

# White Paper

---

## **IBM Tivoli Storage Manager for Virtual Environments**

**Advanced Data Protection for VMware ESX Environments**

*By Lauren Whitehouse*

**June, 2011**

---

This ESG White Paper was commissioned by IBM  
and is distributed under license from ESG.



## Contents

Introduction .....	3
Virtualization and its Impact on Data Protection .....	3
Optimized VMware vSphere Backup & Recovery with TSM for VE .....	6
New Functionality with TSM for VE .....	6
Tivoli Storage Manager for Virtual Environments .....	7
The Bigger Truth .....	8

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.


## Introduction

Server virtualization is taking hold in companies of all sizes, and [VMware](#) is one of the more popular hypervisors adopted by IT organizations. While VMware server virtualization continues to gain momentum, IT organizations still have some hurdles to overcome if they are to deploy virtualization more widely across the enterprise. Backup and recovery of virtual server environments ranks highly as a top initiative and area of investment—a major focus for a growing population of corporate IT organizations expanding the use of virtualization to incorporate more tier-1 production applications. This paper introduces Tivoli Storage Manager for Virtual Environments (TSM for VE), [IBM's](#) solution for optimizing backup and recovery in VMware installations.

## Virtualization and its Impact on Data Protection

Virtualization is a major initiative; ESG research indicates that organizations are increasing their use of the technology, making it the number one IT initiative for 2011.<sup>1</sup> What's the attraction? There are many. Consolidating multiple virtual machines on a single physical system in lieu of a one server/one application model delivers CAPEX and OPEX savings. It also makes IT more agile. These and other benefits explain why ESG research found that within the next two years, 58% of organizations say they will virtualize 40% or more of their x86 servers (see Figure 1).<sup>2</sup>

Figure 1. Current and Planned Adoption of Server Virtualization



Source: Enterprise Strategy Group, 2011.

VMware is a market leader in virtualization—in use by more than 170,000 organizations. ESG research confirms its popularity as a preferred hypervisor vendor: over half of research respondents have a VMware hypervisor in use (see Figure 2).<sup>3</sup>

Deploying virtualization in production environments isn't taken lightly. For an application to be production-ready, a system and/or data protection strategy needs to be in place—or an organization risks suffering the consequences of

<sup>1</sup> Source: ESG Research Report, [2011 IT Spending Intentions Survey](#), January 2011.


<sup>2</sup> Source: ESG Research Report, [The Evolution of Server Virtualization](#), November 2010.

<sup>3</sup> Ibid.

data loss or downtime. This task has not been an easy one. While simply doing what was always done can work, it's not always the most optimized—and isn't great optimization the idea behind virtualization?

Figure 2. Server Virtualization Vendors

Which of the following would you consider to be your organization's primary server virtualization vendor? (Percent of respondents, N=441)



Source: Enterprise Strategy Group, 2011.

