

Realtime
publishers

Virtualization Backup and Recovery Solutions
for the SMB Market
The Essentials Series

10 Must-Have Features for Every Virtualization Backup and Disaster Recovery Solution

sponsored by



Jaime Halscott

Introduction to Realtime Publishers

by **Don Jones, Series Editor**

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers..... i

10 Must-Have Features for Every Virtualization Backup and Disaster Recovery Solution 1

 Exploring the Benefits of Solutions Designed for Virtualization 1

 1. Near-CDP 1

 2. Dramatically Better RTO and RPO..... 2

 3. Virtualization Vendor Support..... 3

 4. Integration with the Virtualization Management Platform..... 3

 5. Ease of Installation and Use 3

 6. Verification and Security of Backups 4

 7. Flexible Backup Storage Options 4

 8. Efficient Use of Storage..... 5

 9. Migration Support 5

 10. Cost Effectiveness 5

Is Cross-Platform Support Needed? 6

Copyright Statement

© 2011 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

10 Must-Have Features for Every Virtualization Backup and Disaster Recovery Solution

Virtualization platforms are different from physical platforms. The virtualization benefits extend to management and support for tightly integrated backup and recovery solutions. Solutions designed exclusively for virtual infrastructure backup provide a host of key benefits and should be leveraged whenever possible.

Exploring the Benefits of Solutions Designed for Virtualization

The nature of virtualization means that hosting several different guest virtual servers running everything from Windows to Linux to custom appliances is not only possible but so much easier than on bare metal. It also means that the ability to centralize backup at the host level is going to be possible.

A virtualization backup solution is ideally going to be host based. The previous article discussed the problems with cost and management of maintaining guest-based backup solutions. Licensing and support costs negatively impacted return on investment (ROI) for guest-based solutions. Host-based virtualization backups offer a tremendous benefit in being able to provide consistent images without the need to “touch” the guest servers.

Ideally, a solution designed for virtualization should have the same feature-richness associated with an image-based backup solution for physical servers, plus direct integration with the virtual platform to take highest advantage of it with the least processing and administrative overhead. There are several must-have features when looking into virtual server backups. These can mean the difference between just surviving a disaster and thriving during it. Consider the following ten must-have features when looking for a solution to back up virtual infrastructures.

1. Near-CDP

CDP stands for continuous data protection. CDP is one of those features that has driven the adoption of new backup solutions. In the early days of backups, you had periodic full backups followed by incremental or differential backups. Although the incremental and differential backups served the functions of shrinking backup sizes and time to complete backups, they still weren't backups that occurred with any more frequency than daily. Yes, there are certain applications that companies were backing up more frequently, such as SQL databases using transaction logs, but for general purpose backups, the full backup scheme supplemented with incrementals or differentials was the norm.

The challenges with daily backups were obvious from the beginning. Any need to recover would be limited by the time since the last backup. This is where CDP, or realistically near-CDP, came in.

By changing to much more frequent backups using near-CDP, the backup sizes didn't grow appreciably and the impact on servers decreased. Changed files or parts of an image were backed up only as they changed. Virtualization platforms such as VMware's vSphere include changed block tracking (CBT), which means they are ideally suited for near-CDP backups. Platforms using CBT determine which parts of a disk have changed so that the backup program knows what needs to be backed up. Although the actual implementation of CBT is more complex than that, the benefits are simple enough to understand.

The virtualization backup solution should support some mechanism for scheduling that enables near CDP. This could be nothing more than a setting to allow backups every minute or every few minutes. This is close enough to be considered continuous.

2. Dramatically Better RTO and RPO

Recovery time objective (RTO) and recovery point objective (RPO) are functions of how often and how quickly backups and restorations can occur. In order to achieve dramatically better RTO and RPO, a virtual backup solution needs to be able to back up virtual servers almost as quickly as data changes, keep that data in a rotation scheme that allows recovery from as far back as possible, and be able to restore service as fast as possible.

