

Growing Pains with Virtual Machine Backups

Author: Alex Mittell
Director of Engineering
phdvirtual.com
Published: April 2009

1.0 Introduction

It's been 10 years since VMware released their first virtualization product, VMware workstation. Since then, Virtualization technology has exploded into the market place with increasingly more organizations migrating their physical servers onto a virtualization platform. Virtualization offers many benefits including less power and cooling requirements, better utilization of computing resources and an easier environment to manage. However, even as organizations have virtualized with problems solved and costs reduced, so too have new challenges emerged, specifically the issue of how to back up consolidated server data. Backup combines the most I/O intensive application in the datacenter with the system least able to handle intensive I/O loads, namely virtual servers. This whitepaper looks at the growing pains that organizations have experienced searching for an effective and efficient backup solution for virtualized environments.

2.0 History of Virtual Machine Backups

In The Beginning...Traditional Backups

In the early days of virtualization there weren't any specialized backup solutions designed to take advantage of the new virtual architecture. The only option was to back virtual machines up as if they were physical servers, a method built on traditional hardware backup practices. This typically involved agent-based backups using products such as Symantec NetBackup, Backup Exec, Legato Networker and HP Data Protector, among others, and although they provided the traditional level of functionality, they did not take advantage of the new architecture or allow for the constraints brought about by virtualization.

Their limitations were:

- Inability to back up powered off virtual machines
- Inability to back up VM templates
- Placing a heavy load on CPU and Memory in the virtual machine
- Requiring backup agents to be deployed and maintained on the individual virtual machines
- Inability to take image level backups without purchasing additional "bare metal restore" licenses per virtual machine

And Then...Custom Scripts Began to Appear...

After a few years, custom scripts started appearing, written by people who sought a way to back up virtual machines that took advantage of the fact that they were essentially just a handful of files. It seemed senseless treating virtual machines as if they were physical when clearly they were different and had some unique advantages. Having to purchase a bare metal recovery license for a server restore also seemed unnecessary when all you had to do was backup a few (albeit large) files. A few innovative people started to realize that there was another way to approach this, so the development of scripts that could suspend a VM, copy the virtual machine files, turn the virtual machine back on and then send the files off to a backup target. This was a significant step forward; having to take the virtual machine offline, however, was a major drawback.



Figure 1: Early VM Backups

Table 1: First VM Backups

Pro's	Con's
Relatively simple to script	Takes VM's offline
Crash Consistent Backup	Backs up full VMDK files each time
Worked!	

The problem of performing hot backups was soon resolved as VM snapshot technology advanced. By utilizing snapshots and Redo logs the backup could take place online and the redo logs would take care of any changes that occurred during the initial backup process. This resulted in the first online Crash Consistent backups.

Table 2: Early Hot Backups

Pro's	Con's
Server stays online	Backs up full VMDK file each time
Crash Consistent Backup	Limited options
	Difficult to manage and maintain scripts in large environment

Early hot backup processes were fairly effective and performed the fundamental job of backing up and restoring virtual machines, but had some limitations in terms of management and flexibility. The scripts only did full backups each night, but were unable to safely abort backups without leaving the redo logs, and they didn't have the flexibility to include or exclude individual machines or retry failed backups. Once again customers were demanding more...

In 2004 a product called esxRanger was released by a start-up company named VizionCore which focused on backing up VMware virtual machines. This product offered definite improvements over the existing scripts including scheduled backups of only specific VMs and pruning backups after a user-defined number of days. Although esxRanger was an improvement, there were a number of drawbacks: reliability issues, random backup failures, inability to safely abort backups and ability to perform only full backups. The search for an effective hot back up solution was still on.

Genesis of PHD Virtual's esXpress ...

Whilst VizionCore was trying to develop a stable backup solution a couple of engineers at PHD Virtual had begun to develop their own solution, as a result of the work they were performing for clients. As Unix/Linux consultants back in the early 2000's they were early adopters of VMware and quickly realized that virtualization was destined to play a big part in the future of IT. As early adopters they had to find ways to perform certain tasks that were critical to a production implementation such as backing it up. Initially, they wrote custom VMware backup scripts to backup virtual machines, also working out ways to perform P2V migrations long before any commercial tools were available. Even as P2Vs were keeping them busy, their clients were demanding better backup solutions and so, seeking a more effective method, the PHD Virtual engineers redoubled their efforts to write backup scripts. These evolved into hot backups via the



Figure 2: Early Hot Backups

use of the redo logs.

Initially, the new method of hot backups seemed to work well, but some of the advantages of backing up virtual machines also started to create problems. The fact that a virtual hard disk is just a single .VMDK file made it easy to backup and restore, however it also meant that files were very large and if there were any changes, then the complete file would need backing up again: 100 virtual machines with 50GB hard disks would require 5TB of data to be backed up each night even if only a few bytes had changed!

Knowing there must be a better way PHD Virtual decided to try and engineer delta (or incremental) backups for virtual machines. Delta backups weren't a new idea; agent-based backup solutions had been doing them for years. However these solutions still only backed up to the file level which for large virtual machine files wasn't sufficient. A block level solution was needed. The PHD Virtual engineers focused on creating the industry's first product that would perform delta backups down to the block level for virtual machines.

In 2006, after five years of work the PHD Virtual enterprise backup product for VMware was released and named esXpress. esXpress included delta backups as well as block and file level data backup.

