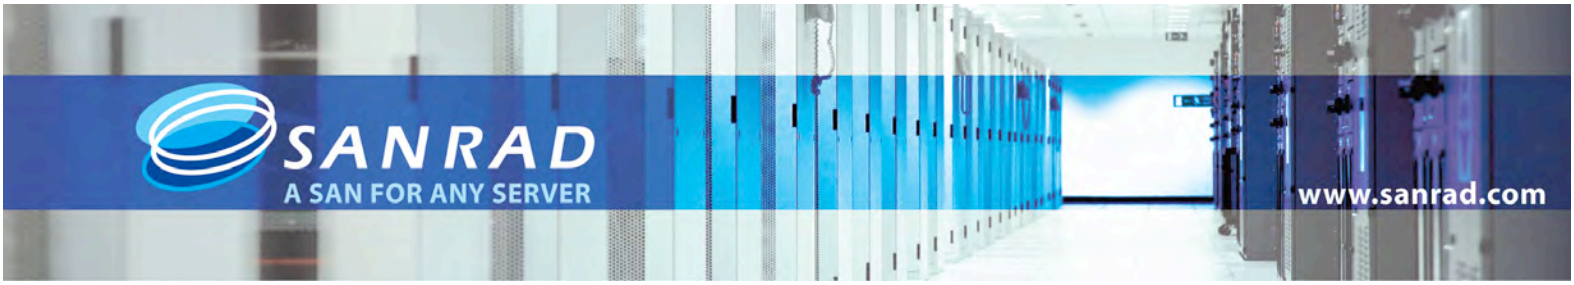

Server and Storage Virtualization:

A Complete Solution

SANRAD White Paper

Server and Storage Virtualization: A Complete Solution

A SANRAD Whitepaper

The Benefits of Combining Server and Storage Virtualization

Server virtualization is an important step toward improving overall IT efficiency. Virtual machine technology reduces the complexity and management of disparate server hardware and OS platforms. Nonetheless, server virtualization is only one component in a truly virtual enterprise infrastructure. Another critical component is storage virtualization.

Similar to server virtualization, storage virtualization creates a logical layer of storage from physical storage devices. The full benefits of server virtualization can only be realized in combination with a virtual storage layer that works in conjunction with and complements the virtual server layer. For example, dynamic virtual machine failover, a key benefit of leading server virtualization solutions, can facilitate disaster recovery. It is incomplete, however, without dynamic storage failover, a storage virtualization feature. Implementing storage virtualization, in other words, extends the benefits of an investment in server virtualization and builds upon them, providing simplified storage management, improved storage utilization and application performance, a bullet-proof disaster recovery solution, and a diminished need for proprietary vendor solutions.

This whitepaper describes what to look for in a storage virtualization solution that best serves your server virtualization project, and how to make sure the combination offers you a complete solution for your IT Needs.

Business Drivers for Storage Virtualization

The business drivers for storage virtualization are much the same as those for server virtualization. CIOs and IT managers must cope with shrinking IT budgets and growing client demands. They must simultaneously improve asset utilization, use IT resources more efficiently, ensure business continuity and become more agile. In addition, they are faced with ever-mounting constraints on power, cooling and space.

Storage virtualization provides companies with tools to address the underutilization of resources and the poor economics of silo-based storage, as well as the flexibility to respond to changing business requirements. In a storage virtualized environment, organizations achieve the full benefits of consolidation, improved resource usage and comprehensive disaster recovery. Storage virtualization also dramatically reduces power and cooling costs.

In a recent survey ESG asked IT manager what benefits they expect from increasing the use of networked storage for server virtualization. As can be seen the primary three benefits include virtual machine mobility, cost effective disaster recovery, and increased uptime and availability. The choice of the right storage virtualization to complement

your server consolidation project will significantly impact your capability to support these business goals.

