needs of a good DR solution, such as central management, flexible volume selection and efficient use of available bandwidth.

SANRAD's Global Data Replication and Disaster Recovery (GDR) resides at the network level of your SAN and provides you with a centrally-managed, comprehensive, easily scalable solution that enables selective volume replication; prevents vendor lock-in; eliminates repeated host software installations and management overhead; enables both synchronous and asynchronous replication for optimal Recovery Point Objective (RPO) and uses guided failover and fallback wizards in conjunction with a preconfigured, server-ready secondary site for optimal Recovery Time Objective (RTO).

This white paper presents you with questions you should consider when designing your GDR solution and explains how SANRAD's solution works and what benefits it can offer you over the traditional host and storage-based solutions.

## WHAT ARE SANRAD'S BENEFITS OVER OTHER SOLUTIONS

Before examining the options available to you within the SANRAD GDR solution, let's review what SANRAD offers over other DR solutions.

While there are various DR solutions on the market, all of these solutions can be divided into three basic categories according to where the solution resides:

- **Host-based**
- **Storage-based**
- **Network-based**

### Host-Based Solutions

Host-based solutions are OS dependent and require an agent installed on each host – an agent that generally has limited OS support. This limits their scalability and increases management overhead as IT must configure and manage its DR solutions per OS and not centrally at the network layer. Further, because some host-based DR solutions operate at the file level, a single block change means replicating an entire file to incorporate the block change. This consumes CPU and memory resources, creating longer replication times and less efficient replication, which results in data copies that are less than up-to-the minute at the secondary site.

### Storage-Based Solutions

Storage-based DR solutions are vendor-specific leading to vendor lock-in or increased management overhead as IT tries to manage more than one storage vendor. Vendor lock-in typically eliminates the option of using lower-cost storage media at the secondary site, which would enable a company to contain its DR solution costs.

### SANRAD's Network-Based Solution

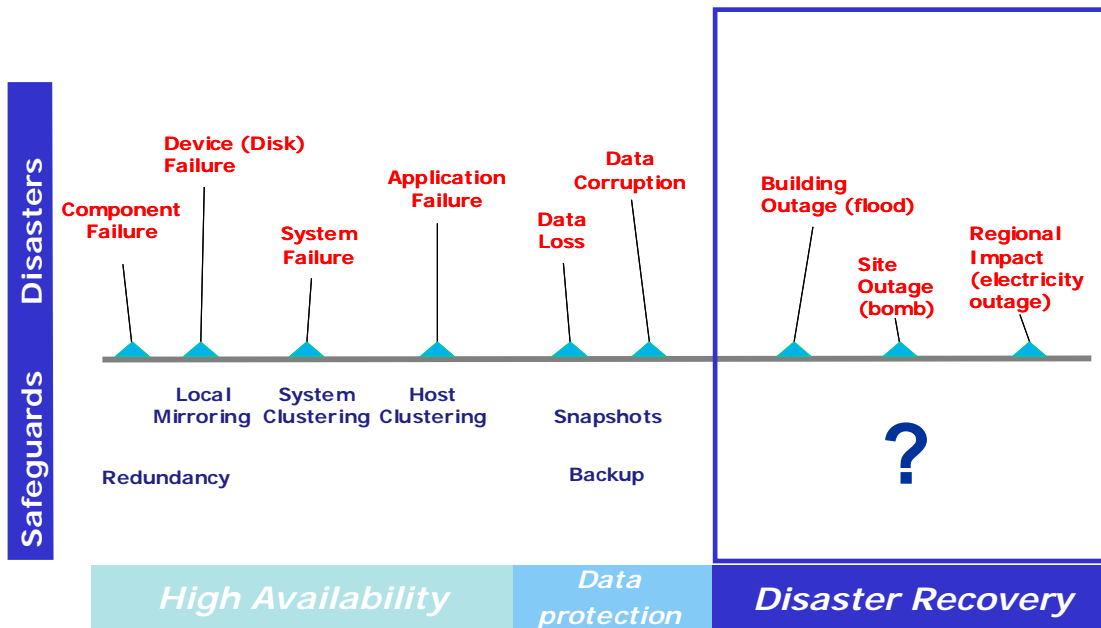
SANRAD's GDR solution lives at the network level. Based on SANRAD's volume virtualization technology, its GDR solution is both OS and storage agnostic making it easily scalable across multiple OS and varying-cost storage media. SANRAD's GDR solution replicates data at the block level providing the most efficient data replication thereby decreasing traffic on network links and resulting in the most up-to-the minute data copies at the secondary site.

