



Replicate with Integrity: Protecting VMware Data to Ensure Fast Recovery, Business Continuity

WHITE PAPER



By Gary Lamb
Chief Technology Officer
AccessFlow, Inc.

Abstract: As organizations increasingly adopt server virtualization, they have begun to explore the notion of 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. However, utilizing simple replication in a virtual server disaster recovery plan provides only “crash consistent” virtual disk images, which can result in inconsistent file systems or loss of application consistency, and ultimately longer recovery cycles. This white paper addresses these issues and discusses how FalconStor technology enables organizations to overcome these limitations and deliver application-consistent replication in a virtual infrastructure environment for improved disaster recovery.

Introduction – Virtualization and Replication

VMware's Virtual Infrastructure 3 (VI3) is the most advanced and commonly used enterprise server virtualization system for the x86 platform. As users become more familiar with server virtualization and the possibilities it enables, they quickly begin exploring replication to another virtualization server as a simpler and less expensive alternative to using traditional individual physical servers for disaster recovery and high availability. The portability provided by virtualization greatly simplifies restoration of replicated virtual machines (due to the elimination of hardware compatibility issues). However, simple replication in a virtual server business continuity plan 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 result in inconsistent file systems and loss of application integrity in databases and other "complex" applications. This makes recovery very difficult and laborious, as systems administrators must engage consistency checks and may even need to rebuild entire systems.

FalconStor Software, Inc., the industry leader in proven data protection technology, provides a comprehensive solution for recovering replicated virtual machines with application-consistent transactional integrity. The FalconStor Application Snapshot Director for VMware (ASD) solution coordinates the quiescing of database applications and file system caches with the creation of snapshot(s) prior to data replication, resulting in complete application consistency in the replicated virtual machine images as well as in the remote copies.

The following document discusses the FalconStor solution in greater detail and illustrates how coordinated replication can significantly accelerate recovery times and increase the number of available recovery points in virtual infrastructures.

Virtual Infrastructure and Disaster Recovery Challenges

A shared storage infrastructure, most commonly a Storage Area Network (SAN), is needed to implement many of the most compelling features of VMware VI3. This shared storage requirement is driving first-time adoption of SAN technology in many mid-sized organizations. Consolidating the virtual servers on one or more networked storage arrays provides opportunities for cost-effective and simplified disaster recovery and business continuity through replication of the LUNs containing the virtual servers to a virtual infrastructure at a remote site. Conversely, the promise of "instant disaster recovery" has been known to drive the adoption of a virtual infrastructure.

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 or the file system caches being flushed. Recovering these servers to an operational state often requires file system and database consistency checks, importing of transactional logs, and, in some cases, rebuilding of applications.

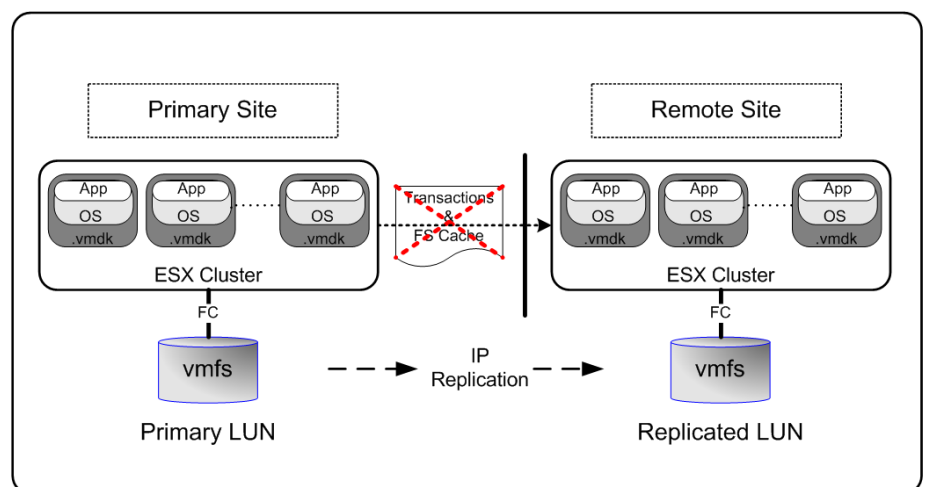


Figure 1: A traditional, standard virtual replication architecture, which results in "crash consistent" images.