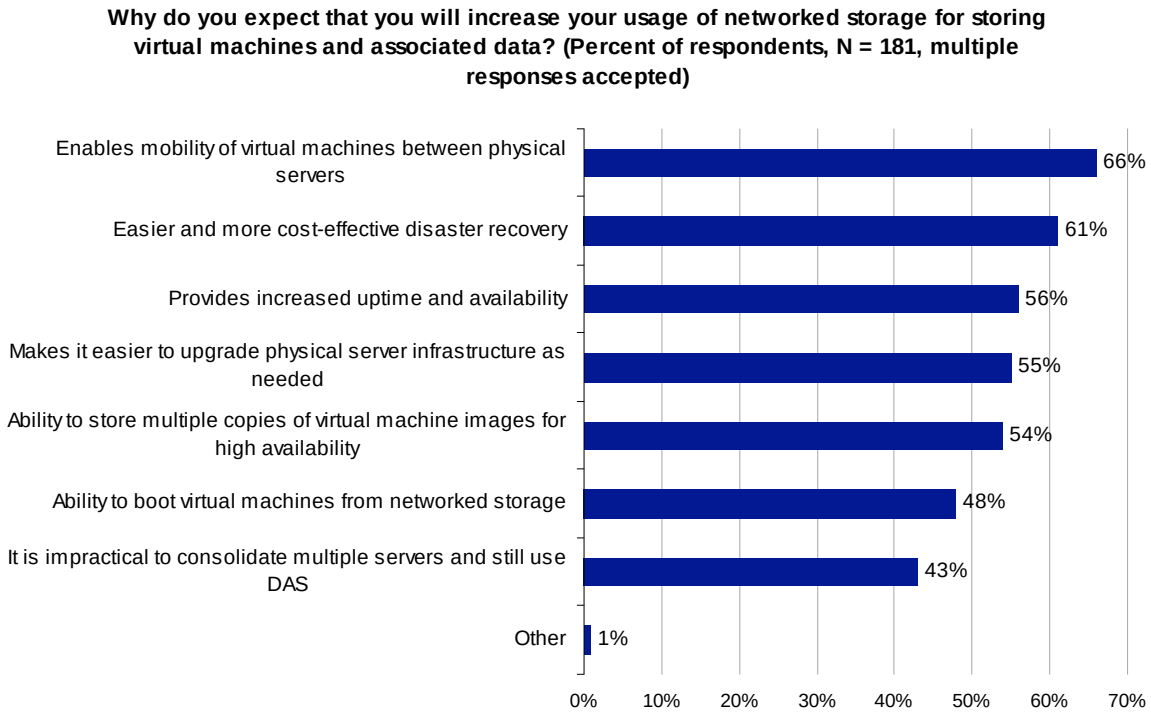

Source: *ESG Research, The Impact of Server Virtualization on Storage, December 2007*

Different Approaches to Storage Virtualization

IT managers have different solution alternatives for implementing a virtual storage infrastructure to complement their virtual server environment. Each of these has its merits and drawbacks, as indicated by the table below.

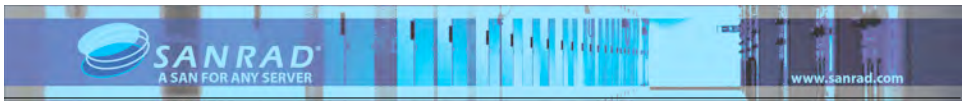

Different Approaches to Disaster Recovery

	Host-Based	Storage-Based	Network-Based
Pros	Storage independent Specialized SW	Dedicated vendor Array-specific Complements server virtualization	Storage and server OS independent Single license covers all storage Complements server virtualization
Cons	Multiple licenses Server performance	Vendor lock-in Multiple licenses	Solutions vary in functionality
Best Fit	Multi-vendor storage environment Familiarity with specific replication application	Single storage vendor preference No need for technology upgrade	Multi-vendor storage environment FC and iSCSI storage management Live data migration

Host-based solutions are storage agnostic, providing IT managers complete freedom to choose any storage to match the different needs of the enterprise. Vendors offer specialized application software, such as backup and replication, which are often rich in features to address a broad range of needs.

However, because agents and/or software must be loaded on each server, costs can be high and administration is more complex. Be careful when evaluating pricing schemes from different vendors, as they may be convoluted and challenging to predict if based on capacity. Another tradeoff is that host-based solutions consume server resources and can impact overall server performance. A host-based solution can be appropriate when IT managers need a multi-vendor storage infrastructure or have a legacy investment or internal expertise in a specific host-based application.

A storage-based alternative provides the benefit of an integrated solution from a dedicated storage vendor. Most storage vendors have also tailored their products to complement server virtualization and utilize key features such as virtual machine storage failover. Some enterprises may also have a long-standing business relationship with a particular storage vendor, and in these cases, a storage solution may have a relevant fit.

