

Top 5 Requirements for Backup of Virtual and Physical Servers

by Greg Shields, Microsoft MVP and VMware vExpert





Greg Shields

Microsoft MVP and VMware vExpert

Greg Shields is a Senior Partner with Concentrated Technology. With fifteen years of IT experience, Greg is one of the world's leading experts on virtualization, cloud, and systems management technologies. He is a Contributing Editor and columnist for Microsoft TechNet Magazine and Redmond Magazine, has written over fourteen books, and contributes regularly to online publications like TechTarget and MCPMag.com. He is also a highly sought-after and top-ranked speaker for both live and recorded events, and is seen regularly at conferences like TechMentor, the Microsoft Management Summit, Microsoft Tech Ed, VMworld, Connections, among others. Greg is a multiple-year recipient of Microsoft's Most Valuable Professional (MVP) award and VMware's vExpert award.

Virtual and Physical – Protecting Both

Reports by leading industry analysts like Gartner, IDC and Concentrated Technology suggest virtual servers in 2011 will eclipse physical servers in total server deployments. The majority of today's business computing environments already have both virtual and physical servers at the same time.

This hybrid reality drives the need for protecting both virtual and physical environment needs. In so doing, the following list of "Must-Have" requirements should serve as a basic roadmap for IT managers as virtualization use continues to outpace physical server utilization.

- **Must Have Requirement #1 (Complete Protection):** Protecting both physical servers and virtual servers
- **Must Have Requirement #2 (RTO):** Reducing recovery time from hours or days to an RTO of minutes. Today's datacenters are no longer willing to accept downtime of any measurable quantity. Any selected data recovery solution must meet this need.
- **Must Have Requirement #3 (RPO):** Completely eliminating backup windows. Doing so reduces the RPO to just minutes. Losing time is an impact, but losing data is a mission critical failure. Solutions that prevent data loss also prevent business impacts.
- **Must Have Requirement #4:** 100% Reliability / assurance of recovery. The notion that backups must be manually tested has become quaint to the point of ridiculousness. The right solution automates this activity for you, eliminating manual verifications while ensuring data is correctly captured.
- **Must Have Requirement #5:** Universal recovery across virtual or physical environments (anywhere to anywhere). It includes restoring to dissimilar hardware. This also includes the capability for granular recovery that ranges from files and folders to email messages and database rows, to application objects, and ultimately full bare metal server recovery.

Achieving the Must Haves by Solving the Perspective Problem

Accomplishing these can't be achieved with just any legacy backup products. The approaches used by legacy backup software suffer because they were developed many years ago for physical servers and were merely retrofitted to accommodate virtual environments.

These legacy backup products also suffer from performance problems, they don't scale well, and can even actually corrupt data as it is backed up. They worked adequately in the early days in the physical-only environments they were designed for, but frankly altogether fail when applied to virtualization.