## WHAT IS A DISASTER

Every well-planned network has built-in safeguards against localized failures. See Figure 1. High availability safeguards against technological failures. Components with dual power supplies to safeguard against a power supply failure; data mirroring to safeguard against a disk failure and host clustering to safeguard against application failures are three examples of high availability.

Data protection safeguards against human failure. Data backups to safeguard against data loss and volume snapshots to safeguard against data corruption are two examples of data protection.

A disaster recovery solution safeguards against a site-wide failure.



**Figure 1. Potential Points of Failure**

A disaster is some event that will shut down a network, leaving its data inaccessible to the people for whom it was intended. A disaster can be caused by several different elements that fall under the categories of natural, man-made or regional disasters. To safeguard against all of these disasters, a company must implement a disaster recovery solution. A disaster recovery solution enables a network to recover from a disaster and continue to operate normally.

## WHAT ARE MY GDR TOPOLOGY OPTIONS


When planning your disaster recovery solution, you need to be aware of the topologies possible and which can best serve your needs. A disaster recovery solution has two sites: the primary site that will be replicated and the secondary site that will hold the replicated data. With the SANRAD GDR solution, there are three possible topologies:

- **Active/Passive Configuration: One primary site and one secondary site**
- **Active/Active Configuration: One primary site and one secondary site**
- **Star Configuration: Many primary sites and one secondary site**

In all three topologies the primary site pushes the replicated data to the secondary site. The secondary site does not pull data from the primary site.

### Active/Passive: One Primary Site and One Secondary Site


In an active/passive topology, the secondary site normally acts only as a data warehouse. Volumes are replicated from the primary site to the secondary site. The secondary site sits in wait to take over network functions in the event of a disaster. See Figure 2.



**Figure 2. Active / Passive Topology**

### Active/Active: One Primary Site and One Secondary Site

In an active/active topology, both sites are productive (primary) data sites as well as replicated (secondary) sites for the other site. Volumes are replicated both from primary to secondary and from secondary to primary. See Figure 3.



**Figure 3. Active / Active Topology**

***Storage systems used at the primary site and remote site do not need to be the same brand or type. Select the storage type that best fits your requirements (FC or SCSI systems using FC, SCSI, SATA or ATA drives).***

### Star Formation: Many Primary Sites and One Secondary Site

A star topology is similar to an active/passive topology but, in this case, there are more primary sites but still only one secondary site. This topology is suited to a multi-branch company in which each branch (primary site) replicates its data to the headquarters (secondary site). See Figure 4.