One drawback with a storage-based approach is that IT managers become locked in to a specific storage vendor. Some storage vendors have compatibility restrictions within their own storage array product line, making technology upgrades and data migration potentially expensive. When investigating storage alternatives, IT managers should pay attention to the total cost of ownership, since the cost of future license fees and support contracts will impact expenses longer term. A storage-based approach can make sense

when IT managers have a strong relationship with a particular storage vendor or in cases where heterogeneous storage is not a major requirement.

A network-based approach with intelligent switches has the advantage of being both storage and server operating system (OS) independent. IT managers gain the flexibility to choose the storage which best meets their cost, performance and reliability needs in different situations. For example, in the data center, IT managers may want to deploy a high-end storage array but in remote branches implement low-cost SATA storage. Because the intelligent switch sits at the network layer, a single management interface and software license covers all storage. Server OS independence provides economical benefits (no additional infrastructure needed) as well as operational advantages (no server overhead, no disruption to server operations, live data mighration).

Network-based solutions offer varying degrees of embedded functionality. When investigating network-based solutions, it is important to validate how well the vendor integrates with virtual server environments. If the IT project requires data migration, it is wise to confirm whether the vendor provides data migration as part of the standard product.

The Intelligent Network Switch: Complete Storage Mobility

Intelligent network switches with built-in storage management services, such as SANRAD's V-Switch, have changed the virtual storage paradigm from a model of expensive proprietary vendor lock-in to one of low-cost open support for total storage flexibility. Open storage virtualization enables IT managers to improve overall storage utilization by allowing capacity from any storage array to be combined in centrally managed, virtual storage pools. IT managers can dynamically reduce capacity for applications that are not growing and reuse that capacity for those which are. This can eliminate the need to procure new storage or, at a minimum, delay acquisition.

Intelligent network switches also enable organizations to easily link storage hardware and software solutions from multiple storage vendors, removing the burden of proprietary vendor lock-in and simplifying tasks such as data replication, mirroring and data migration. For example, when a storage network is comprised of disparate systems, even from the same vendor, data migration is time-consuming and application-disruptive. In a typical scenario, the application is taken offline, data is moved to tape and then restored. Intelligent integrated network switches permit data to be migrated in real time, from any storage array to any other storage array, without taking applications offline. Storage virtualization utilizing intelligent switches, such as the V-Switch, provides non-disruptive online migration that eliminates downtime and greatly reduces administrators' involvement in data or server migration.

High Availability Incorporating Storage Virtualization with Server Virtualization

When integrated with server virtualization, storage virtualization enables one storage system to fail over to another storage system with minimum disruption. However, with

most storage virtualization solutions the failover is not instantaneous and manual intervention is required. Therefore, the benefit of overall instantaneous system recovery is lost. SANRAD V-Switch, with its network-based intelligent switch architecture, as depicted in the illustration below, addresses this discrepancy between the server and

