

Virtual Disaster Recovery

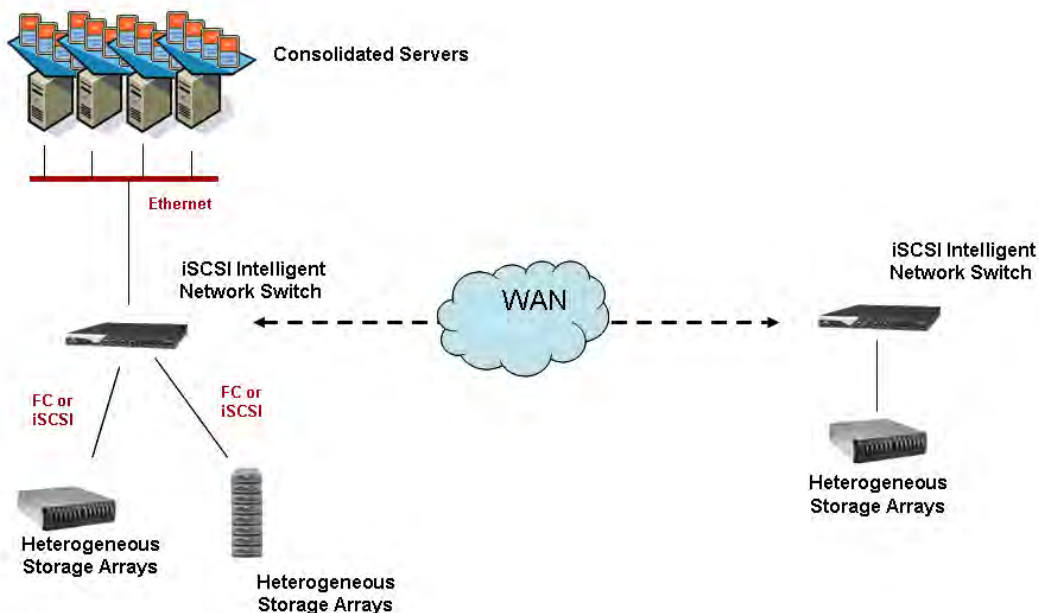
IT managers live in a world where disaster could strike in different forms: natural calamities, massive power outages, or a complete site failure. These failures can result in an entire company being unable to run its daily business operations or provide critical services to internal and external customers. If the only copy of important corporate data resided in the failure site, the consequences could be fatal.

Global Data Replication or GDR is the accepted best practice for creating remote copies of corporate data.

Though disaster recovery (DR) may be top-of-mind for IT and business executives, many may not be aware of the new, important considerations required when implementing DR in a virtual server environment.

This article highlights best practices in using the new features of VMware's ESX Version 3.5 and the SANRAD intelligent V-Switch to create seamless in-band or out-of-band GDR that will enable not only data replication but also reinstating the virtual machines in a consistent state at a remote site.

Figure 1
In-band GDR




In-Band versus Out-of-Band GDR

The storage industry has traditionally required the IT manager implementing GDR to select between two methods to achieve disaster recovery. In-band GDR (shown in Figure 1) calls for placing a SANRAD intelligent switch in the data path from the servers to the storage. The intelligent switch splits the data requiring replication and creates two copies. One copy continues to local storage. The second copy is transferred to a remote site over the WAN in one of two ways. Under out-of-band GDR (presented in Figure 2) the data path of the replicated data is separate from the main data path. Regular read and writes to the data are done through an "operations" data path, and the replicated data is transferred to remote location using a different "replication" data path.

Two main advantages are offered to the user by in-band replication. First, by replicating the data at the network layer, the user gains efficiency and performance. This holds because the network layer is doing what networks do best, sending data to where it is needed with minimal latency and performance impact. Second, by performing the replication on a central network, there is no need to modify the operation of the servers or install any additional software on them. The servers continue to function as usual and the replication effort is transparent to all of the servers.

Figure 2
Out-of-band GDR



Out-of-band replication has its own distinct advantages. Because the replication uses dedicated lines, the IT manager does not need to use the same storage networking technology for the main data path and the replication data path. For example with out-of-band replication the user can select Fibre Channel or NAS connectivity for the main data path, while using a SANRAD iSCSI intelligent switch for the Global Data Replication. This allows optimizing the budgetary investment by creating tiers of storage and networking investment to fit the actual requirements. Also, when

architected correctly, out-of-band replication can be designed to minimize the impact on the main data path. However, up until recently, out-of-band replication required installing software agents on all servers running applications with data to be replicated. These agents needed maintenance and updating and inevitably affected the overall server performance.