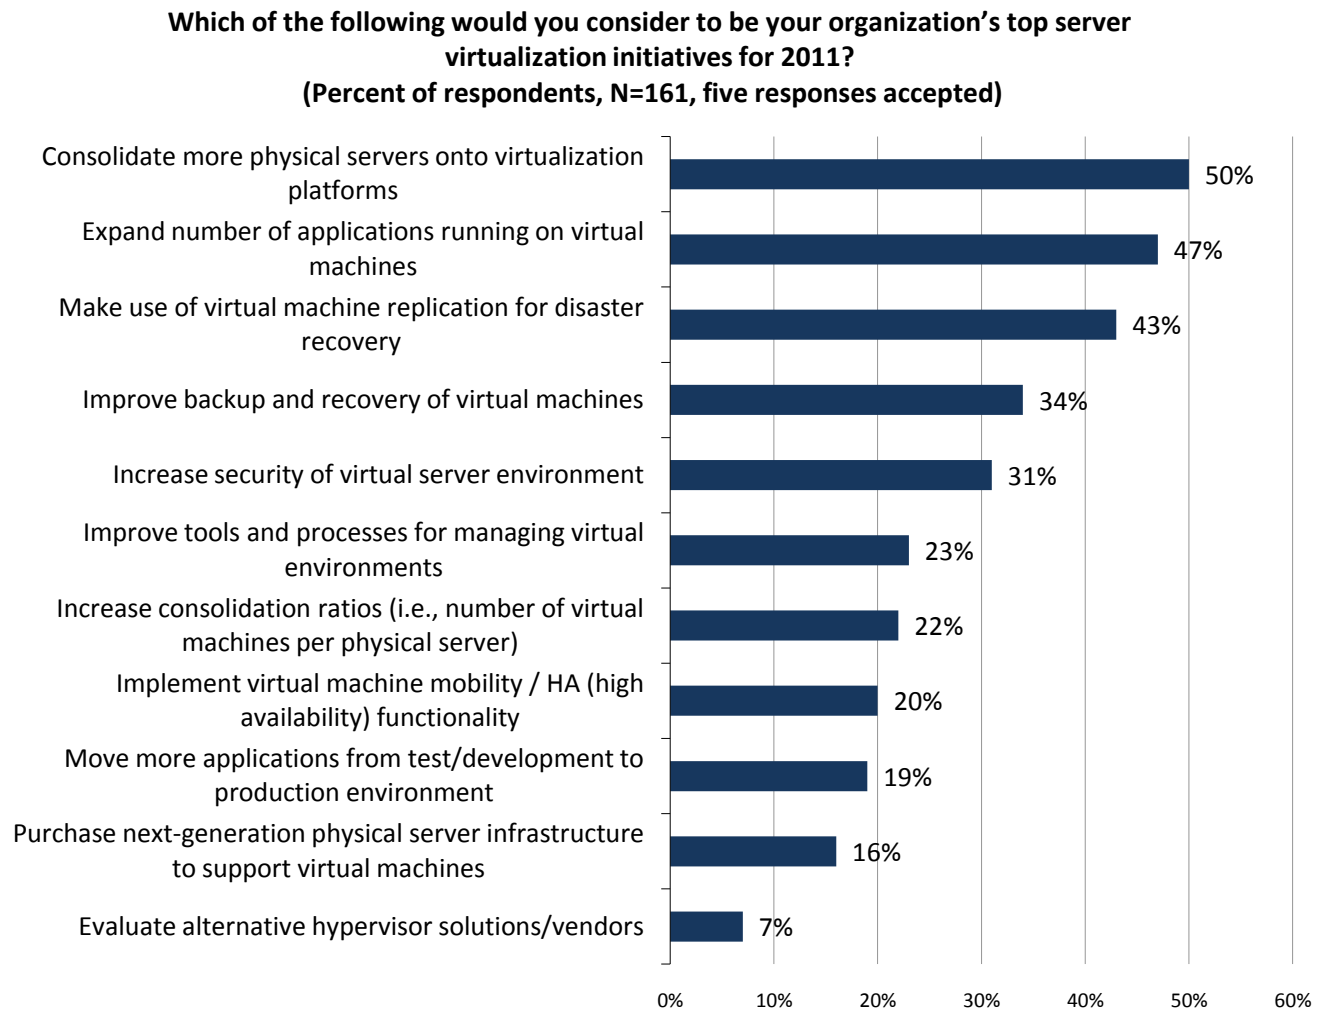
Performing backup of virtual machines with strategies and technologies built for physical environments (i.e., installing a backup agent in the guest operating system) is not optimized. An “in-guest” approach places a burden on the VMware host server’s CPU, memory, disk, and network components. While that may be okay in a 1:1 virtual-to-physical consolidation ratio, it’s likely not to be in a 10:1 configuration. Resource contention is possible since typical backup jobs demand significant processing power—use of the required resources needed to execute backup jobs may compromise the performance of other virtual machines sharing the same resources. ESG research respondents concur: the need to improve backup/recovery for virtual machines is a top initiative in virtual environments (see Figure 3).<sup>4</sup>