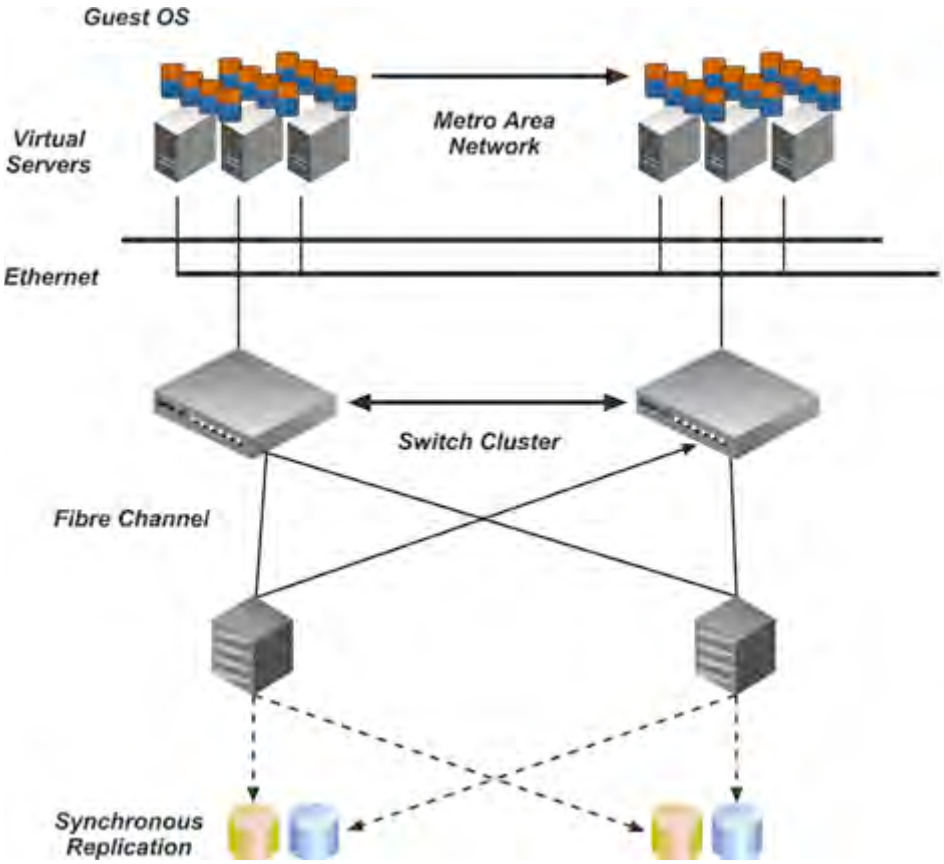

storage, providing instantaneous failover.

In the event of a complete disaster at one location, the virtual server there fails over to the server at the remote location without any interruption. If the primary storage system is affected by the same disaster, the other dynamically responds to the virtual server. This ensures zero hours RTO (recovery time objective) and zero hours RPO (recovery point objective). Thus, the new generation of open virtualization solutions delivered via the V-Switch provides enterprise-class high availability without the expenses imposed by proprietary storage architectures.

Snapshots in a Server Consolidation Environment

High availability functionality alone, however, can not completely guard your server virtualization environment from all calamities. For example if your environment is attacked by a malicious virus, synchronous replication will only copy the virus to both

sites and the secondary site will also be infected. To protect against attacks of this type, one must use the storage virtualization capability of creating "Snapshots" of the content of a given volume at a certain point in time.

Snapshots essentially provide an image of a virtual machine "frozen" to the instant the snapshot was taken. What makes snapshots such an important capability is the wide array of critical functionality they provide both to traditional servers and to virtual servers.

Using snapshots, for example, an IT manager can "keep" copies of the server configuration at given times in the day. If a virus attacks the system (or a software upgrade ruins the configuration) the user can revert to the last snapshot taken before the attack. Essentially, the IT manager simply "rolls back" to the system configuration before the virus attack. Another important functionality snapshots provide is the capability to perform consistent backups of data without interruption to a working application. The IT manager can create a snapshot of a volume and backup the data from the snapshot while the application keeps running on the original volume. In this way a backup copy of the data, consistent to a point in time, is created without the need to take the system offline.

The advent of virtual servers also created a third high value snapshot usage: snapshots of "golden masters". IT managers of virtual server environment often need to keep "golden masters" of virtual servers containing a given configuration and certain installed applications. Under certain circumstance one or more instances of these virtual servers need to be created to support running the applications without changing the "golden master" itself. Snapshots provide the best practice for providing this exact functionality.

The Benefit of Integrating Server and Storage Virtualization Snapshots

Today snapshots of virtual machine data can be taken at the server virtualization layer, at the storage device layer, or at the network layer. Using snapshot functionality at the server virtualization layer certainly has its advantages, but also has some drawbacks that must be taken into account. The main advantage is the ability to use the server virtualization capability of quiescing the virtual machines. Quiescing a virtual machine assures that the data in the snapshot is crash consistent. This is especially important when the snapshot is to be used for backup and recovery. Another important advantage of using the server virtualization layer for snapshots is that the virtual server center has the best knowledge of where the files owned by a certain virtual machine reside. For example VMware's VirtualCenter can be queried for the location of the vmdk files of specific virtual machines.

However, there are also some drawbacks associated with using the server virtualization layer for snapshots. For example, server virtualization snapshots do not support all the data accessible to a virtual machine. For example if you are running VMware and your virtual machines are accessing raw disks or RDM physical mode drives, you can not take a snapshot of their data via the VirtualCenter. Another critical drawback is that taking a snapshot is is highly CPU and IO intensive. Using the server to perform many snapshots can easily stress the performance load on the servers and competes with the resources of the virtual machines themselves.

