

Realtime
publishers

The Essentials Series: Managing
Workloads in a Virtual Environment

Virtual Disaster Recovery

sponsored by

Novell[®]

by Jaime Halscott

Virtual Disaster Recovery	1
Virtual Versus Physical Disaster Recovery.....	1
Platform Independence	2
Time to Recovery.....	3
Summary	5

Copyright Statement

© 2009 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 or 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.

Virtual Disaster Recovery

Disaster recovery is a tremendous challenge to businesses of all sizes. Maintaining system uptime and availability is difficult whether the servers are physical or virtual.

Virtual Versus Physical Disaster Recovery

Traditional disaster recovery options have always been limited for a physical server infrastructure. A physical server has several drawbacks, the most difficult being that it presents a specific set of hardware devices to the operating system (OS) that require drivers tailored to those components. This factor greatly limits the flexibility of backup and recovery options. Traditionally, backup solutions focused on the data and applications on the server, so hardware wasn't such a concern. Backup solutions were fairly effective in this task, but the time required to recover an entire server was painstakingly long. A server would have to be rebuilt and then restored to. This process could take several hours even in a best-case scenario.

The success rates for backups were also not particularly high in a physical world. Recovery required highly reliable media and a system that worked correctly within the operating system itself. Rebuilding a server operating system didn't always guarantee that media libraries from the previous server install would even be recognized, much less easily restorable. Tape systems represented low cost per gigabyte and acceptable speed, but the media was prone to stretching and errors.

To address these concerns, image-based backup solutions were introduced. These solutions offered an increased chance of restoration because the entire server could be restored directly to bare metal. Traditionally, these solutions haven't always supported all available hardware, but in recent years, broader platform support has been built-in. When considering using an image-based recovery solution, it is extremely important to use one that is designed to work with the hardware and software platforms in use. A solution should support both physical and virtual platforms to allow seamless transitions between the two.

Physical servers often employ third-party software backup solutions to supplement the native tools found in the OS. These third-party backup products have costs associated with both the licensing and training required to master them. The backup tools built-in to the OSs usually don't scale well to larger organizations and are meant to offer basic protection and business continuity for individual servers. It becomes additionally costly and complex to manage backup of multiple platforms such as businesses that also have Linux and UNIX-based servers.

Platform Independence

Virtual servers lend new flexibility to the concept of disaster recovery. Virtual servers are platform independent. Virtualization platforms such as Microsoft's Hyper-V or VMware's ESX Server present themselves to the virtual server as the same hardware whether or not the server is from the same generation or the even the same vendor. Using virtualized hardware, the virtual server becomes a portable "image" that can easily be copied between physical servers running the same virtualization platform. The virtual servers appear to the underlying OS or virtualization platform as a single file or a small group of files.

Virtualization platforms also support the use of snapshots. Snapshots use delta files to capture the changes between the virtual server virtual hard disk baseline and all changes made since the last previous snapshot. This gives a sort of instantaneous undo functionality that isn't possible in a physical server. Snapshotting technology is a core feature of all major virtualization platforms. Snapshots give IT personnel the ability to deploy updates and software with more confidence knowing that software-related problems can be instantly undone. Unlike technologies that run inside the OS, snapshots are part of the virtualization platform and are not affected by the stability of the OS itself. Live continuous backup solutions designed for physical servers still suffer from this downside.

Virtual disaster recovery is only limited by the tools that are available to manage the virtual server files. Although this can be as simple as copying the files to another location and copying them back, true disaster recovery solutions need to be more robust and support automation. Third-party tools offer a number of enhanced features to automate the process of backup and recovery as well as provide advanced features for replication.

One key feature to consider with a virtual disaster recovery solution is the ability to keep a warm spare environment. This lets the primary production virtual environment stay synchronized with a backup virtual environment. In the event of a primary failure, the secondary system can perform where the failed one left off. Ideally, the solution selected should be able to do this based on predetermined thresholds, such as a heartbeat or other means of detecting that the primary environment is offline. Also, depending on the size of the primary environment and how severe the failure is, the solution needs to be aware of the best means for recovery. It can do so by first profiling the workloads of the servers and giving IT management a picture of exactly what the environments look like over a particular time period. Knowing this information is vital to a successful failover. In this scenario, the time to recovery is extremely fast.

Time to Recovery

A recent trend in the virtual world has been to boot from the Storage Area Network (SAN). The ever-decreasing cost of shared storage products such as iSCSI SANs have enabled the virtual server disk files to exist on shared storage. This greatly reduces the time to recover. When a physical host has a virtual server disk mounted and running, it is accessing the file directly from the SAN. If this physical host fails, a backup physical host needs only start the virtual server and it can continue from where it left off. Recovery is seconds or minutes, not hours. Third-party tools exist to make this process transparent. Solutions should support thresholds that can easily be set to create an automatic failover.

Physical disaster recovery solutions have typically used tape and more recently disk-based storage for backup targets. The failure of hardware was more of a threat than software-related disaster. Tapes presented a cost-effective storage media with good speed, but reliability of the media was always a concern. Disk-based backup addressed these concerns in the past few years when disk drive storage costs dropped to a level competitive with tape. Additionally, options for shared storage using iSCSI instead of Fibre Channel have greatly reduced the entry costs for high-performance shared storage. Utilizing off-the-shelf adapters and switches, iSCSI solutions are quickly becoming the dominant method for accessing large volumes of shared storage. Also, iSCSI software support is built-in to the major OS platforms.