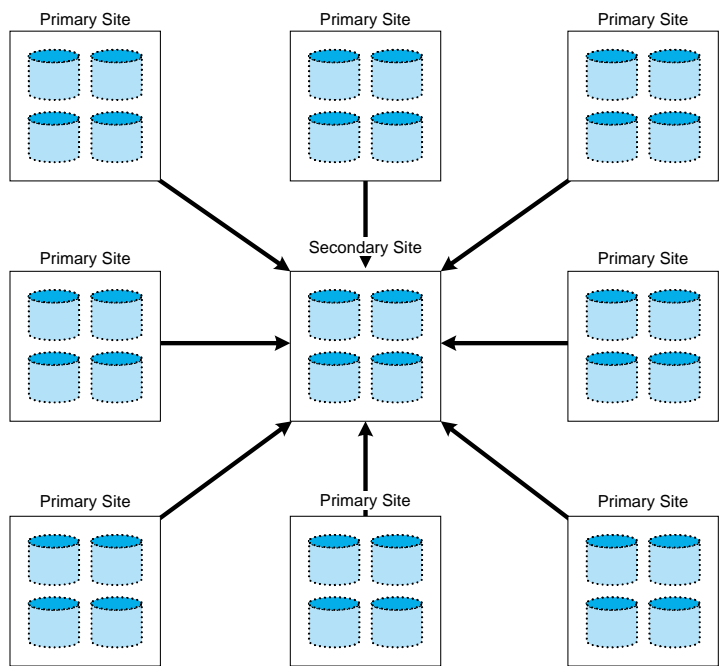


Figure 4. Star Topology

## GDR BASIC BUILDING BLOCKS

Once you have chosen the appropriate GDR topology you can concentrate on the actual replication. Next you need:

- Space for replication – DR journal volume
- Volumes for replication – DR pairs
- Groups for replication – Consistency groups

### DR Volumes

Physical storage space is needed at the secondary site to hold replicated data from the primary site. This is obvious. Beyond this physical storage space at the secondary site, another space, a *Disaster Recovery (DR) journal volume*, is needed per V-Switch. This volume is needed for administrative journaling functions of GDR replication. Volume journaling includes maintaining all changes made to the volume at the block level. A DR volume is expandable to accommodate growing journaling functions. See Figure 5.

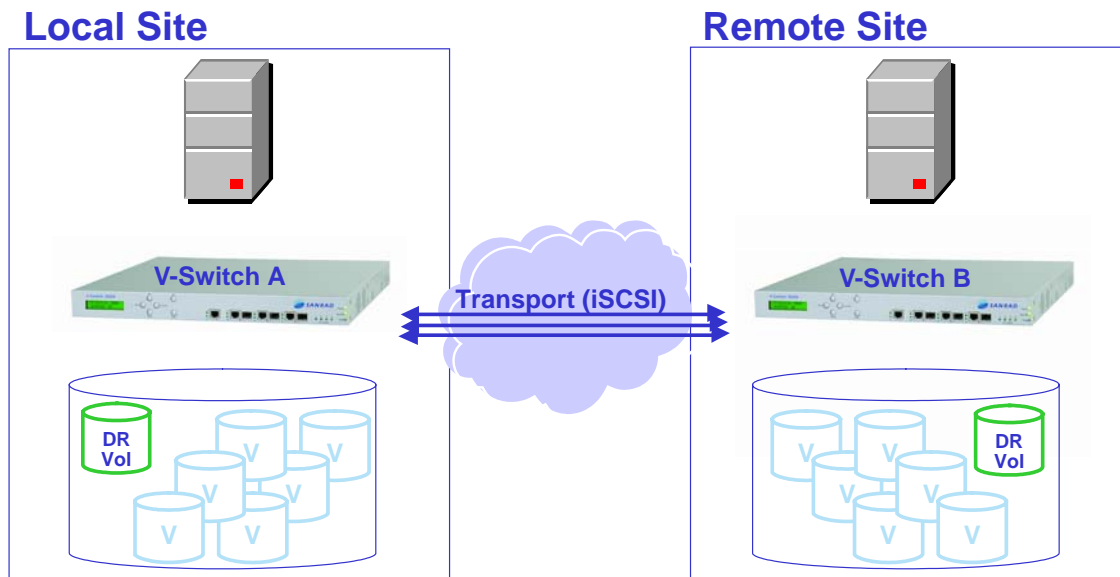



Figure 5. DR Volume

## DR Pairs

Your primary data is already virtualized into volumes on the V-Switch. These virtual volumes are your basic building blocks in your GDR solution. Each virtual volume can be replicated to a dedicated virtual volume at the secondary site. These two volumes together, the primary and secondary, form a *DR pair*. See Figure 6. SANRAD provides the flexibility of choosing to replicate all or only specific volumes.