In comparison, taking snapshots within a storage array is much more efficient and will

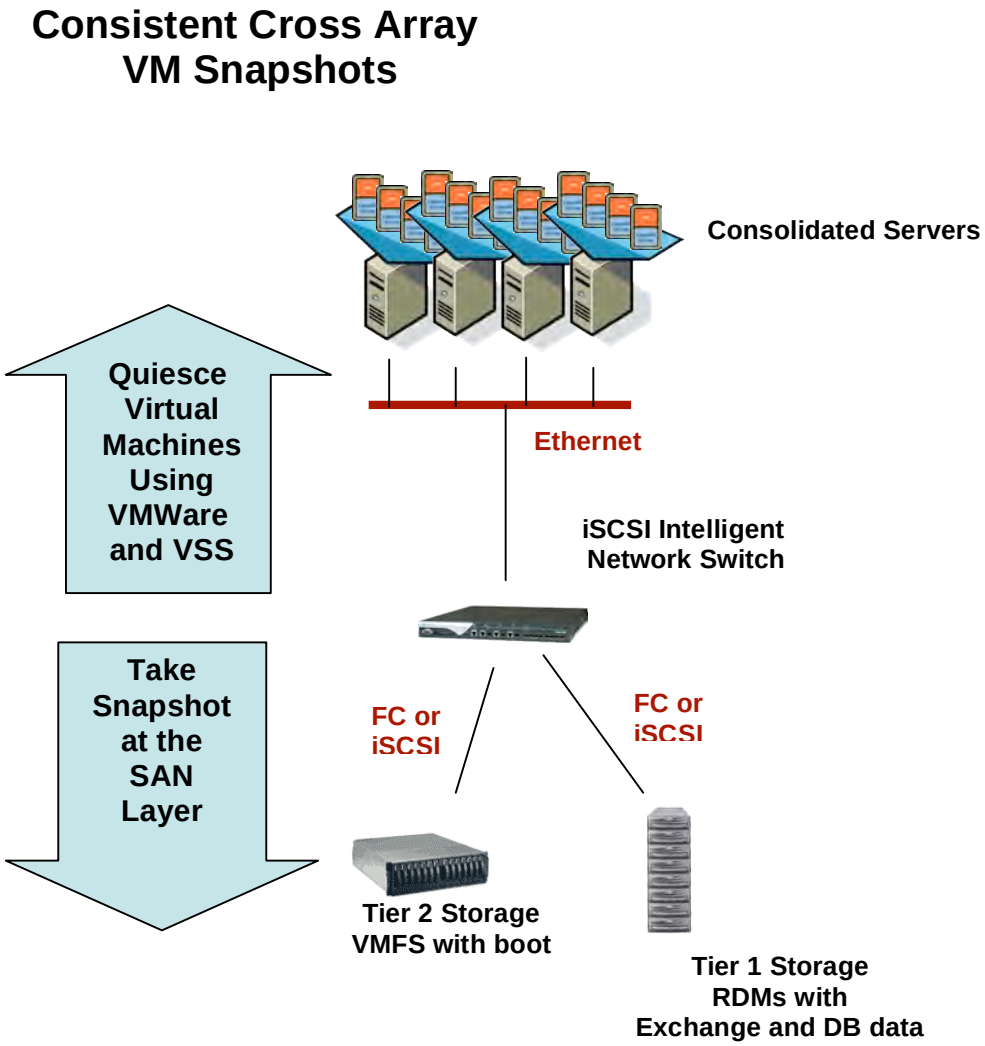

afford the user much higher performance. But storage array snapshots also have their drawbacks. In many cases, such as storage high availability scenarios and resource tiering, the data of a virtual machine may be spread across several storage arrays. In such cases using one array's snapshot capability will create only a partial image of the state of a virtual machine. The same issue occurs when a virtual machine uses VMFS for its boot disk but an RDM for its application data. A user will not be able to reconstruct the state of the virtual machine at a given point in time from the partial data collected in the

array snapshot. Furthermore, array-based snapshots in themselves can not grantee a state consistent snapshot. Someone has to make sure that the data in the array at the time the snapshot was taken is crash consistent.