However, both tape- and disk-based backup solutions are slow to recover large server workloads or enterprise-wide failures. For companies needing faster reaction to an outage, the only option was to keep similar, or ideally, an exact copy of the production environment in terms of hardware and software. This is an expensive proposition. Duplication of the production physical environment had to be as identical as possible to prevent problems such as hardware compatibility conflicts. In most cases, this redundant hardware was underutilized or in a worst case scenario, completely dormant. This supporting hardware can help form the backbone of a virtual server infrastructure.

Once the transition has been made to a virtual environment, the complex and costly process of duplicating hardware and maintaining same-generation spares is no longer an issue. Hardware for the recovery environment does not need to be identical or even as powerful as the production systems. Due to the vendor- and platform-agnostic nature of virtualization, the servers for recovery can be those that represent the best bargain for the needed output. Companies are not locked into any single hardware provider either. Additionally, the physical servers used for the virtual recovery can themselves be considerably less powerful than the production models. Because resources such as memory and processor can be dynamically provisioned in a virtual server, IT can choose just how much a given workload needs. Resources can be scaled up or down to ensure the proper level of service for any given server workload. Once the initial workload profiling has been done, this process becomes much simpler.

Depending on the service level agreements (SLAs) or internal policies, allowing diminished performance in the warm spare virtual environment can greatly reduce costs. The number of standby servers can be far less than the primary network calls for. Processors can be slower and memory can be less as well. In this scenario, business continuity is of greater importance than quality of service.

One of the greatest advantages of a virtual infrastructure is platform independence. Although virtual servers do require some specialized hardware in order to see the greatest benefits in terms of speed and integration, these specialized components are built-in to almost all modern servers. Modern, hypervisor-based virtualization platforms require processors that have virtualization offload support built in, such as Intel's VT or the AMD-V technologies. Servers should also utilize the x64 technologies from Intel or AMD in order to support a high number of virtual machines on the same physical hardware. The 4GB memory limitation of 32-bit hardware limits the usefulness of a physical server to support virtual servers by allowing it to run only a small number of virtual instances. Almost any modern server will have hardware support for these virtualization technologies. Increasing the amount of RAM available to virtual servers on the physical host will be the single largest factor in determining how many virtual guests' workloads can be run at any given time.

Any server designed for virtualization can be from any vendor. A business doesn't need to be locked into any model or generation of server from any single provider. As long as the servers meet the requirements for the virtualization platform selected, they can be utilized with the same benefits. Performance under a given workload will be the only variable. Businesses will find it easier to transition between server hardware providers or select a best-of-breed solution using servers from multiple companies. Servers used for disaster recovery can come from just about any vendor.

There are still software considerations to take into account when planning virtual disaster recovery. Native tools from the major virtualization platform vendors provide little more than basic recovery options. Some of the native backup solutions will integrate with the virtualization platform itself to allow live backups, but in the medium-sized business or enterprise space, a customer will require a greater level of control. This control must be centralized in order to ensure consistent and cost-effective management of the disaster recovery process. Tools provided by the virtualization vendors only support their respective platform. Disaster recovery between different hypervisor platforms isn't possible with native tools. When choosing a software solution to assist with backup and recovery, it is vital to select one that supports multiple virtualization and hypervisor platforms, multiple OS platforms, and multiple hardware solutions.

Summary

Time and cost to recovery in a virtual server environment is greatly reduced versus a strictly physical server environment. The chances of recovering a functioning virtual infrastructure also greatly increase over a physical server infrastructure. New options for disk-based shared storage lend flexibility to recovery. Now recovery isn't simply occurring within the server itself but becomes a function of a system of distributed workloads. Keeping the workloads running is the most important goal. Virtualization provides disaster recovery solutions to accommodate.

When making the transition to virtualization, the disaster recovery solution requires stringent planning. The first step is to establish a service level for the servers in the organization. This doesn't need to be the same for each server but does need to accurately reflect the roles of the servers. Consider establishing a Service Level Agreement (SLA) that clearly defines the required uptime.

The next step is to determine the recovery objectives for your servers. This breaks down into three areas:

- Recovery Time Objective (RTO)—This is the measure of downtime
- Recovery Point Objective (RPO)—This is the measure of data loss
- Test Time Objective (TTO)—This is the measure of testing ease

Using these measures combined with the goals of your SLAs, you should select the best disaster recovery solutions for your virtual environment. The following table illustrates the appropriate technologies and the impact of choosing each:

Solution	Cost	RPO	RTO	TTO
Server Clustering	\$\$\$\$\$\$	Near Zero	Near Zero	Near Zero (Impacts production data, ads risk)
Consolidated Recovery (virtualization)	\$\$\$\$	Minutes	Minutes	Minutes (No impact on production data)
Imaging (virtualization)	\$\$\$	Hours	Hours	Minutes (No impact on production data)
Image Capture	\$\$\$	24h	Hours	Hours (Requires additional hardware)
Tape/Manual Rebuild	\$	24h+	Days	Days (Not practical)

Properly planning a virtual disaster recovery solution requires an understanding of the potential downtime, the time to recover, the cost of the solution and the platforms supported. By taking these factors into account and leveraging the best features of virtualization, an organization can get faster and more cost-effective solutions for disaster recovery.