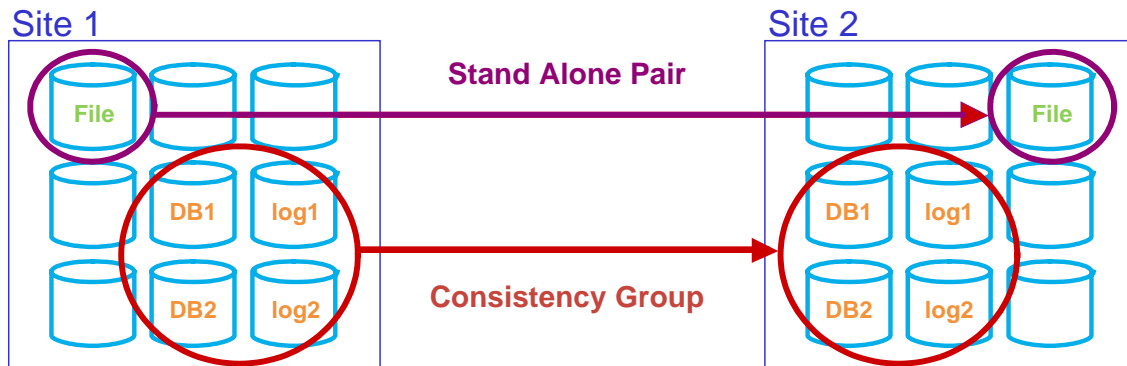


**Figure 6. DR Pair**

Now the DR pair needs replication parameters to direct its behavior.

## Consistency Groups

Because some applications create or require more than one volume with inter-volume dependencies, e.g. databases and email applications that have several volumes for databases and several volumes for logs, you can create a consistency group. A consistency group applies a given set of replication parameters to a group of DR pairs to ensure that they are replicated together. See Figure 7. Their PiTs are transferred together and only when all PiTs have been successfully replicated are they merged to the secondary volumes. This maintains cross volume data consistency. DR pairs not in a consistency group are stand alone pairs.

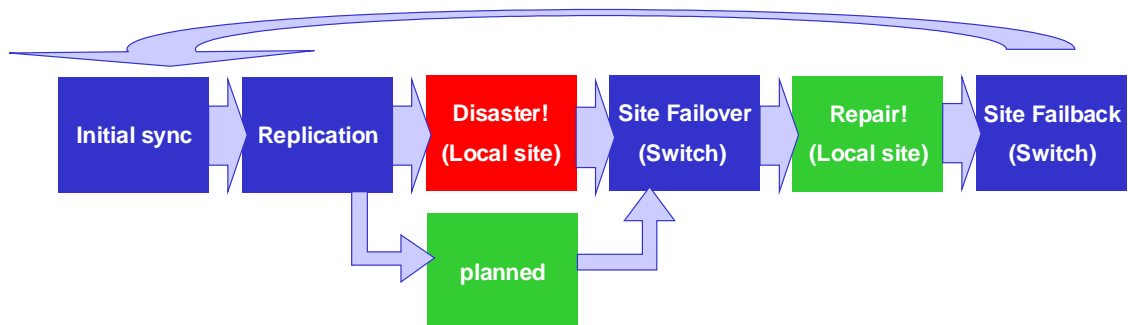


**Figure 7. DR pairs within a Consistency Group**

***Consistency groups ensure that related volumes are replicated at the same time to maintain data flow integrity for databases and other multiple-volume applications.***

## GDR WORKFLOW PROCESSES

Now that you have all the elements of the story, how does the story go? What is the actual GDR workflow? See Figure 8.



**Figure 8. GDR Workflow**

SANRAD's GDR wizards move you through all stages of your GDR solution implementation and execution, including failover and fallback.

First, volumes marked for initial synchronization are replicated to the secondary site. Next, routine GDR replication follows, as per each stand alone DR pair or consistency group.

If or when the primary site goes down, due to scheduled maintenance or disaster, the primary site is manually failed over to the secondary site and all DR volumes are exposed accordingly to their designated hosts. At this stage the secondary site acts as the primary.

The secondary site continues to function as the primary site while the primary site is repaired. When the primary site is repaired or scheduled maintenance is finished, the secondary site is manually failed back to the primary site.

## GDR INITIAL SYNCHRONIZATION OPTIONS

Before you begin you must set your data replication baseline. Do your primary volumes already contain active data? If so, you need to perform an initial volume synchronization action to replicate primary data to the secondary site.

