

## VIRTUALIZATION AND BUSINESS CONTINUITY: WHAT'S MISSING?

### AN INTRODUCTION TO VIRTUALIZATION AND THE ROLE OF DATA PROTECTION

Virtualization is one of the most compelling technologies of the last decade. The ability to lower information technology infrastructure capital expenditure and operating expenditure due to the ability to dramatically reducing physical infrastructure is profound and has a demonstrable impact on the bottom lines of companies as varied as small businesses to global 1000 enterprises. At the same time, there's been an abrupt shock by many IT professionals as they realize that they are either much more exposed to data loss or that their budgets for infrastructure have only shifted from one type of expenditure to another. And there has been a profound disappointment as the hopes for a simpler, more efficient infrastructure are dashed against the harsh reality of environments that are more, not less, complex and heterogeneous.

What IT professionals are rapidly realizing is that data protection is one of the most difficult challenges in a virtual environment and can't be solved effectively only using virtualization offerings such as VMware. Simple replication, for example, turns out to be not only capital and operationally intensive and expensive but doesn't address data protection issues such as retention - thus for example replication alone simply assures that you've replicated your virus infected data in two places instead of just one.

There are a number of options associated with virtual infrastructures, each with rabid proponents who fanatically tout the advantages while frantically attempting to hide the disadvantages of the approach. In this paper we'll discuss virtualization and business continuity without trying to hide either the advantages or the disadvantages of the various options. We'll refer to VMware since it currently offers the broadest suite of ancillary offerings beyond the core virtualization platform.

### BUSINESS CONTINUITY

VMware notes that there are three core components of business continuity:

- **High Availability:** Maximizing high availability means decreasing both unplanned downtime and planned downtime. It also means that recovery from downtime to a functioning system is critical.
- **Disaster Recovery:** Disaster recovery means the ability to rapidly recover from when a location-wide disaster occurs. A disaster may be natural (e.g., floods) or it may be artificial (e.g., malicious attack.)
- **Data Protection:** The rapid recovery of structured and unstructured data and systems.

In order to illustrate each of these approaches, we'll use VMware's tool suite since it offers the broadest suite of ancillary offerings. To do this, a brief introduction of each of these VMware offerings is discussed in the sections that follow with respect to high availability, disaster recovery, and data protection.

## HIGH AVAILABILITY

**VMware HA (High Availability):** This offering emulates a traditional active/passive cluster using VMware's platform. VMware HA allows an automatic virtual machine restarts when a guest operating system fails. VMware HA also allows a guest operating system to be restarted on another physical server if its primary physical server fails.

**VMware DRS (Dynamic Scheduling of System Resources):** This offering emulates a traditional active/active cluster using VMware's platform. VMware DRS balances computing capacity to deliver performance, scalability, and availability.

**VMware VMotion:** This offering allows the live migration of a virtual machine from one physical server to another without incurring downtime. As such it is emulating a traditional active/active cluster using VMware's platform. VMotion allows customers to proactively and responsively manage both planned and unplanned potential downtime.

Taken together, HA, DRS, and VMotion are a powerful and comprehensive set of tools to manage high availability. These tools are also much less complex than their predecessor high availability technologies of traditional active/passive and active/active clustering and are equally powerful. VMware has also created an entire suite of management tools that allow these technologies to appear as a cohesive set of capabilities to the user.




Figure 3: VMware VMotion

## DISASTER RECOVERY

So VMware offers a wonderful set of high availability options. What about disaster recovery? Once again, there's not much new under the sun. VMware offers a set of geographically separated active/passive clustering options that are packaged within a suite of management tools. VMware offers two fundamental disaster recovery mechanisms: host (server) replication and storage array-based replication. Both rely upon the underlying concept of replication. Replication simply means that data is replicated from one storage device to another. Typical mechanisms for replication include SANs and explicit replication software.




Figure 1: VMware HA




Figure 2: VMware DRS

There are two categories of replication: asynchronous and synchronous replication. Asynchronous replication is more common and does not have the performance impact on the two (or more) systems being replicated but may take from a few minutes to a few hours to switch from the live to the secondary system; synchronous replication is more rigorous and has a more severe performance impact due to communication latency but has the capability of switching from the live to the secondary system in seconds rather than minutes.