Another issue plaguing backup processes is the amount of data redundancy within and across virtual machines. Virtual environments have natural data redundancy because each virtual machine maintains its own copy of operating systems, applications, and data. The byproduct is often an increase in the volume of data to store, manage, back up, and recover. About one-quarter of organizations note that the impact of virtualization on the volume of data storage is a significant challenge<sup>5</sup> and, if not optimized, it can exceed storage quotas. Similarly, unless data for virtualized environments is optimized, the volume of data can also exceed available network bandwidth. Finally, an increase in the volume of data could result in more time required for backup processes—taxing already-stressed backup windows.

<sup>4</sup> Source: ESG Research Report, [2011 IT Spending Intentions Survey](#), January 2011.

<sup>5</sup> Source: ESG Research Report, [The Evolution of Server Virtualization](#), November 2010.

Figure 3. Server Virtualization Priorities



Source: Enterprise Strategy Group, 2011.

The volume of data also impacts recovery. The more data transferred from a backup repository to a recovery site/system, the longer the transfer takes. Users are looking for optimization here, too. For example, recovery of a whole virtual machine (i.e., the .vmdk file containing the virtual machine image) is often desired to accelerate recovery. While whole-virtual machine approaches are optimized, IT organizations still require the ability to recover data at a file or item level to streamline recovery. Flexibility in data recovery options is key to meeting the new recovery demands of virtualized systems.

Another factor to consider is that the on-demand provisioning enabled by virtualization could cause virtual server sprawl and create an unwieldy management situation as well as vulnerability. Since virtual machines can be provisioned more rapidly than physical ones, the risk of unprotected systems (a virtual machine not being incorporated into the backup strategy) can be introduced.

## Optimized VMware vSphere Backup and Recovery with TSM for VE

IBM Tivoli Storage Manager (TSM) software is an integrated backup, recovery, and archiving solution supporting business continuity and long-term data retention efforts. It is a comprehensive data protection solution for both physical and virtual server environments via a single point of control. While TSM provides an effective solution for backup and recovery of VMware environments, IBM introduced TSM for VE which offers features that enhance and optimize backup and recovery processes for VMware vSphere hypervisors.

### New Functionality with TSM for VE

TSM for VE leverages several features of vSphere, VMware's flagship server and infrastructure virtualization platform. In its latest release, VMware introduced a number of features that data protection vendors, such as IBM, can take advantage of to optimize backup and recovery. Many of the hypervisor-specific features are available to data protection vendors through special application programming interface (API) commands.

#### VMware vStorage APIs for Data Protection (VADP)

VMware made a big improvement in its backup and recovery APIs in its latest iteration of VMware infrastructure. The company introduced its vStorage APIs for Data Protection to enable direct access to VMFS file system drivers, allowing TSM for VE to access data on VMFS volumes—optimizing data backup. The APIs enable snapshot-style backup, providing an “off-host” method of protecting data. By offloading backup workloads from a VMware ESX- or ESXi-based server to a centralized vStorage backup server, backup processing on the virtual machine is minimized considerably and the need to install a backup agent in every virtual machine is eliminated. This approach also greatly accelerates backup and recovery.

#### Changed Block Tracking

With Changed Block Tracking (CBT), incremental backups occur at a block level. CBT keeps track of changes in virtual disk blocks for virtual machines at the virtualization layer. This information is available to TSM for VE, so data transmitted during backup is limited to only the blocks that changed since the last backup.

#### Direct backup from shared storage

A backup system in the backup data path can create network bottlenecks and inefficiency. Now, a direct-to-target architecture is employed, allowing backup data to go from the ESX or ESXi host directly to the storage target. By leveraging the vStorage APIs, virtual disk data (.vmdk) is read directly from the SAN storage device.