Snapshots at the Intelligent Network Layer – The Best of Both Worlds

The advent of intelligent network switches, such as SANRAD's V-Switch, with snapshot capabilities enables users to get the best of both worlds. For the first time users are able to obtain cross-array state-consistent snapshots of their virtual machines. On top of this, by managing the snapshots at the SAN network layer through the intelligent switch, rather than at the server layer, the high efficiency of storage- based snapshots is retained, and the server resources are not strained by an extra effort.

Consider the example presented in the above figure. In this example a cluster of virtual servers is connected to a SAN using an iSCSI intelligent network switch. To maximize their return on investment, the SAN managers are using two tiers of storage. The data from Microsoft Exchange servers and databases are stored on tier 1 storage as RDMs. VMFS along with boot volumes for the virtual machines are stored in tier 2 storage on the SAN.
The SAN manager is required to provide consistent snapshots of the virtual machines and their data at given intervals. Under this architecture the task becomes easy. By using the scripting capabilities of the intelligent network switch (such as SANRAD's V-Script capabilities) the SAN manager runs a simple two step script. In step 1 the script calls the virtual machine (using the VMware scripting and VSS) to quiesce both the applications and the virtual machines themselves. In step 2 the script activates snapshots of the virtual machine data on both tiers of the storage.

The goal has now been achieved. The SAN manager has created a cross-array snapshot of a virtual machine at a consistent recovery point in time.

Easy Cost Effective Disaster Recovery

High availability in combination with snapshots will protect an environment from nearly all local or campus level failures. However, in our world one must always be prepared for the absolute worst: a comprehensive site failure. Natural calamities, acts of terrorism, or even massive utility outages can all cause an entire site or campus to discontinue services. If the affected location housed the only copies of important corporate data, the results can be disastrous for a company, no matter what the size. Global Data Replication or GDR is the accepted best practice for creating remote copies of corporate data. SANRAD V-Switch in combination with server virtualization environments such as VMware's ESX Version 3.5 enables seamless, cost-effective in-band or out-of-band GDR that allows you not only to replicate the data but also to reinstate the virtual machines in a consistent state at the remote site.