If the WAN connection between the replicated systems is of sufficient quality and there is a large amount of bandwidth available, replication works very well. VMware even refers to "dark fiber" in their replication literature - that basically means unused WAN capacity that is seen in some global 1000 companies (e.g., AT&T.) However, for most SMB customers and even larger enterprises that aren't yet in the global 1000, WAN bandwidth is an expensive and precious commodity. Optimizing that WAN bandwidth is incredibly important for these companies. Finally, the techniques of asynchronous and synchronous replication tends not only to be expensive in terms of WAN resources, but it also is very greedy in terms of shared WANs (the most common type - after all, few companies can devote their entire WAN bandwidth during operating hours for replication - they bought the capacity for the applications they use to operate their business.

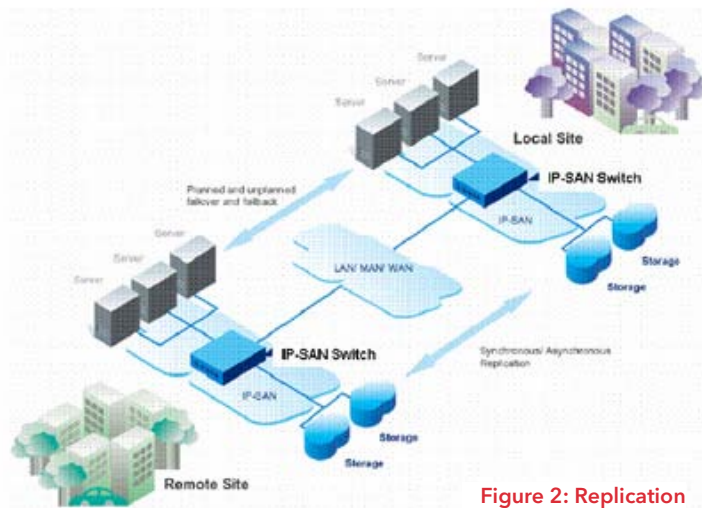


Figure 2: Replication

## DATA PROTECTION

Data protection may be achieved on a VMware system through either **VMware VCB** (VMware Consolidated Backup) or by backing up individually each virtual machine. There are a number of challenges associated with data protection on VMware systems - here are just a few:

- The proliferation of virtual machines means that data protection is not only potentially more time-consuming but is also potentially more complex.
- Virtual machines within an existing information technology infrastructure increases the heterogeneity since existing legacy operating systems and applications must be supported as well as new virtual machines implemented in VMware.
- Backing up each virtual machine independently can be I/O and CPU intensive; simultaneous backup operations for multiple virtual machines can overload the host physical machine.
- The recovery of granular items within a virtual machine can be not only complex but manually intensive.
- The use of VCB without third-party tools requires manual scripting.
- The use of VCB without third-party tools requires integration modules.
- The use of VCB requires entire .VMDK recovery.




Figure 3: VMware Consolidated Backup

# Unitrends™

- The use of VCB requires manual setup.
- Retention must be supported in an optimized fashion or the user will quickly run out of space meeting even relatively minor retention requirements, let alone more stringent regulatory requirements.
- Virus proliferation must be contained across retained backups.

This is the reason that all VMware data protection is designed to work with a third-party backup tool, as per VMware's own recommendations and literature. And of course this is the reason that Unitrends supports virtual infrastructures as well as physical infrastructures.

## VIRTUALIZATION, BUSINESS CONTINUITY, AND UNITRENDS

As noted previously, there are three components to business continuity from the standpoint of virtualization: high availability, disaster recovery, and data protection. VMware does a superb job of handling high availability. VMware has some good disaster recovery offerings on its platform, but these suffer unless a great deal of unused WAN bandwidth is available and is vulnerable to human error, viruses, and the like. In terms of data protection, VMware has designed their platform such that third-party data protection tools are needed.

Unitrends offers a suite of appliances that are designed to both handle disaster recovery in a WAN-optimized fashion taking into account the constraints of bandwidth found at our customer's sites and that are designed to support data protection at both the VCB and the virtual machine level. Our fundamental architecture is depicted in the diagram below this text - the acronym "RVA" means "Retention Virtualization Agent" and enables advanced data protection at the physical and virtual levels on both non-virtualized and virtualized systems.

For more information, please contact us at [800.648.2827](tel:800.648.2827) or [sales@unitrends.com](mailto:sales@unitrends.com).




Figure 4: Unitrends Retention Virtualization Agent Topology