

Replicate with Integrity:

Protecting VMware Data to Ensure Fast Recovery, Business Continuity

By:

Gary Lamb

Senior Director of Data Center
Virtualization Practice
INX Inc.

and

Peter Eicher

Director of Product Marketing
FalconStor Software, Inc.



***Abstract:** As organizations increasingly adopt server virtualization, they have begun to leverage remote replication as a means of increasing system availability. This is a simpler, more cost-efficient alternative to using traditional physical servers for disaster recovery (DR). However, utilizing simple replication in a virtual server DR plan provides only “crash consistent” virtual disk images, which can affect the consistency of file systems or applications and ultimately create longer recovery cycles. In addition, managing and coordinating resources between sites can be complex and challenging. This document addresses these issues and discusses how technologies from FalconStor Software and VMware work together to enable organizations to overcome these limitations and deliver application-consistent replication for automated DR in a virtual environment.*

Introduction: Virtualization and Disaster Recovery (DR)

VMware Infrastructure 3 is the most advanced and commonly used enterprise server virtualization system for the x86 platform. As organizations become more familiar with server virtualization and the possibilities it enables, they have begun exploring replication to another virtualization server as a simpler and less expensive alternative to using traditional individual physical servers for DR and high availability. The portability provided by virtualization greatly simplifies restoration of replicated virtual machines, due to the elimination of hardware compatibility issues.

However, there are still complexities involved in virtual-to-virtual (V2V) DR. Coordinating the associations between storage volumes and virtual machines at the recovery site can be a complex and error-prone manual task. Furthermore, simple replication of virtual servers only provides “crash consistent” disk images akin to those of a system after a sudden power failure. Taking snapshots and replicating a LUN containing several virtual disks without the knowledge and cooperation of the operating systems and applications supported by those disks can foster inconsistent file systems and loss of application integrity in databases and other applications. This makes recovery very difficult and laborious, as system administrators must engage consistency checks and at times even rebuild entire systems.

FalconStor Software, the premier provider of TOTALLY Open™ network storage solutions, has teamed with VMware to provide a comprehensive solution for managing and recovering replicated virtual machines with 100% application-consistent transactional integrity. This solution coordinates replication to significantly accelerate recovery times while decreasing WAN bandwidth costs for virtual infrastructures.

Managing the Virtual Environment

As virtualized server environments grow, it becomes increasingly challenging to track server and storage resources between primary and recovery sites. The fact that a recovery site may often have

fewer physical servers and higher rates of virtual machines per server further complicates matters. Traditionally, IT departments have used detailed “runbooks” to document the steps and resources required for the recovery process. Since it can be very labor-intensive to maintaining runbooks, they are often out-of-date or seriously incomplete.

To simplify this process, VMware has released the VMware Site Recovery Manager (SRM), a DR management and automation solution for VMware Infrastructure. VMware SRM simplifies and automates the recovery process, maintaining coordination between virtual machines and storage across sites.

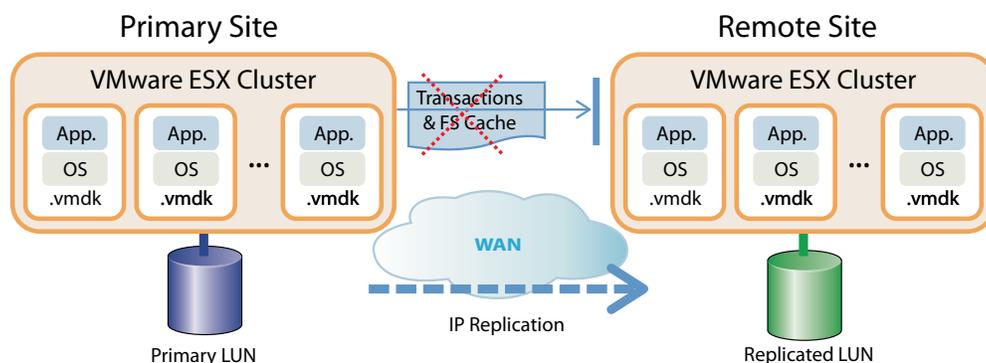
FalconStor technology integrates with VMware SRM through the FalconStor® Storage Replication Adapter (SRA), which provides a layer of software that communicates between VMware SRM and the underlying disk resources controlled by the FalconStor Network Storage Server (NSS) disk virtualization platform. Together, VMware SRM and FalconStor SRA ensure that all the virtualized applications at the recovery site connect to the proper storage resources, allowing business operations to resume quickly and reliably.