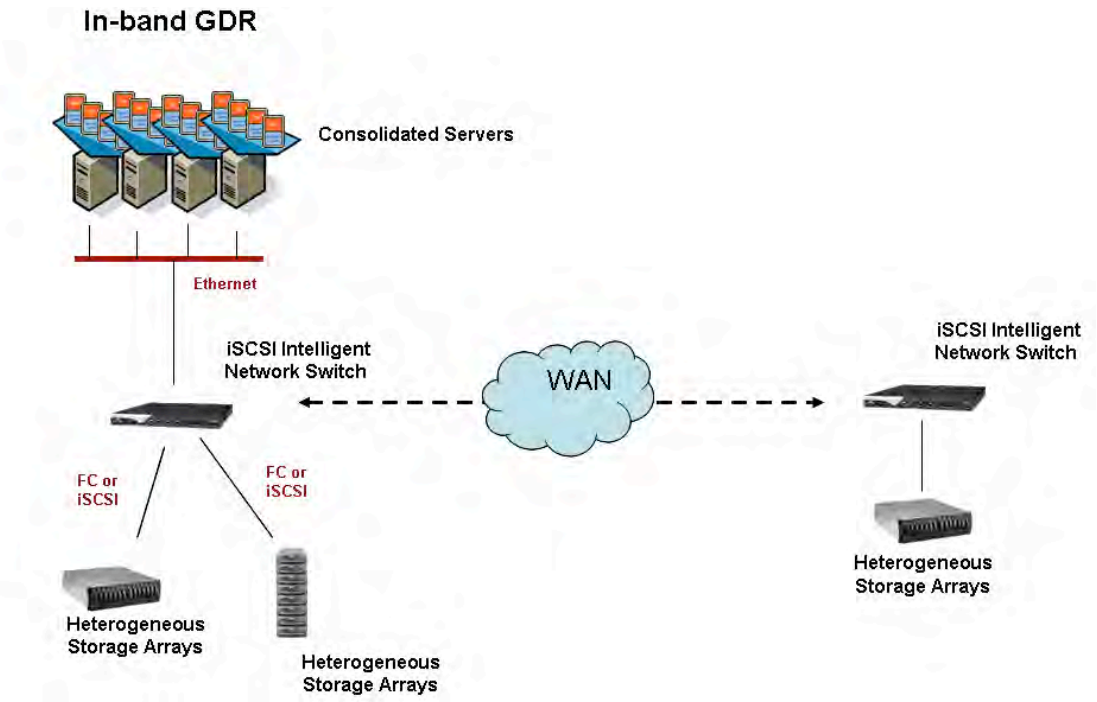

In-Band versus Out-of-Band GDR

The storage industry has traditionally required the IT manager implementing GDR to choose between two methods to achieve disaster recovery. In-band GDR (shown in the above figure) calls for placing an intelligent switch such as SANRAD V-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 on to local storage. The second copy is transferred to a remote site over the WAN. Under out-of-band GDR (presented in the figure below) 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.

In-band replication offers users two main advantages. First, by replicating the data at the network layer, the user gains efficiency and performance. This is possible 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 location, there is no need to modify the operation of the servers or to install any additional software on them. The servers are then able to continue functioning as usual while the replication effort continues transparently.

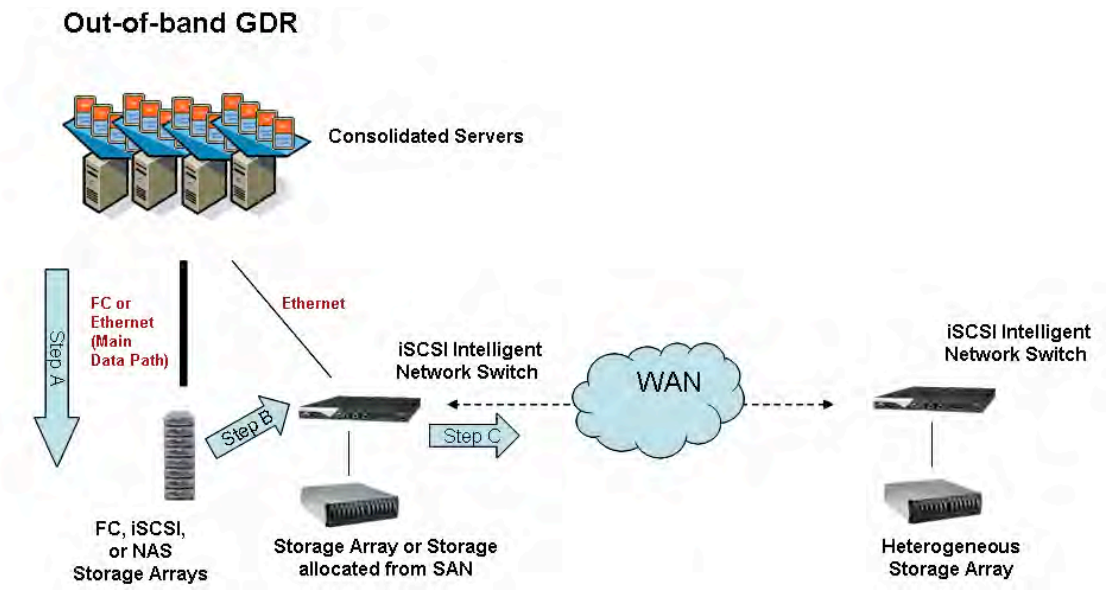

Out-of-band replication has its own advantages. Because the replication uses dedicated lines, the IT manager is not required to use the same storage networking technology for the main data path and the replication data path. For example, with out-of-band replication the user can select Fibre Channel or NAS connectivity for the main data path, while using iSCSI for the Global Data Replication. This allows optimizing the budgetary investment by tiering the 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 overall server performance.

The issue of agent management is further 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 running multiple storage replication agents on the same physical server. These all require maintenance and can affect the overall server performance by essentially duplicating efforts.

One factor that makes the choice of replication technology critical to an organization is that, until recently, IT managers were forced to make an up front choice, committing their infrastructure to one of these 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. Furthermore, 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 (such as SANRAD V-Switch) have lifted these restrictions, and enable IT managers

complete flexibility in both the choice of technology and complete portability between them as demonstrated in the next section.

The Freedom to Choose – Agent-less Out-of-Band GDR

Out-of-band replication is useful when IT managers prefer to keep their NAS, iSCSI or Fibre Channel storage in the main data path. The key reasons for choosing this option would be minimal change to the existing data path.

