VEEAM
IT JUST WORKS!™

Top Ten Best Practices for vSphere Application Backups

Greg Shields

Microsoft MVP and VMware vExpert

Modern Data Protection
Built for Virtualization | **#1 VM Backup**

Virtualization has become the *standard* for Windows servers. The proof is in its widespread acceptance. Some analysts estimate that nearly 60% of all IT workloads will be virtualized by the end of 2013. A large portion of those will be running atop VMware's vSphere virtual platform.

But while virtualization eases the activities in server management, it can sometimes complicate those for enterprise applications. One activity requiring extra attention in a virtual environment is data protection. In a virtual world where different backup solutions can gather data from different locations, it can be challenging to figure out the very best ways to protect your enterprise applications.

There is a bit of good news. With so many servers now virtualized, vSphere administrators are beginning to agree on the best practices for backing up critical business applications. What follows are ten of the very best practices you might consider implementing in your vSphere virtual environment.

#1 – Ditch the inside-the-guest backup solutions

Virtualizing a server atop VMware vSphere automatically creates new locations where backups can be sourced. These new locations go beyond the traditional approach where agents are installed into a server's OS with the goal of backing up files and folders directly.

With vSphere, backup data needn't necessarily be sourced from inside the VM. Your backup data can be the VM's disk file itself. You can capture that disk file from a variety of locations, such as the hypervisor, another co-located VM or even directly from your storage device.

Each location offers advantages to the backup process that don't exist in traditional inside-the-guest solutions. Using these newer approaches, entire VMs can be more easily restored and can be seamlessly migrated and/or replicated to new virtual hosts. Tests against backup data can be better automated and disaster recovery better assured when backups and offsite replication are integrated into a single solution

> **The best practice:** Start by replacing your entire data protection approach. The files-and-folders approach is dying and can't keep pace with today's datacenter demands.

#2 – Leverage snapshots, but not hypervisor snapshots

The term *snapshot* has gotten somewhat overused in today's IT vernacular. Your hypervisor comes equipped with snapshots. Backing up applications inside a VM leans on snapshots. And backing up that VM's disk file often leverages actions that some call snapshots. With so many very different activities sharing the same name, it isn't easy to keep the terminology and the activities straight.

vSphere virtual disks operate almost like databases in that they contain data (a VM's files and folders) and require special efforts to be backed up online. Your backup solution should leverage a snapshotting—or quiescence—process to ensure data integrity during backups. That process must occur for both the VM's disk file, *as well as the applications installed inside the VM*. These are much different activities than clicking "Take Snapshot" inside the vSphere Client

> **The best practice:** The *Take Snapshot* action in the vSphere Client is not a backup solution. Using it can actually create significant downstream problems if snapshots are kept for too long. On the other hand, snapshots in a backup solution are a fundamentally important capability. Fear the click, but don't fear the term.

#3 – Separate backup data from production data

Virtualization and its new backup approaches are quickly bringing about the end of the tape backup era. Disk is replacing tape as the medium for storage. Disk is easier to work with, offers greater performance and can enjoy special benefits like data deduplication and compression.

Yet there's a catch: Not all disks are *the right disks* for backups.

Be careful about consolidating your backup data with production data on the same storage device. Doing so might seem like a good idea, because…well… disks are disks. But by combining backup data with production data you're setting yourself up for failure should that storage device fail catastrophically.

> **The best practice:** Don't let a storage device failure take down your VMs *and all their backups*. Use separate storage for backups to protect against a single point of failure (and a Resume Producing Event).

#4 – Beware of physical RDMs and guest iSCSI connections

In vSphere's early days, virtual disks were limited to relatively small sizes. A large and growing database, for example, could easily hit the maximum size of a VMDK (Virtual Machine Disk) file. Large VMDKs could also be sources of performance problems in certain use cases. Eliminating these limitations required eliminating the VMDK file and storing data in a non-encapsulated form using Raw Device Mappings (RDM).

Today RDMs come in two flavors: physical and virtual. Virtual mode offers full virtualization of the mapped device and, with the right backup solution, most of the benefits of a VMFS (Virtual Machine File System) virtual disk. Physical mode offers no such virtualization.

As a result, physical mode RDMs are sometimes incompatible with vSphere backup solutions. These solutions focus their attention outside the virtual guest and, as such, can't always recognize connected RDMs inside. The same holds true for iSCSI connections initiated from inside the guest. As with physical RDMs, these guest iSCSI connections often can't be seen—and, thus, backed up—by a vSphere backup solution.

> **The best practice:** Eliminate the physical RDMs and guest iSCSI connections from your vSphere VMs. They don't enjoy the benefits of virtualization and sometimes won't enjoy the benefits of data protection either.

#5 – Choose the correct level of protection

These days, data protection has evolved far past its roots in mere server backups. Protecting IT servers means protecting IT services, many of which are classified as highly critical to business operations. Protecting critical services requires extra effort beyond simple server backups.

That extra effort can come from replication, where backed up server disks are replicated to an alternate location. The replicated server disks wait in that location for the primary server to experience a problem and for an administrator to kick off a recovery, which, if comes from clustering, greatly automates the return to service for a dead or dying server. Recovery can also involve fault tolerance, which links the disks of two separated VMs in lockstep so that one can immediately take over for the other.

What's most important to recognize is that each added level of protection adds to your management cost. Choosing not just the best, but the correct, level of protection is a key decision in today's data protection activities.