There are three initial synchronization policies:

- **No Sync – no initial volume copy**
- **Online – copy volumes over the network**
- **Offline – copy volumes using backup and restore operations**

If your primary volume is a new volume and contains no data, there is no need for an initial data synchronization. There is no data to synchronize between the primary and secondary volumes.

In both online and offline synchronization a snapshot is taken of the primary volume. In online synchronization, the snapshot of the primary volume is pushed to the secondary site in real time. In offline synchronization, the snapshot of the primary volume is replicated to a selected volume and the media is transported to the secondary site as you see fit.

Once the initial synchronization replication is complete, the routine GDR replication schedule, either synchronous or asynchronous, will run as scheduled.

---

***Select only the volumes you wish replicated to conserve bandwidth costs and efficiently utilize storage capacity.***

## GDR REPLICATION OPTIONS

The method you choose for replicating your data will depend on several factors but the most important are available bandwidth and latency between sites.

The GDR solution offers you two methods of replication:

- **Synchronous – suited to high bandwidth; low latency**
- **Asynchronous – suited to low bandwidth; high latency**

A site can use a combination of both synchronous and asynchronous replication methods. When deciding which method is most appropriate for each DR pair, consider these issues: How quickly does the data change? When is the data most active? What data shares interdependency with other data and, therefore, should be replicated together? What kind of impact do small changes in a certain data chunk have on business functioning, i.e. what is the data loss tolerance?

### Synchronous Replication

Synchronous replication means that every write operation is written to the primary volume and to the second volume before sending back a write acknowledge to the server. A synchronously replicated DR pair functions similar to a regular mirrored volume over iSCSI, with each volume as a child. However, if the local volume fails, the volume must be manually failed over to the remote volume for the remote volume to take over regular volume functions. See Figure 9, page 14. Synchronous replication provides zero data loss. If your network enjoys high bandwidth and a low latency, e.g. campus LAN, synchronous replication is a viable option for your GDR solution.