The near-CDP function discussed earlier begins to help with the RPO aspect. In particular, this helps prevent a front-end gap in data protection. The often more difficult challenge is recovery to a point further back in time. Frequently changing systems may only be able to keep a limited number of days for backups. Once recovery is required from weeks or months back, there becomes a problem. In order to support far better RPO, a virtual backup solution needs to address this challenge.

In order to meet long-term retention goals in support of distant past restorations, a virtual server backup solution needs to support built-in schemes for retention such as grandfather, father, and son. In this manner, a solution can keep near CDP for the day, then daily backups for a given time period, then weekly backups for a given time period, then monthly backups for a given time period, and so on. This allows for frequent backups and historical recovery.

RTO is dramatically shortened by the unique ability of a virtual server backup solution to run a live copy of the backed up virtual machine. In the event of a host or guest failure, the backup copy of the virtual server can be run without the need for restoration. Ideally, the solution should allow the changes that occur to then be saved back to the backup image. This technology allows RTO to be nearly instantaneous.

3. Virtualization Vendor Support

The vendor of the underlying virtualization platform must support the backup software's methods of backup and restoration. Vendors like VMware have carefully designed support for application programming interfaces (APIs) that enable third-party virtual backup solutions to back up and restore virtual machines properly. This ensures that the running virtual machines are in a consistent state, that snapshots are done correctly, and that applications running inside the virtual machines aren't damaged.

The selected virtual backup solution must use these APIs correctly and handle errors gracefully so as not to break the virtualization platform. The importance of using a vendor-supported backup method goes even further. Trying to get support for the operating systems (OSs) inside the virtual platform or the virtual platform itself may not be possible if an unsupported backup method is used. To choose a solution without proper vendor support is not worth the massive risk.

4. Integration with the Virtualization Management Platform

Integration with the management platform can be done in one of two ways. Integration can mean a plug-in that sits inside of the management console, such as an application inside VMware's vCenter Server, or a complimentary solution that runs as a virtual machine itself and is thus managed inside of the virtualization platform console.

Although direct integration with vCenter is a nice feature to have, it could be one more item to break. Being able to run a backup solution inside of a self-contained virtual machine is going to provide a host of benefits. The ability to deploy an open virtualization format template (OVF) means that your deployment will be the same as all of the others out there. This makes supportability much better. It also leads to the next must-have feature.

5. Ease of Installation and Use

Deploying backup solutions seems like one of the toughest challenges that an administrator faces. Get it wrong, and it will lead to support challenges at best and failed backups at worst. There is also a tremendous cost with setting up backup servers, patching them, installing backup software, and managing it.

Being able to deploy a backup solution as an appliance makes the deployment little more than a few clicks. It also virtually guarantees that the backup solution deployment is consistent and supportable. Ease of use is the other side of the usability equation. Too many options really don't make any sense in a virtualization backup solution. All you want to do is ensure that the virtual servers are backed up fully and quickly with fast restoration as needed.

Administrators want to use products that don't have a steep learning curve. It also means that administrators can jump right in and use the product with support and training costs that are minimal. Everyone knows famous tales of software that costs next to nothing with training that costs many times more, or even worse, software so complex to use it becomes "shelfware."

Many administrators prefer to use standard Web browsers to manage software. This gives administrators flexibility to access and administer backups from a wide range of devices such as smartphones and tablets. Proprietary management consoles tend to be more difficult to use, especially given the rich Web experience that has become prevalent in management tools.

6. Verification and Security of Backups

Backing up is fine, but the days of taking the software's word for its validity are long gone. A virtualization backup solution has to be able to not only tell you when backups are successful but also provide a means of knowing through positive action that the backups will work.

Organizations are also under increasing regulatory pressure to keep data secure. Internal security guidelines should include a mandatory encryption requirement to keep backups secure. Remember that this is a complete copy of your critical infrastructure. In the wrong hands, it can do as much damage as having no backup could in a disaster recovery scenario. Encryption of the backups also needs to use a powerful cypher like AES so that protection isn't just a false sense of security.

