

Realtime  
publishers

## The Essentials Series: The Evolving Landscape of Enterprise Data Protection

### Virtual Machine Environments: Data Protection and Recovery Solutions

*sponsored by*

**syncsort**

by Dan Sullivan

---

Virtual Machine Environments: Data Protection and Recovery Solutions.....	1
Virtual Machine Environments in the Enterprise .....	1
Adapting Backup Strategies to Virtual Machine Environments.....	2
Conventional Backups .....	2
Specialized Virtual Machine Backups .....	3
Source-Side Deduplication Backups.....	3
Leveraging Virtual Machines for Disaster Recovery .....	4
Summary .....	5

---

## Copyright Statement

© 2009 Realtime Publishers, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers, Inc. (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, Inc or its web site sponsors. In no event shall Realtime Publishers, Inc. 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](mailto:info@realtimepublishers.com).

---

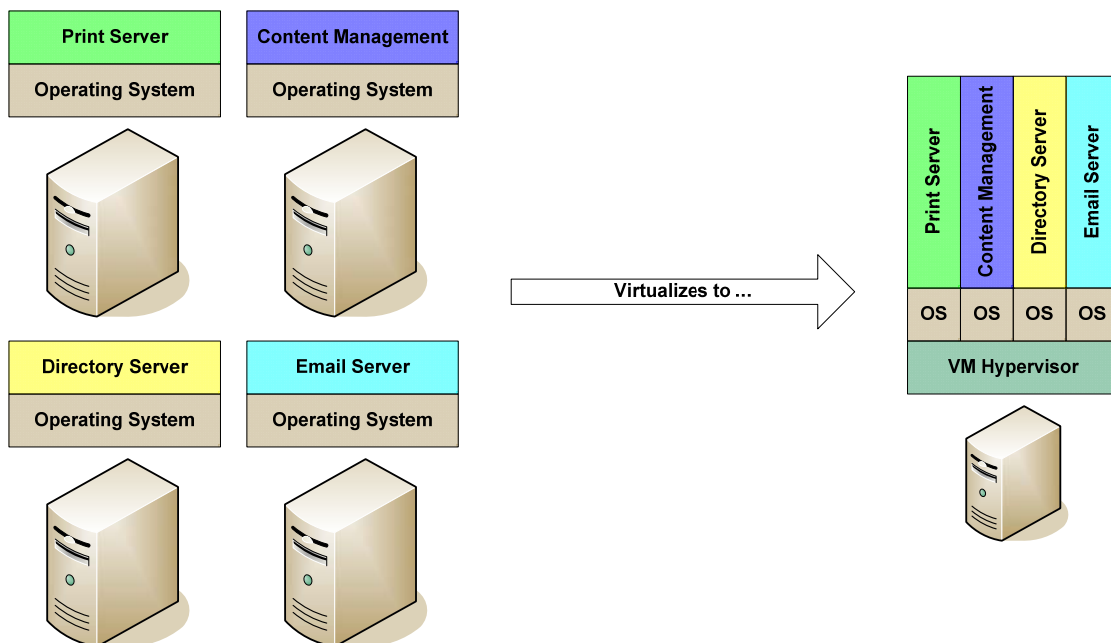
# Virtual Machine Environments: Data Protection and Recovery Solutions

---

Virtualization technology, widely adopted for its cost efficiency and ease of administration, has implications for data protection and recovery solutions. The data and applications running on virtual machine environments must be protected just as those on physical servers, but differences in physical and virtual servers have implications for your backup strategies. At the same time, virtual machine technology can significantly contribute to streamlining disaster recovery operations. In both cases, proper planning and deployment of virtual machine environments and backup solutions can positively influence the return on investment (ROI) in these technologies.

## Virtual Machine Environments in the Enterprise

Virtual machine technologies such as VMware's Infrastructure 3 and VMware Server and Microsoft's Hyper-V technology, available in Windows Server 2008, are gaining in popularity. Virtualization technologies such as these insert a software layer between the hardware and the operating system (OS), which traditionally ran directly on that hardware. The new software layer, known as a hypervisor, can act as a mediator between CPUs, memory, storage, and other hardware components and one or more OSs.



**Figure 1: Virtual machine environments enable server consolidation without sacrificing the benefits of running applications in dedicated OSs.**

---

Some of the key drivers behind this adoption are

- The ability to consolidate multiple applications on a single physical server while retaining the advantages of running applications on separate OSs
- Reduced power consumption due to fewer servers and less demand for cooling
- Reduced cost and time to provision application servers on a virtual machine versus on a physical machine
- Ability to reduce cost of disaster recovery

With these advantages, though, come differences in the way typical IT operations are performed. Backups, in particular, require special attention.

## **Adapting Backup Strategies to Virtual Machine Environments**

Moving to virtual machine environments does not change the need to protect these servers at the same level that physical machines are protected. Although virtual machines can be installed from images to new hardware (“raw iron”) faster than OSs and applications can be installed and configured from scratch, there is still a need to protect the data on virtual servers. It is reasonable to presume that virtual machines will require similar service level agreements (SLAs), recovery point objectives (RPOs), and recovery time objectives (RTOs) to those of physical servers.