The issue of agent management is even aggravated when the servers are virtual machines. With virtual machine portability and new technologies such as VMware's storage V-motion, it becomes more and more difficult to track the location of the data to be replicated and the physical location of the virtual machines. If the agents are run within the virtual machines, one ends up having multiple storage replication agents running on the same physical server all needing maintenance and effecting the overall server performance by essentially all doing the same thing.

One of the factors that made the choice of replication technology critical to the organization, was that up until recently, the IT manager needed to make an up front choice, committing his infrastructure to one of the alternatives. Few systems supported the capability to do both in-band and out-of-band replication for different virtual machines on the same servers concurrently, and the investment in either porting the storage technology or installing all of the agents was a formidable effort not easily reversed. Fortunately, recent advances in both server virtualization technology and intelligent network switches from SANRAD have lifted these restrictions, and enable IT managers complete flexibility in both the choice of technology and complete portability between them. The next section explains these benefits.

The Freedom to Choose – Agentless Out of Band GDR

Intelligent network switches from SANRAD in combination with VMware's new storage features in version 3.5, allow for the first time, to perform agent less out-of-band GDR, with complete freedom to choose what volumes to replicate on which virtual machines. To explain how this works, one needs a short dive into the new multiple snapshot capabilities of VMware's ESX version 3.5. Figure 3 presents the file structure of a typical multiple snapshot on VMFS. When the first snapshot is taken ESX freezes the "original" vmdk file containing the virtual machine's volume at its state during the time the snapshot was taken. It then opens a new vmdk file "snapshot 1". All **new** data from the virtual machine is stored on this new file, which now contains the changes between the time the snapshot 1 was taken and the present. The file "snapshot 1" continues to grow until the point in time when a second snapshot is taken ("snapshot 2"). At this time ESX server freezes the "snapshot 1" file, and starts writing **new** data to a new "snapshot 2" file. The combination of the "original" vmdk file and the now frozen "snapshot 1" file represents the state of the data in the system at the point in time snapshot 2 was taken. This process continues with each snapshot. At any point in time, only the last snapshot file is written to and the data in the (now frozen) preceding files contains the state of the data in the system at the different points in time the snapshots were taken. What makes these files so important is that they can serve as excellent point in time recovery points for global data replication.

Figure 3
Multiple Snapshot Files



By employing a simple script, using SANRAD's scripting capability with V-Script, one can combine the power of the storage capabilities of ESX version 3.5 with a SANRAD Intelligent Network Switch to provide agent less out-of-band GDR with complete crash consistent recovery for virtual machines.

The script itself is a surprisingly simple three step process (as illustrated in Figure 2):

- In Step A, VMware's virtual center's API is used to quiesce the selected virtual machine and activate VMFS snapshot.
- In Step B, virtual center's API is used to copy the now frozen snapshot file to out-of-band storage managed by the intelligent network switch (Note that this does not need to be a physical copy, it can also be achieved by "exposing" a snapshot of the volume to the SANRAD intelligent network switch).
- In Step C the script instructs the intelligent network switch to replicate the data of the WAN to a remote location.

Thus in three simple steps using SANRAD intelligent network switches and scripting capability the IT manager achieved out-of-band GDR with no need to employ agents on any of the servers.

The Best of Both Worlds

GDR is a key component in the disaster recovery strategy of any enterprise. The combination of ESX Version 3.5 and SANRAD intelligent networks switches provides storage managers with complete flexibility in the implementation of GDR.

The storage manager can choose either in band or out- of-band GDR for any virtual machine, and even combine the two methods for different virtual machines running on the same server. The combination allows for complete flexibility to optimize the solution to the specific virtual machine requirements, as well as easy migrations between the two methods. In one example if a user wanted to implement GDR out of band (for minimal impact) as a first step, and later migrate to in-band GDR for the extra efficiency, SANRAD's architecture provides the perfect flexibility to do so.

Thus, the dynamic innovation in both server virtualization and SANRAD intelligent network switches delivers complete end-to-end disaster recover solutions to fit the needs of each and every corporation.