Many Perspectives for where a Backup Agent Might Sit

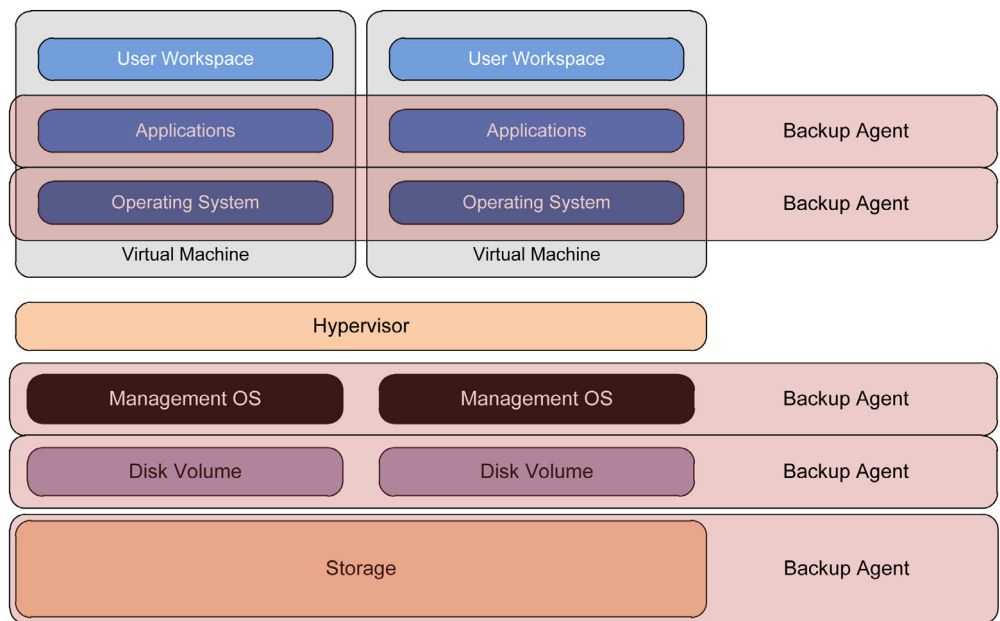


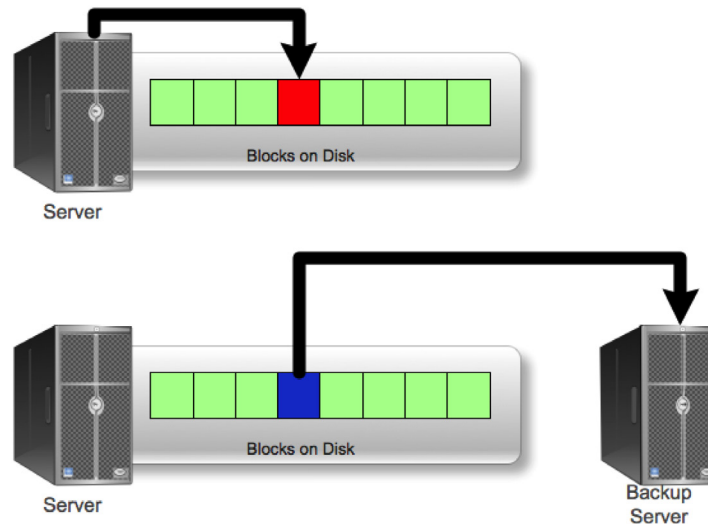
Figure 1 shows a graphical representation of exactly those possible perspectives. A backup agent with application awareness is necessary when applications require quiescence. One is needed in the operating system and yet another for files and folders and to interface with application-aware agents.

Virtualization adds further complexity, as well as new perspectives from where backups can be sourced. An agent might sit in the hypervisor's management partition, within its external storage, or even the VM itself. Each of these positions comes with its own strengths as well as weaknesses.

There is, however, one location across both physical and virtual infrastructures that presents a single, unified place for a smartly-designed agent to capture everything: The disk volume layer. Far better than dealing with individual files and folders in a file system, the disk volume layer sees everything as blocks on disk. When data changes, so does the disk block.

Capturing those changes and replicating them elsewhere means backing up those changes as they occur. That backup happens in near real time, and in a fashion that can be replayed back elsewhere whenever necessary.

Replaying changes in a block-by-block fashion.



Such a solution is elegant in its simplicity. We know that at the disk volume layer data is treated as blocks. Thus, attaching a backup agent to this process means watching each block as it changes. Once that block changes, as Figure 2 shows, the change is replicated to the backup server. This eliminates the entire notion of a backup window, because data is constantly being backed up. **Must Have #3, solved: Completely eliminating backup windows and reducing the RPO to just minutes.**

Once at that server, all kinds of useful things can be done with the data. Stored on disk, views of that data can be created almost instantaneously by simply replaying the data into a workable user interface. **This solves Must Have #2: Reducing recovery time from hours or days to an RTO of minutes.**

Testing the backups can be automated by running a series of automated tests. Those tests can validate file system integrity, as well as database functionality across a range of supported applications. **That fixes Must Have #4: 100% Reliability / assurance of recovery.**

Data within VMs is also far easier to restore. Such data can be simply gathered using similar views, without requiring a painful double-restore process. In fact, with the right intelligence built into such a system, restoring any kind of supported data structures can be accomplished via the exact same process. This eliminates the variety of competing steps that are necessary in other backup architectures, and wholly resolves **Must Have #5: Universal recovery across virtual or physical environments (anywhere to anywhere).**

Above all, a disk volume focus enables virtual and physical computers to be treated equally from both the backup and recovery perspective. No longer must backups for virtual machines involve different steps and technologies than those for physical ones. From the perspective of the backup, they're all just blocks on disk. That's the solution for **Must Have #1: Completely protecting both physical servers and virtual servers.**

The Best Backup Protects Both Virtual & Physical

Ultimately, the goal in eliminating these problems is in finding that one backup that protects both virtual and physical servers and is built for virtual and physical from the ground up (not just built for physical and retrofitted later for virtual). Focusing attentions on disk volumes, as you've seen here, represents one approach that might just fit every Must Have. With a remarkably simple and yet powerful architecture and a reasonable fulfillment of today's needed requirements, this approach is the only way to go with in today's mixed virtual and physical IT environments.