A backup solution for virtual servers must include the ability to automatically verify backups to ensure that they can be restored. There is no room for the excuse that the backup said it completed successfully when a restoration fails. Ideally, automated lab testing gives further piece of mind.

7. Flexible Backup Storage Options

Backups used to be almost exclusively to tape. Then disk-based backups came along with more speed and high reliability. The selected virtual backup solution needs to be even more flexible. Backup to disk and over the network to SAN or NAS storage are must haves. However, there are some other storage options that are now required to meet organizational needs. These include LAN-free SAN backup and cloud backup.

LAN-free backup is no different than other disk-based storage, but with LAN-free backup, the backup traffic stays inside the SAN. Network traffic is virtually eliminated. This speeds the time to back up and reduces the overall load on the network.

Cloud-based backup should also be considered a must-have option for storage flexibility. Flood, fire, theft, and a host of other natural disasters mean that virtual server backups must be stored away from the virtual infrastructure. Historically, this meant physically taking disks or tapes off site. This approach is prone to human error and the risk of loss or theft of media is ever present. The cloud has emerged as a much more reliable and cost-effective option for getting virtual server backups to a safe offsite location. This makes it a must-have feature.

8. Efficient Use of Storage

Features like compression and deduplication have long been found in physical backup and recovery products. Compression has been part of solutions for longer than deduplication, but deduplication has proven to be a stable addition to backup solutions. These are only now beginning to see the light of day in virtual environments. Disk space may be cheap as a general rule, but recent upticks in hard drive pricing and the high cost per gigabyte of SAN storage means that efficient use of space is crucial. This can often be cost justification enough to buy a compression and/or deduplication solution. If it saves enough space in cost per gigabyte, then a cost model can be built to justify the expense. The savings of deduplication aren't as easy to quantify as compression, but every little bit helps.

9. Migration Support

Products like VMware Converter are fantastic for migrating servers into a vSphere infrastructure, assuming that all of the factors are met. In a lot of cases, a server might be too old or of the wrong platform to work correctly with this P2V scenario. A virtual server backup solution should support some means of migrating the physical servers into the virtual infrastructure. Although this isn't a part of the backup role per se, it does factor into supportability.

If a backup vendor provides a means of migrating from physical to virtual (P2V) or virtual to virtual (V2V), that vendor certainly will not have problems supporting the migrated virtual servers. By including a migration support method into the backup solution, a backup vendor will certainly stand behind the technology used and that the stability of the system wasn't compromised by the migration.

10. Cost Effectiveness

Yes, your data is the lifeblood of your business, and as much as 50% of IT budgets go towards backup and disaster recovery solutions, but that doesn't mean that a solution should be overly expensive. IT budgets are shrinking, but the volume of data and the number of virtual servers is increasing. Reconciling this is going to require a solution that provides standard features at a lower price than competitive offerings.

Even saving a few hundred dollars per license starts to add up quickly. Cost effectiveness is also a function of the virtual backup solution's licensing model. A virtual backup solution should be licensed on a per-host or per-processor model. This way, a business can get the best ROI for a highly consolidated infrastructure. This also encourages a business to invest in hardware that supports higher virtual machine density and the resulting lower support costs. It is a win-win. Avoid any solution that licenses on a per virtual machine scheme.

Is Cross-Platform Support Needed?

One feature that is missing from the 10 Must-Have Features is cross platform support. This might at first seem like a glaring omission, but in practice, it really isn't. The VMware ESX/ESXi hypervisor is the most mature of the technologies out there, and it is estimated that as much as 85% of virtual server deployments occur on this platform.

Organizations also are working diligently to standardize on a single hypervisor in order to keep ongoing support costs down whether that is using VMware, Microsoft's Hyper-V, or Citrix's Xen Server, among others. It just makes good economic sense to stick with a single virtualization platform.

Although a virtualization backup solution would certainly be thorough if it could support various different hypervisors from different vendors all under the same console, it would quickly get complex, costly, and difficult to support. This would defeat the purpose of implementing many of the earlier discussed must-have features for a virtual server backup solution.