VMware Consolidated Backup (VCB)

VMware's answer to virtual backups was released in VMware Infrastructure 3 as a 'free' backup solution. Although late to the party, VCB offers users and backup vendors the ability to perform some of the actions required to backup their virtual machines. The VCB name is a little misleading in itself as it is not a backup solution but rather a series of scripts which trigger a snapshot of a virtual machine, thereby allowing the VMFS file system to be presented to a Windows server so third party backup products can back up the files. In many ways, these packaged VCB scripts perform similar tasks to the custom scripts of prior years. VCB has allowed the traditional, so-called enterprise backup products to finally realize a product that takes advantage of a virtual machine's encapsulated architecture. These can now simply "bolt on" to the VCB scripts and claim that they are designed for VMware backups. This would be fine if it fully solved the problem. But does it?

While there's no argument that VCB-based backup solutions allow backing up of virtual machines which in many cases meet the customer's requirements, it remains inefficient. Following the example above, consider a situation with 100 virtual machines each with 50 GB hard disks (.VMDK files). Each time one of these files changes, (which is almost anytime the servers are online) the .VMDK file changes and hence the whole file needs to be processed by VCB. 100 VM's x 50 GB disks = 5TB of data each backup cycle!

There are legitimate reasons why customers use a VCB-based solution. The processing occurs on a dedicated VCB Proxy server and customers who are used to traditional backup products may be reluctant to venture outside of their comfort zone to look for a more purpose built solution. In a situation where existing third party backup licenses exist in the organization, and the time and bandwidth that the virtual machine takes to backup is not important, then a VCB based solution is sufficient.



Figure 3: Delta Backups

Even so, VCB has other drawbacks, the primary being cost. VCB is free right? Wrong if the user considers what is actually needed to back up a virtual machine with VCB. First there is the VCB proxy server itself. One of the advantages that VCB advertises is that it's a LAN free backup and this can be true. However, to execute LAN free backups, the VCB proxy server must connect directly to the SAN requiring the customer to provision VCB Proxy server hardware, a Microsoft Windows license, and an expensive Host Bus Adapter. Far from "free". In addition, while the third party backup software backs up the data, costly storage must also be provisioned for the VCB Proxy on which the virtual machine data will reside. By comparison, and as explained below, the cost of a single esXpress Host license, which includes as many CPU's as the customer needs, may actually end up being much cheaper. It can become cheaper still if users consider the cost in saving management time and oversight.

Hot Full Backups Using esXpress

Returning to the world of PHD Virtual, once esXpress was released, word quickly spread amongst the VMware community and it became renowned because of its robustness and speed. esXpress approaches backups from a different angle in that it doesn't require any additional hardware or Windows software licenses to run - it uses small dedicated Linux Virtual Appliances called VBAs (Virtual Backup Appliances). Since the VBAs are virtual appliances they run outside the service console and their resource usage can be constrained by running them within resource pools. esXpress uses the virtual environment to back up the virtual environment. The product provides valuable features not seen elsewhere: it is finally possible to reliably Pause, Stop and Abort backups and have the snapshots cleaned up in the process. esXpress' pioneering delta technology saves disk space, network requirements and time. Not only does it provide delta backups, it also compresses the data at the source providing further network bandwidth savings. The backups can be sent to multiple backup failover targets, and the alerting function provides administrators the visibility they require. Replication features also provide a multi-tier approach to storing the data and an extra level of data redundancy. From the outset, the philosophy at PHD Virtual was always "It's Your Data", and consequently, the backup files can be restored without an esXpress installation present since the backups are stored at the target as compressed files using industry standard compression techniques.

Within only two years of its release esXpress won the prestigious Gold Award for Data Protection at VMworld in 2008 and was named Searchservirtualization's Data Protection Product of the Year in 2008. esXpress has continually evolved and, at the time of writing the next version of esXpress (3.5) is in Beta, due for release in early Q2, 2009. This release offers exciting new features such as built-in de-duplication at the source, a central management console integrated into Virtual Center that allows the installation to be centrally pushed out and configured, as well as advanced replication and disaster recovery capabilities. The revised file level backup option which indexes all files as part of the backup provides a web-based graphical console. With only a few clicks IT staff can easily restore any file or folder within a backup.

With the growing number of virtual machines, users increasingly appreciate that a reliable and efficient backup solution is critical to ensuring disaster recovery capabilities of their virtual infrastructures. It explains why more and more people are turning away from traditional backup products towards solutions like esXpress that are designed from the ground up to cost-effectively protect data in a virtualized environment.

3.0 Summary

When designing a virtual infrastructure organizations should keep backup firmly in mind. Administrators shouldn't think about backups until it's all finished! A well-designed virtual environment can fall apart if it cannot handle the task of backing up massive amounts of data in a limited time.

From the initial growing pains, virtual backups have come a long way - the development remains ongoing. The pending release of vSphere, the next version of VMware ESX and the growing uptake of other hypervisors such as Microsoft's Hyper-V and Citrix XenServer mean that the development team at PHD Virtual is constantly making sure they are ahead of the curve. Backing up other non VMware Hypervisors as well as backing up directly to tape may be additional solutions to look for in the future from PHD Virtual.

4.0 Useful References

1. PHD Virtual <http://www.phdvirtual.com/>
2. esXpress 3.1 Trial <http://www.phdvirtual.com/download>

About PHD Virtual

As the pioneer of virtual backup appliances (VBAs), PHD Virtual Technologies has been transforming data protection in virtual IT environments since 2006. Its award-winning data protection solution for virtual infrastructures, esXpress, is used today by more than 1,600 enterprises worldwide to achieve unlimited dynamic growth, high availability, no single point of failure and scalable performance. PHD Virtual is committed to helping our customers and provides free, easy-to-use virtualization utilities to assist with the administration and management of virtualized environments.

North America Headquarters
111 Howard Blvd. Suite 104B
Mount Arlington, NJ 07856
Phone +1 (973) 288-7000
www.phdvirtual.com