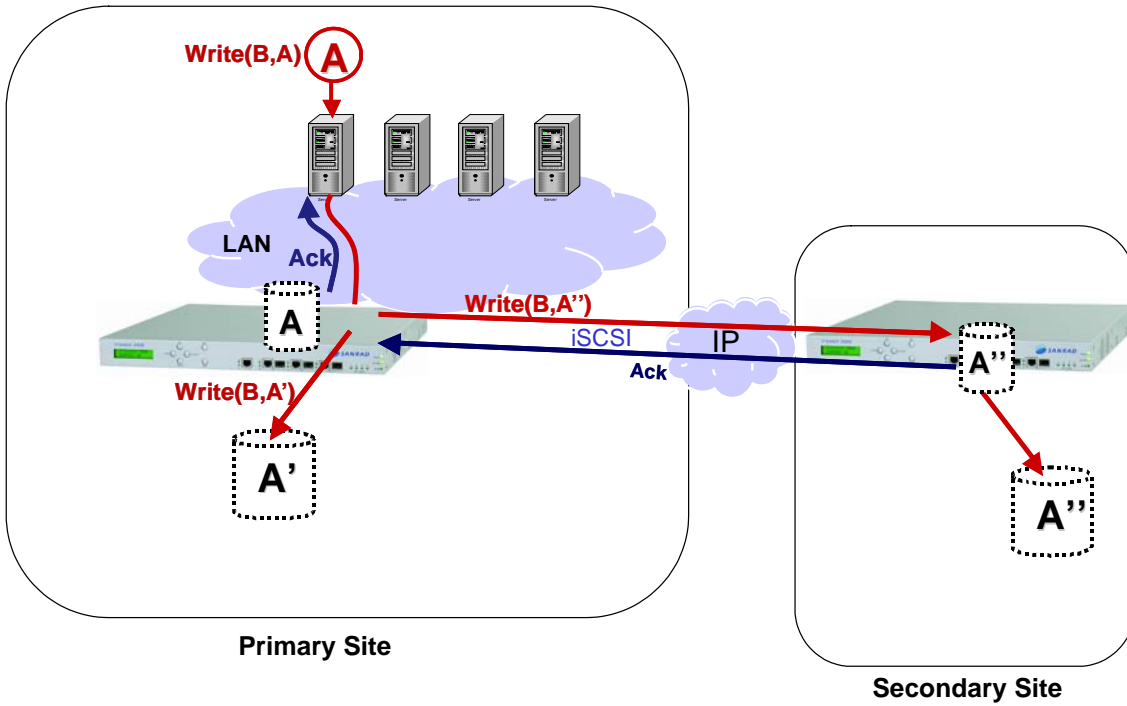
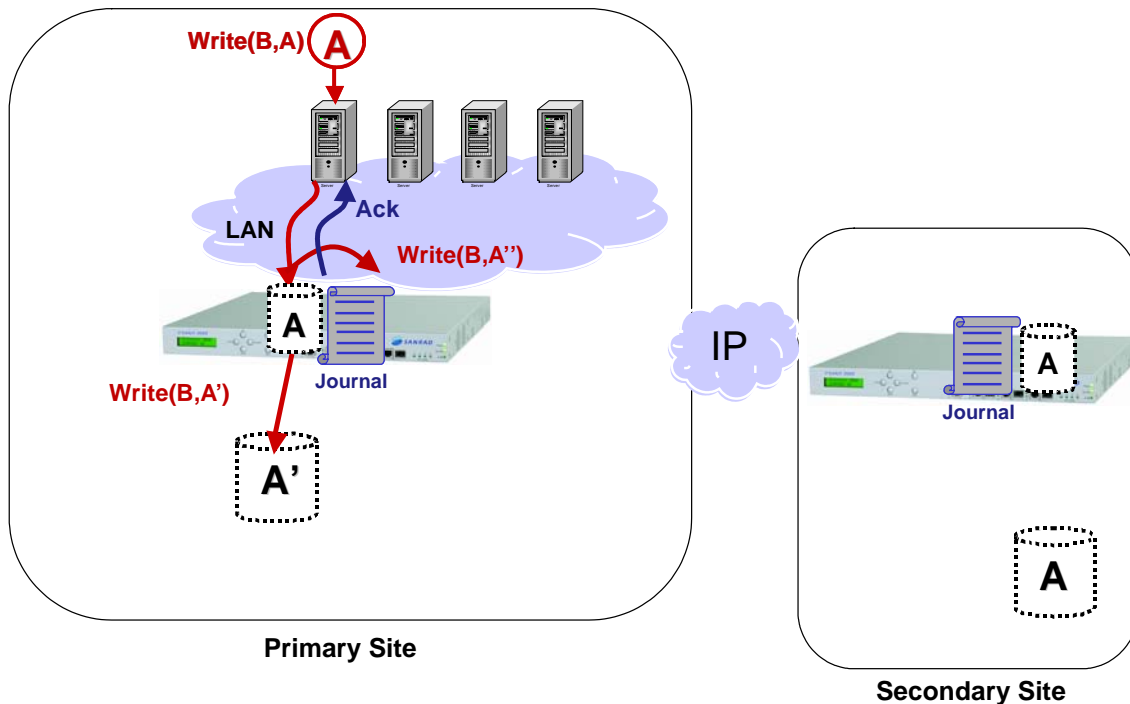


Figure 9. Synchronous Replication Method

## Asynchronous Replication

With asynchronous replication, every write operation is written to the primary volume. However, instead of continuing on to travel to the secondary site, the data is written in the primary journal then a write acknowledge is sent back to the server. See Figure 10, page 15.



**Figure 10. First Step of Asynchronous Replication Method**

At a configured time according to policy, a Point in Time snapshot (PiT) is created of the primary journal and the PiT data is pushed to the secondary site journal. When this PiT write operation is finished, an acknowledgement is sent to the V-Switch in the primary site. At the secondary site, after the entire PiT has been transferred, the PiT is merged with secondary volume. See Figure 11, page 16. Because the PiT data merge does not begin until the complete PiT is successfully transferred to the secondary site there is no potential for data loss in the secondary volume if network conditions interfere with the successful transfer of a PiT to the secondary site. The volume copy from the previous successful transfer remains intact. If your network does not have the bandwidth needed for synchronous replication or the latency between sites is too high to support synchronous replication, asynchronous replication is the solution.