#### Hot-add virtual disk support

Hot-add virtual disk support bypasses the network layer and streams data directly to the target backup storage, reducing LAN traffic and improving performance. By leveraging APIs, the TSM for VE components on the vStorage backup server can access the virtual machine image on shared storage to read the data and back it up.

#### Virtual or physical backup vStorage backup server

In vSphere, the backup server can be a virtual machine or a physical one. Using a virtual machine as the backup server eliminates the need to invest in physical infrastructure for backup systems, reducing backup infrastructure costs.

#### Granular recovery from whole virtual machine backup

Multiple restore options are available from the virtual machine-level backup via the proxy server. TSM for VE can rapidly restore data at the file-, disk volume-, or virtual machine-level from a single-pass whole-virtual machine image backup. The speed and flexibility introduced in this feature allows organizations to meet aggressive recovery SLAs.

## Tivoli Storage Manager for Virtual Environments

Tivoli Storage Manager for Virtual Environments is a virtual-only backup and recovery solution that introduces tremendous efficiency in data protection for VMware vSphere installations by capitalizing on the vSphere APIs. However, TSM for VE has several additional features that introduce another layer of backup optimization.

Transferring and storing less data introduces cost and time savings. In an era where data growth rates translate into data doubling every three to five years,<sup>6</sup> introducing efficiency can really pay off. For backup operations, data growth issues are compounded by the fact that multiple copies of data are retained for extended periods of time for recovery or business intelligence needs.

TSM has native, efficiency-minded features that when combined with the network- and storage-optimization features in vSphere—such as CBT, hot-add virtual disk support, and direct-to-shared storage backup—deliver an even higher degree of optimization. Deduplication is one such feature. It identifies and eliminates redundant data in backup data streams at the block level. Only unique data is transferred and stored, while a pointer to unique data replaces duplicate blocks. Deduplication works in concert with CBT to further optimize the capacity of data transmitted and stored.

Speeding recovery is equally, if not more important, for many companies. TSM for VE offers near-instant restore of Windows and Linux disk volumes, allowing data to be made immediately available while it is recovered behind the scenes.

The issue with virtual server sprawl discussed earlier is that new virtual machines can be provisioned and go unprotected, effectively orphaned in backup processes. TSM for VE, however, can automatically discover new virtual machines and automatically assign backup policies to them. This lowers the risk of data loss and downtime.

TSM for VE is a fully-integrated member of the Tivoli Storage Manager family. From TSM, management of physical and virtual system backup is available from a single console. It also enables more storage tiers for backup since TSM supports physical tape media. Therefore, TSM for VE-invoked virtual system backup can be incorporated into TSM's policy-based tiered storage targets, including tape.

TSM's unified recovery management capability provides a single point of management for all restore operations. By tracking the various recovery points, managing them efficiently, and providing access to them wherever they reside, administrators can improve productivity, and, importantly, recover data faster and with fewer errors.

---

<sup>6</sup> ESG research found that 58% of [2010 Data Protection Trends](#) survey respondents cite annual data growth in the 20% to 30% range.

## The Bigger Truth

Virtualization is not a fad; its benefits are not over-hyped. Implementing virtualization delivers value in minimizing time, cost, and risk due to the introduction of “dynamic” provisioning, CAPEX and OPEX savings via consolidation, and business continuity benefits. The downside is that virtualization often necessitates an infrastructure refresh to reap these optimization rewards.

When it comes to backup and recovery, the infrastructure refresh need only come in the form of adding TSM for VE to the environment. As a complement to IBM’s flagship backup, recovery, and archiving solution, Tivoli Storage Manager, TSM for VE enhances support of VMware virtual server environments. By taking advantage of several vSphere features and upping the ante with its own optimization, IBM is meeting one of the biggest challenges and impediments to expanding the virtualization footprint: *optimized* backup and recovery.





Enterprise Strategy Group | **Getting to the bigger truth.**