Virtual Infrastructure and DR Challenges

While VMware SRM eliminates the management challenges of V2V DR, there are still important data consistency issues to consider. A shared storage infrastructure, most commonly a SAN, is needed to implement many of the most compelling features of VMware Infrastructure 3. Consolidating the virtual servers on one or more networked storage arrays provides opportunities for cost-effective and simplified DR and business continuity through replication of the LUNs containing the virtual servers to a virtual infrastructure at a remote site.

In Figure 1, the virtual server disks replicated to the DR site are in a crash consistent state. The servers will most likely boot up; however, there is risk of file system and application corruption caused by the snapshot being taken without the applications being quiesced (i.e., in a state where no data is being written to disk) or the file system caches being flushed. Recovering these servers to an operational state

Figure 1: Replication without quiescence compromises transactional integrity



often requires file system and database consistency checks, importing of transactional logs, and, in some cases, rebuilding of applications.

In order to ensure consistently viable DR of virtual machines using replication, organizations must adopt a method of coordinating the quiescence of all of the applications and file systems concurrently with storage replication.

Scripting the coordination between an array and the supported applications on an individual physical server can be done, but it is not a trivial or simple task. Replication agents in the operating system make this task easier; however, most array replication agents are not virtualization-aware. Things get much more complex in the virtual world where several virtual servers share one or more LUNs. Coordinating the quiescence of all virtual machines and their applications concurrent with the snapshot process is difficult and can cause application downtime while disk I/O is suspended waiting for affected virtual machines to complete replication or reapply redo logs.

Once an organization opts to integrate systems and applications to quiesce applications and flush the file system cache along with the wide variety of potential disk configurations, it needs a comprehensive system for extending the built-in functionality of VMware technology by managing the coordination and replication of virtual machines. To accommodate these needs, FalconStor has developed a simplified and consolidated DR automation solution for virtual infrastructures.

The FalconStor Solution

The FalconStor Application Snapshot Director (ASD) for VMware coordinates between FalconStor NSS appliances and application-aware FalconStor Snapshot Agents in virtual machines to provide

an integrated solution for non-disruptive replication to remote sites with complete transactional and file system integrity.

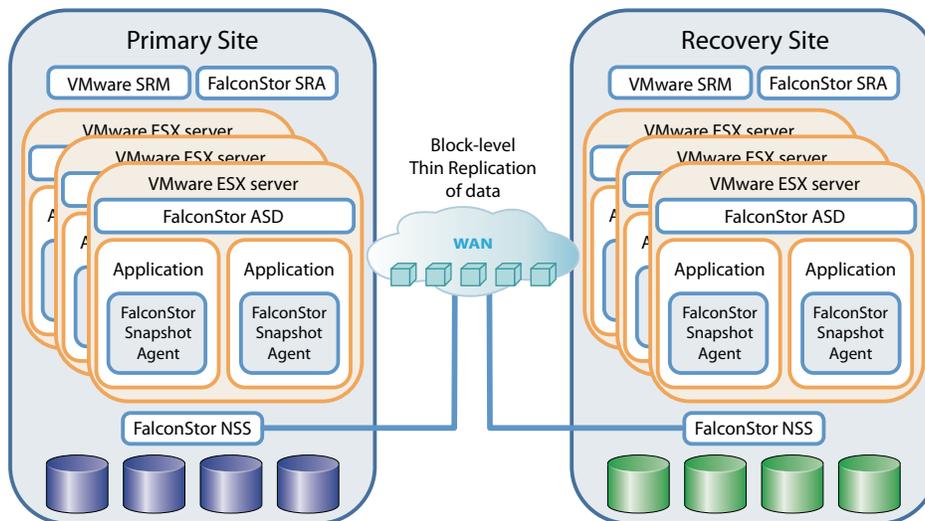
Replication can be configured for a variety of virtual infrastructure storage configurations. For ease of management, several virtual machines can be grouped on a single virtual file system. In order to take a snapshot of the file system, each virtual machine is quiesced, and a snapshot is taken of the underlying disk on which the virtual file system resides.

For virtual machines where I/O performance is critical, the ideal configuration is built using VMware ESX Raw Disk Mapping (RDM) in a physical compatibility mode mapped to the VMware ESX Server 3.5 console and directly to the guest operating systems. Leveraging physical compatibility, RDM ensures that the guest operating system is configured and replicated in almost exactly the same manner as a physical server.