In order to ensure consistently viable disaster recovery of virtual machines using replication, organizations must adopt a method of coordinating the quiescence of all of the applications and file systems concurrently with the 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 different virtual servers share one or more LUNs. Simultaneously coordinating the quiescence of all of the virtual machines and their applications concurrently with the snapshot or replication TimeMark[®] is difficult and can cause application downtime while disk I/O is suspended waiting for all the affected virtual machines to complete replication or reapplying the redo logs.

Once one considers the need for system and application integration to quiesce the applications and flush the file system cache along with the wide variety of potential disk configurations, it is clear that organizations need a comprehensive system for extending VMware's built-in functionality by managing the coordination and replication of virtual machines. To accommodate these needs, FalconStor has developed a simplified and consolidated disaster recovery and business continuity replication solution for virtual infrastructures.

The FalconStor Solution

FalconStor's *Application Snapshot Director for VMware* coordinates between IPStor[®]-powered appliances and snapshot agents in the virtual machines to provide an integrated solution for non-disruptive replication of virtual machines 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 needs to be quiesced, and a snapshot 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 ESX Raw Disk Mapping (RDM) in a physical compatibility mode mapped to the ESX 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, the IPStor-powered server 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 IPStor-powered system, 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 TimeView[®] (an image of the underlying volume as it was when the snapshot was taken). These TimeViews 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 existing functionality of VI3, which can only process one virtual machine at a time.

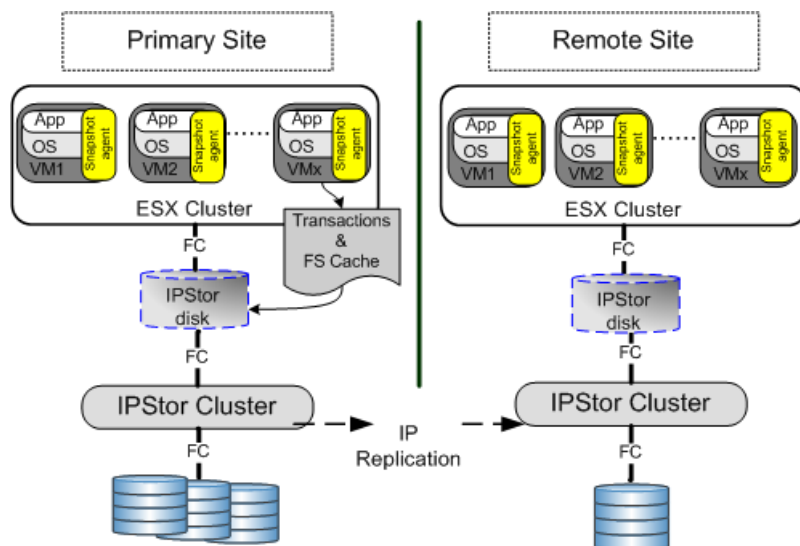


Figure 2: A typical implementation of Application Snapshot Director for VMware.

A typical enterprise solution will involve a clustered primary site and a single node disaster recovery site. For large enterprises, this is achieved using IPStor software on commodity x86 hardware, typically utilizing fiber channel (FC) disk storage (perhaps SATA storage at the disaster recovery site). In this solution there are no limits on storage capacity, and support for up to 256 snapshots per IPStor LUN. The Application Snapshot Director for VMware software is licensed on a per-node basis, and replication can be upgraded to include encryption for WAN data.

For smaller solutions (such as VMware environments with less than 4TB of data), FalconStor offers the *Storage Replicator for VMware Appliance*, which is a turnkey solution that replicates production data in real-time, and provides a hot standby copy that can be activated at any moment for fast, accurate recovery. It communicates with snapshot agents in each virtual machine via the Application Snapshot Director for VMware, which coordinates the quiescing of applications and the taking of snapshots. The Application Snapshot Director can take up to 64 application-coherent snapshots per underlying disk (LUN). Encryption and compression capabilities ensure fast, secure disaster recovery protection.

Conclusion

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. FalconStor's Application Snapshot Director for VMware and Storage Replicator for VMware complement the virtual infrastructure to provide an unprecedented level of versatility, integration, and coordination among virtual storage, operating systems, and applications, for completely consistent, replication of VMware storage and offloaded file-level backups.

AccessFlow, Inc.
1100 N. Market Blvd.
Suite 204
Sacramento, CA 95834
916-339-7047

www.accessflow.com

FalconStor Software
2 Huntington Quadrangle
Suite 2S01
Melville, NY 11747
631.773.5859

www.falconstor.com