SANRAD V-Switch, in combination with VMware's new storage features in version 3.5, for the first time allow IT managers to perform agentless out-of-band GDR, making the entire operation simple with no impact on the virtual servers. IT managers also have complete freedom to choose any volumes (from different storage arrays) to replicate on which virtual machines, thereby gaining true virtual storage mobility.

To explain how this works, one needs a short explanation of the new multiple snapshot capabilities in VMware's ESX version 3.5. The figure below 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 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, 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.

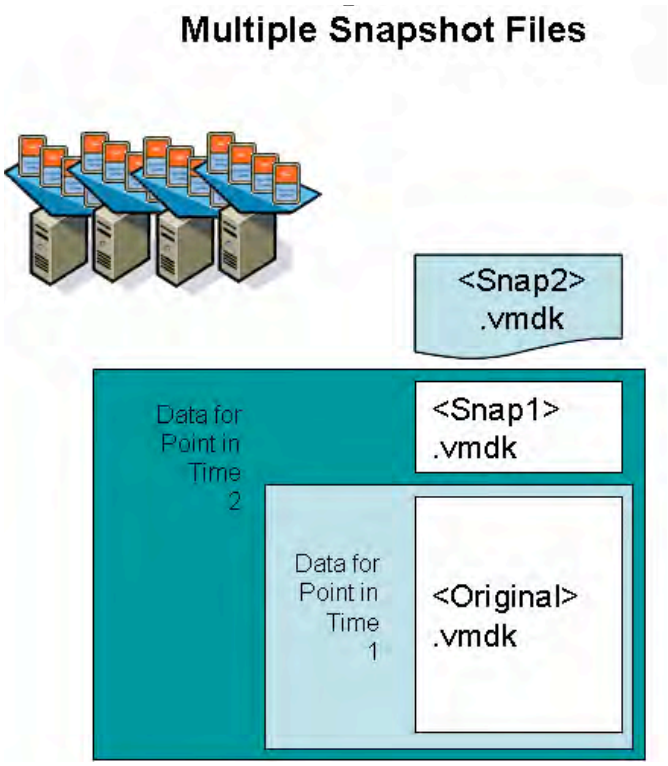

By employing a simple script (using scripting tools such as SANRAD's V-Script) one can combine the power of the storage capabilities of ESX version 3.5 with SANRAD's V-Switch to provide agent-less out-of-band GDR with complete crash consistent recovery for virtual machines.

The script itself is a simple three-step process:

- In Step A, VMware VirtualCenter's API is used to quiesce the selected virtual machine and activate a VMFS snapshot.
- In Step B, VirtualCenter's API is used to copy the now frozen snapshot file to out-of-band storage managed by the V-Switch.
- In Step C the script instructs the V-Switch to replicate the data over the WAN to a remote location.

Thus, in three easy steps, the IT manager achieved out-of-band GDR without employing agents on any of the servers.

The Benefits of GDR at the Intelligent Network Layer

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

The storage manager can choose either in-band or agent-less 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 migration between the two methods. In one example, a user wanted to implement GDR out of band (for minimal impact) as a first step, and later migrate to in-band GDR for extra efficiency, with the above architecture providing him the perfect flexibility to do so.

Thus, the dynamic innovation in both server virtualization and intelligent network switches delivers complete end-to-end disaster recovery solutions to fit the needs of enterprises everywhere.

Conclusion

The full benefits of server virtualization can only be attained by combining it with storage virtualization. The combination of the two enables the IT manager to provide a complete solution enabling a true reduction in TCO due to achieved hardware and energy efficiencies.

An IT manager must carefully select the best storage virtualization architecture to complement the planned server virtualization environment. A storage virtualization solution for server consolidation must assure the overall new environment is not only cost effective but can also protect the users against the disruptions that can affect IT infrastructures.

The combination of server virtualization with the storage virtualization provided by SANRAD's V-Switch uniquely addresses those exact needs. Cost effectiveness is achieved by iSCSI connectivity at the network layer along with complete flexibility in storage array connectivity and expansion. SANRAD V-Switch's network layer-based virtualization also allows achieving enterprise class high availability, snapshots, and global data replication tightly integrated to the complementary server virtualization offerings.