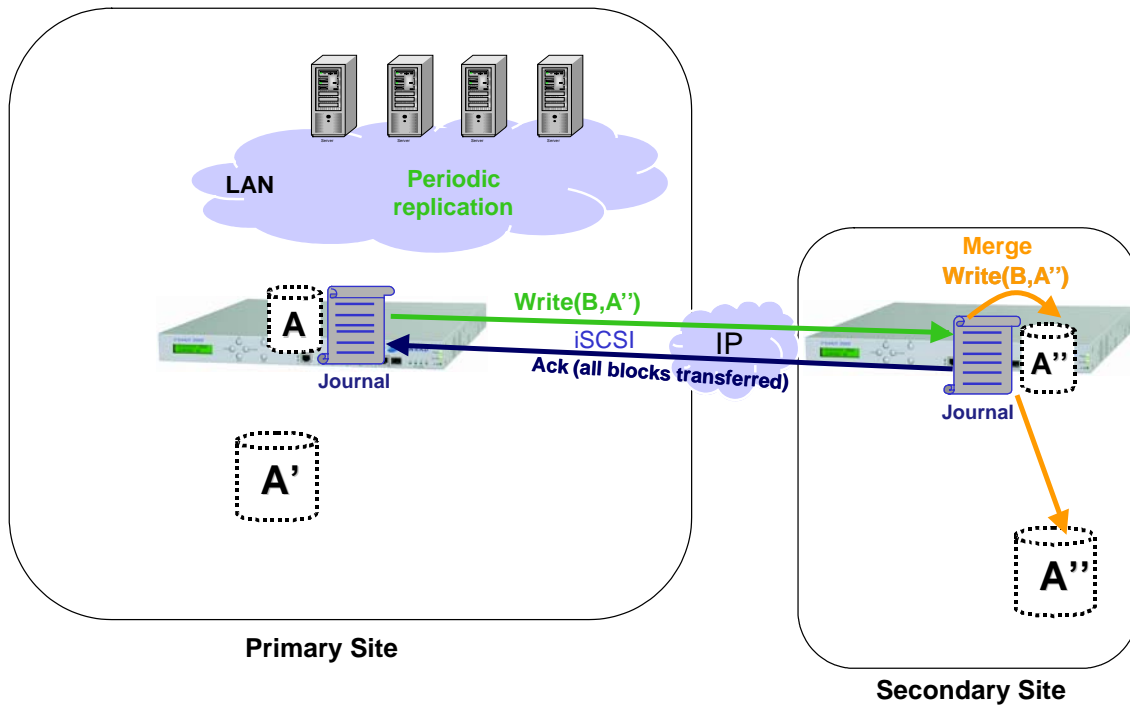


Figure 11. Second Step of Asynchronous Replication Method

*Replicate data independent of the bandwidth to the remote site to reduce the cost of using high-speed connections.*

## PLANNED FAILOVER & FALLBACK

A planned failover can be used when a site must go down for primary maintenance or for site relocation. The unique feature of a planned failover is its ability to continue to keep a journal of data changes, eliminating the need for an initial data resynchronization when falling back to the primary site.

SANRAD's Failover Wizard guides you through the failover process. First you need to deactivate the necessary applications at the primary site. Once all write operations are stopped and all caches are flushed, PiTs are taken of all GDR volumes and are replicated to the secondary site. When all data has been transferred, you can switch to the secondary site. The primary site is now available for maintenance.

SANRAD's Failover Wizard continues at the secondary site. At the secondary site the PiTs are merged with the secondary volumes for the most up-to-the-minute data copies. Once all of the data is replicated, the secondary volumes are exposed automatically to their configured hosts. Note that the secondary site continues to keep a journal volume. This journal volume will be instrumental in quickly restoring the primary site.