> **The best practice:** Ensure that your backup solution supports all the protection capabilities your mission-critical servers might need. Implement those that make sense for each workload on a case-by-case basis.

#6 – Replication: Server-centric or storage-centric?

The process of data replication itself can happen via a variety of mechanisms, each with its own benefits and concerns. In a general sense, replication can exist as a server-centric activity or as an activity that happens at the storage layer. Server-centric replication tends to integrate well with running services and applications, but sometimes at the cost of performance. Storage-centric replication tends to perform well, but requires extra effort in orchestrating its activities with those running inside each VM.

> **The best practice:** There is no globally accepted best practice regarding whether to select a server-centric or storage-centric replication process. Many backup solutions now offer replication as a feature, as do many storage devices. Evaluate which works best for your needs.

#7 – Automate application recovery testing

IT has been told to "test the backups" for decades. For decades most of us have been forced to ignore the requirement. In the days of tapes, testing the backups was an activity that was challenging to the point of absurdity. Combine that difficulty with the recognition that backup testing should be an everyday event, and one can quickly understand why this critical activity was rarely done.

Among all their other benefits, vSphere virtual disks offer incredible portability. A disk that runs on one ESXi host can run on just about any ESXi host anywhere. Combine this fact with new approaches in disk-based backup, and you'll quickly see how testing VMDK backups suddenly became very, very simple.

The right backup solution will take much of the backup testing out of your hands, automating the provisioning of backed up virtual disks to special ESXi hosts. In a protected and isolated environment, each backup can be tested for successful boot, data integrity and even application functionality. This process can largely be automated, freeing you to continue "not testing" the backups—as your backup solution handles testing for you.

The best practice: *Don't test your backups*. Let an automated solution do it for you. Seek solutions that can test against the variety of possible failures such as ability to boot, data integrity and even application functionality.

#8 – Verify VSS Writer registration

Microsoft's Volume Shadow Copy Service (VSS) offers a three-part solution for facilitating application backups. One part, the VSS Requestor, is managed by your backup solution. Another part, the VSS Provider, orchestrates activities with storage. The third part, the VSS Writer, registers the installed applications' before-, during-, and after-backup activities to ensure data integrity. For more information, please see Microsoft VSS: What Every VMware Admin Needs to Know.

This separation of "what the backup solution manages" from "what the application manages" is an important evolution in backup compatibility. As long as your application registers itself correctly with the VSS Writer, you can rest assured that your applications will recover correctly.

However, a little due diligence is required to ensure that each application has registered with VSS correctly. In Microsoft Windows, the command *vssadmin list writers* can be used to display information that will assist with this verification. The command produces a list of writers that have successfully registered with VSS.

The best practice: Run the *vssadmin list writers* command on your servers to double-check the registration for each of their installed applications, and then run it again from time to time to verify that your backups are being captured correctly.

#9 – Verify VMware Tools version and functionality

Verifying successful VSS Writer registration is only one step in ensuring Windows applications will back up correctly. The VSS Writer is instructed to perform its duties when a backup is about to begin, and that orchestration initiates from the VSS Requestor. In a vSphere environment, the activities of the VSS Requestor are commonly handled by the VMware Tools. The correct version of these tools must be installed and operational in each VM for backups to function. This process in vSphere is not necessarily automated.

The best practice: The Virtual Machines tab in the vSphere Client includes two hidden columns titled *VMware Tools Running Status* and *VMware Tools Version Status*. Right-click the column header and add these columns to your view. You should keep a close eye on your VMware Tools status.

#10 – Seek the right solution for comprehensive application protection

Gone are the days of file-and-folder backups, backup windows and agents inside VM OSs. Today's best practices leverage disk-based backup solutions that integrate virtual disk backups with application snapshots, replication to offsite locations and all the necessary test automation that ensures every backup is good. The right solution includes all of these features in its integrated console, allowing you the flexibility to set up backups—and then forget them.

About the Author

Greg Shields, Microsoft MVP and VMware vExpert, is an independent author, speaker, and IT consultant, as well as a Partner and Principal Technologist with Concentrated Technology. With 15 years in information technology, Greg has developed extensive experience in systems administration, engineering, and architecture specializing in Microsoft OS, remote application, systems management, and virtualization technologies.

About Veeam Software

Veeam® is Protection for the Modern Data Center™ - providing powerful, easy-to-use and affordable solutions that are Built for Virtualization™ and the Cloud. Veeam Backup & Replication™ delivers VMware vSphere backup, Hyper-V backup, recovery and replication. This #1 VM Backup™ solution helps organizations meet RPOs and RTOs, save time, eliminate risks and dramatically reduce capital and operational costs. Veeam Backup Management Suite™ provides all the benefits and features of Veeam Backup & Replication along with advanced monitoring, reporting and capacity planning for the backup infrastructure. Veeam Management Pack™ (MP) extends enterprise monitoring to vSphere through Microsoft System Center and also offers monitoring and reporting for the Veeam Backup & Replication infrastructure. The Veeam Cloud Provider Program (VCP) offers flexible monthly and perpetual licensing to meet the needs of hosting, managed service and cloud service providers. VCP currently has over 4,000 service provider participants worldwide. Monthly rental is available in more than 70 countries from more than 50 Veeam aggregators.

Founded in 2006, Veeam currently has 23,000 ProPartners and more than 91,500 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland and the company has offices throughout the world. To learn more, visit http://www.veeam.com.

VEEAM

IT JUST WORKS!™

Protection for the **Modern** Data Center

To learn more, visit **http://www.veeam.com/backup**