### **Conventional Backups**

Conventional backup methods, however, can adversely impact some virtual machine environments and slow the performance of other virtual machines. For example, if one virtual machine is running a full backup, the performance of other virtual machines on the same physical server may be so degraded that they no longer meet their response time objectives. In other cases, a systems administrator may be forced to choose between delaying the backup of one virtual machine to allow another virtual machine’s backup to complete or risk degrading the performance of all virtual machines on the physical server.

---

## Specialized Virtual Machine Backups

Specialized solutions can improve upon conventional backup methods but at the cost of increased complexity. Take, for example, solutions that offload the backup process to another server. This type of solution offers the ability to perform full and incremental backups of virtual machines, make full image backups of virtual machines, and perform these operations from a centralized management system. It does require some changes to the typical backup system:

- Needs a backup proxy system that stores a snapshot of a virtual machine
- Conventional backup software performs backup from proxy server to tape or disk storage
- Restoration becomes a two-step process with backup software restoring to the proxy device; that image is then restored to the target server

This method eliminates the overhead of performing backups from the virtual machines and consolidates it on the proxy server. It does, however, introduce another component (the proxy server) to manage. Another approach is to reduce, rather than shift, the overhead of creating backups.

## Source-Side Deduplication Backups

Backup overhead is reduced when the amount of data copied, transmitted, and stored is reduced. There are a number of ways to do so, including source-side decompression and target-side decompression; this article will discuss the efficient source-side deduplication method.

See the first article in this Essential Series for details about source-side deduplication compared with other data reduction strategies.

With source-side deduplication, block-level incremental backups are performed. This is more efficient than file-level backups because only changed blocks rather than entire files are backed up. As a result, backups finish faster with less demand on CPU, disk, and network resources. This approach eliminates the need for a proxy server to reduce the load on virtual servers, which, in turn, eliminates the two-step recovery process required with a proxy-based solution. Fewer recovery steps means less complexity, less room for error, and faster recovery times.

## Special Issues with Databases and Block-Level Backups

It should be noted, though, that some applications, such as databases, may perform block-level operations of their own within the context of logical transactions. Block-level backup systems must take this into account to ensure that backups represent a consistent database state. If this is not possible, backups should be performed only when there are no changes being made to the database.

---

Depending on the backup solution and the hypervisor constraints, a data protection system with source-side deduplication may also offer:

- Bare metal recovery
- Support for multiple OS platforms
- Application-specific recovery options for databases and email servers

The last feature is especially important if you are required to recover a single user's email folder or a single table within a database. Restoring an entire database, for example, in order to restore a single table can add significantly to the time and space required to perform a restore operation.

Solutions that reduce overhead and reduce complexity yield favorable ROI for a variety of reasons:

- Small backup footprint allows for more frequent backups, which provide many recovery points
- Management costs are reduced because there is no need for an additional proxy server
- Reduced backup sizes reduce storage costs
- More virtual machines may be run on a single server because backups do not introduce unnecessary load on the servers

Optimizing backups for virtual machines is just one way to improve overall ROI with virtual machines and it contributes to more effective disaster recovery as well.

## Leveraging Virtual Machines for Disaster Recovery

As the number of virtual machines in business grows, so does the need for disaster recovery support. When selecting a backup solution, look for features that support creating virtual machines from a single backup instance. This approach brings several advantages. Perhaps the most important from a disaster recovery perspective is the ability to rapidly recover both virtual and physical servers with the instant availability of virtual machine images. Certain data protection products allow you to iSCSI-connect and LUN-map from the virtual server to the backup image, creating a temporary fully operational virtual machine without even a data transfer.

There are cost advantages as well. The ability to recover physical servers to virtual machines reduces the cost of maintaining a disaster recovery center. For example, if a production email server running on a dedicated server completely fails, a virtual machine can be rapidly deployed to a shared server in a disaster recovery site. Performance may be slower but email services can be maintained until a production-level server is available. In addition, jobs, such as backup jobs, can continue to complete even after migration.

---

The overall impact is the need for fewer servers in a disaster recovery site. Rather than a worst case scenario of having to support a failover server for every production server that runs a single application, you can consolidate both production and disaster recovery servers. Presumably, virtual machines are already used in production to optimize the number of physical servers needed to meet performance requirements and SLAs. If those SLAs allow for reduced performance during disaster recovery periods, you could further reduce the cost of maintaining a disaster recovery site by running more virtual machines on physical servers during recovery periods. Favorable ROI is also realized by using a shared backup and disaster recovery procedure: the same backups used for data recovery are available for disaster recovery.

## Summary

Virtual machine environments are a staple of enterprise IT infrastructure. The cost and management advantages of virtual environments are compelling, so backup and disaster recovery strategies should adapt to their particular requirements. In the process, though, you find that the same techniques you deploy to provide backups for virtual machines without undo overhead on physical servers yield favorable ROI. The lower overhead performance of source-side deduplication allows for more resources dedicated to virtual machines and their applications and reduced storage costs on backup devices. In addition, the same procedures provide for cost-effective disaster recovery.