In such a configuration, FalconStor NSS can notify the Snapshot Agent in the virtual machine to prepare for a snapshot by quiescing the application(s) and flushing the file system cache into the RDM disk for file system consistency. The snapshot is then replicated to another remote FalconStor NSS appliance, providing a transactionally consistent image of the environment at a known good point in time.

A snapshot can also be mounted to a local or remote machine as a FalconStor TimeView® image, which captures the underlying volume as it was when the snapshot was taken. These TimeView images can then be backed up utilizing existing backup software, completely offloading the backup processing from the production systems. This is an important enhancement to the functionality of VMware Infrastructure 3, which only processes one virtual machine at a time.

Figure 2: FalconStor data protection solutions in a VMware environment



TOTALLY Open Flexibility and WAN Cost Reduction

A typical enterprise solution will involve a clustered primary site and a single node DR site. For large enterprises, this is achieved using FalconStor technology on commodity x86 hardware, typically utilizing Fibre Channel (FC) disk storage at the primary site and lower-cost storage, such as SATA, at the DR site. In this scenario there are no limits on storage capacity, and support for up to 255 snapshots per virtualized LUN. The FalconStor ASD for VMware is licensed on a per-node basis, and replication can include encryption for WAN data.

Because FalconStor NSS can utilize any disk array, users have total freedom to deploy the storage that best fits their cost and performance needs, both at the primary and the recovery sites, without being locked in to a specific disk type.

One of the largest challenges associated with DR is the ongoing cost of bandwidth. Keeping multi-terabyte data centers in synch can consume huge amounts of WAN capacity, often resulting in limited DR plans that leave many applications behind because the budget cannot cover all of them. For this reason, FalconStor offers Thin Replication with patented FalconStor MicroScan™ technology, designed to use the least possible amount of bandwidth when shipping data bits over the wire.

FalconStor MicroScan is a data deduplication mechanism that eliminates any exaggerated block-level changes due to inefficiencies at the application and file system layer to minimize the amount of data transferred. As a result, only real changes at the disk-sector level (512-byte) are transferred. This helps reduce bandwidth and associated storage costs for DR.

The process employed by FalconStor MicroScan is more efficient than file system and disk array processes. Most file systems write a minimum of 4KB or 8KB of data whenever a file is changed, even if the change is as small as one character in a document. Many replication solutions will then send the entire update over the WAN because, from their point of view, this is the “minimum” amount of data that can be sent. Some disk array systems even process at a larger track level and will ship 16KB or even 32KB to accommodate the smallest changes.

In contrast, FalconStor MicroScan monitors data at the disk sector level. In this example, if a minor change is made, MicroScan looks at the 4-8KB of “new” data written by the file system and also looks “below” that level at the actual disk sectors, at a granularity of 512 bytes or 1 individual disk sector; the smallest possible write to a disk drive. If the change is limited to only one sector, then only 512 bytes are replicated. (Keep in mind that 4KB is eight times more data than 512 bytes.) By tracking data changes at the smallest possible disk sector level, MicroScan minimizes data transfer, resulting in WAN utilization rates that are 8 to 64 times more efficient than major array-based replication models.

Summary

Server virtualization is becoming pervasive and provides dramatic increases in flexibility and efficiency in the data center. A centralized and virtualized storage infrastructure is critical to realizing the full benefits of virtualization. The FalconStor solution set for VMware complements VMware Virtual Infrastructure 3 and VMware SRM to provide an unprecedented level of versatility, integration, and coordination among virtual storage, virtualized applications, and virtualization servers. In doing so, the FalconStor solution enables completely consistent, reliable replication of VMware storage and fast, secure recovery for uninterrupted processes and business continuity.

About FalconStor

FalconStor Software, Inc. (NASDAQ: FALC), the premier provider of TOTALLY Open™ Network Storage Solutions, delivers the most comprehensive suite of products for data protection and storage virtualization. Based on the award-winning IPStor® platform, products include the industry-leading Virtual Tape Library (VTL) with Single Instance Repository (SIR) for deduplication, Continuous Data Protector™ (CDP), Network Storage Server (NSS), and Replication option for disaster recovery and remote office protection. Our solutions are available from major OEMs and solution providers and are deployed by thousands of customers worldwide, from small businesses to Fortune 1000 enterprises.

For more information, visit www.falconstor.com or contact your local FalconStor representative.

Corporate Headquarters
USA
+1 631 777 5188
sales@falconstor.com

European Headquarters
France
+33 1 39 23 95 50
infoeurope@falconstor.com

Asia-Pacific Headquarters
Taiwan
+866 4 2259 1868
infoasia@falconstor.com

FalconStor
Software