SANRAD's Fallback Wizard guides you through the fallback process. Just as when you failed over the primary site, you need to deactivate the necessary applications at the secondary site. Once all write operations are stopped and all caches are flushed, PiTs of all GDR volumes are replicated to the primary site. When all data has been transferred, you can switch back to the primary site.

At the primary site the PiTs are merged with the primary volumes. Note that an initial resynchronization of all data is not necessary. Only the PiT data needs to be merged with the primary site data.

Once all of the data is replicated, you can reactivate all applications.

---

***The unique feature of a planned failover is its ability to continue to keep a journal of data changes, eliminating the need for an initial data resynchronisation after planned fallback.***

## UNPLANNED FAILOVER & FALLBACK

A disaster is some event that will shut down a network, leaving its data inaccessible to the people for whom it was intended. A disaster can be caused by several different elements that fall under the categories of natural, man-made or regional disasters.

In an unplanned failover there is nothing to do immediately at the primary site. You can begin working straight in the secondary site. SANRAD's Failover Wizard will guide you through the failover process. Switch the secondary volumes to primary volumes. You can now expose the secondary volumes to their configured hosts and activate any applications that need to be activated. Note that the secondary site does not continue to keep a journal volume of data changes. The secondary site functions as a regular site without GDR capabilities.

When using asynchronous replication, there will be a potential possibility to lose some of the data that was not replicated to the secondary site upon disaster. With SANRAD GDR solution you may keep the data loss to a minimum or completely eliminate it by using optimal replication policies.

Depending on the level of disaster, the need for initial fallback configurations at the primary will vary. SANRAD's Fallback Wizard will guide you through the fallback process. If equipment and virtual volumes are still intact there is no initial reconfiguration that must be done at the primary site. If the site suffered hardware or software damage, primary volumes, volume hierarchies, targets, DR pairs and consistency groups must be reconfigured from the ground up. If volume names have changed you must update volume pair relations at the secondary site.

Once both sites are properly configured, you can begin initial synchronization from the secondary site to the restored primary site of all volume pairs using either online or offline copy. In essence, the secondary site is acting temporarily as a primary site replicating data to the primary site.

Once the initial sync is completed, you can deactivate applications on the secondary site and failover the secondary site back to the primary site. Once all write operations are stopped and all caches are flushed, PiTs are taken of all GDR volumes and replicated to the primary site. When all PiTs have been transferred, you can switch back to the primary site.

You can now reactivate all applications.

---

***Failover applications to the remote site or  
activate remote application servers in the event  
of primary site outage.***

## CONCLUSION

SANRAD's GDR solution lives at the network level. GDR is different than host-based solutions. Host-based solutions are OS-dependent and require an agent on each host. SANRAD's GDR solution is OS agnostic and easily scalable. There is no need to purchase additional data replication software or licenses as your SAN grows and, with unlimited numbers of volumes, servers and storage capacity, SANRAD's GDR solution provides a significant ROI. Unlike storage-based solutions, SANRAD's GDR is storage agnostic, eliminating vendor lock-in and enabling the continued use of all currently operating storage. Furthermore, this enables the use of low-cost SATA disks at the secondary site even when high-cost enterprise-class RAID subsystems are used at the primary site.

SANRAD's GDR solution replicates data at the block level. This allows a single block change to only require replication of the single block to the secondary site, eliminating bandwidth intensive full file replication necessary with host-based solutions. Because only block-level changes are replicated, the replication time is faster, resulting in more up-to-the minute data copies at the secondary site, which provides a better RPO in the event of a site failover.

SANRAD's GDR solution provides complete coverage for all available network bandwidths and tolerable latencies by offering both synchronous replication, where high bandwidth is available or RPO must be zero, as well as asynchronous replication, where low bandwidth would create too high an application performance penalty.

Businesses recognize that remote data replication, remote backup, site failover and disaster recovery are essential to their survival. SANRAD GDR provides a solution that is cost effective, simple to use and easy to scale enabling business to meet their needs for 24 x 7 data availability and business continuance.