



The Benefits of Virtualization for Your DR Plan

Server virtualization is making a positive impact in the area of disaster recovery. Through virtualization, organizations can save money, as well as increase application recovery time. This guide will show how utilizing virtualization technologies can improve your DR plan, and will also serve as a checklist on what to be aware of when choosing the right solutions for your virtualized environment.

Sponsored By:



The Benefits of Virtualization for Your DR Plan

Table of Contents:

[Server virtualization may have big disaster recovery payoff](#)

[Leveraging storage replication for VM disaster recovery](#)

[Resources from VMware](#)

Server virtualization may have big disaster recovery payoff

By Howard Marks

February 20, 2008

While most of the buzz around server virtualization in general, and VMware Infrastructure in particular, have been about server consolidation and greening the data center, disaster recovery may be the IT area where server virtualization technology has the biggest impact.

Disaster recovery (DR) planning for mission-critical applications historically called for replicating the data for these applications and having servers standing by at the DR site ready to take over at a moment's notice.

Most organizations can save money by virtualizing these standby servers. A single offsite server can act as the standby domain controller, SQL server, Exchange server and several more. Not only can you save the cost of all those physical servers, but also the rack space and power charges from your DR site.

Saving money and still providing the same level of protection that your old expensive physical server solution could is a good thing. But the real payoff is improving the recovery time of the applications that you wouldn't dedicate a standby server to. Most organizations soon realize they can move some applications up from the secondary tier to having standby servers, since the standby servers are essentially free.

Solving bare metal restore to different hardware

In the "old days," secondary applications were limited to restore from tape as their protection model, resulting in multiday recovery points and recovery times. Even if you were replicating the application's data, it wasn't always possible to get an identical server to restore the application backup to. You either had to go down the dank dark path of a bare metal restore to different hardware, or pursue a new OS and application install, all the while hoping you had a record of all the patches needed to mount that database.

The "different hardware" problem is solved because virtual machines are indeed virtual machines—they all run with the same set of drivers and can't tell if they've been moved from one host to another. In addition, virtual machine snapshots from VMware or even Microsoft's Virtual Server or Hyper-V are just files, so restoring a virtual machine is just a matter of mounting the files on a new host.

Rather than relying on tape transfers, you can schedule snapshots of your virtual machines and transfer them to the DR site over the replication link. And if your network guys can prioritize traffic properly, it won't interfere with real-time replication.

The real fun comes when a disaster is declared and you have to start switching over to the standby servers. Because the suspenders-and-belt crowd set up their DR infrastructure to be able to take over at full speed the minute the switch was thrown, their DR site has lots of compute horsepower. (Of course lots of horsepower means lots of money.)

The more frugal companies take advantage of VMotion, which moves virtual servers from one host to another dynamically while they're still running and, in addition, DR providers like SunGard's "shared server" offerings. With shared servers, you pay a few shekels to the DR provider every month for the right to claim servers out of their stock at the DR site when you declare an emergency. Once you declare that, you get the servers for your exclusive use and can install VMware ESX on them.

Then, once the new hosts are up, you can use VMotion (or even better VMware DRS) to dynamically allocate virtual servers to hosts based on load and to mount your virtual servers on the new hosts. This will boost your application performance... probably before your users can get to their new workplaces to use the applications.

Note: The same trick—albeit with longer recovery times—can be used with vendor's server quick ship programs that will ship you new servers in the event of a disaster.

About the author: *Howard Marks is chief scientist of Networks Are Our Lives Inc., a Hoboken, N.J., network and storage consulting and education firm. Marks' company specializes in bringing the infrastructures and processes of mid-market firms up to enterprise standards in the areas of systems, network and storage management, with a focus on data protection and business continuity planning. Marks is the author of three books and more than 200 articles on network and storage topics since 1987. He is a frequent speaker at industry conferences.*



With disaster recovery,
speed is the name of the game.



Take your business to a whole new level with VMware Infrastructure.

Disasters come in all shapes and sizes-and it's your job to be prepared for all of them. But until now, you didn't have an option that provided the speed and reliability you needed without prohibitive cost and complexity. That was before VMware.

With VMware virtualization, you get an automated solution that's easily tested, hardware-independent and offers complete infrastructure protection.

All of this, and more, with an unprecedented level of simplicity.

VMware Disaster Recovery. Sounds like a good day to be in IT.

Find out more at www.vmware.com/go/fast

Leveraging storage replication for VM disaster recovery

By Chris Wolf

February 20, 2008


Storage replication is a popular method for synchronizing production and disaster recovery (DR) sites in virtual server environments. If you're either using array-based replication or leveraging a storage virtualization appliance for replication, there are several variables that will influence the efficiency of your storage topology as it relates to DR.

When you're sizing up storage solutions for DR, you should consider five issues:

- Vendor support
- Storage architecture
- Replication options
- Deduplication or single instance storage support
- Recovery options

Of course, there are several ways to get data from a production site to a DR site. Rather than simply give a high-level overview of these alternative virtual machine (VM) replication methods, this article will take a deeper look at specific storage array considerations. However, when it comes to architecting replication for virtual environments, this article can only scratch the surface. Many storage and DR optimization tricks are vendor-specific. Be sure to check your storage and server virtualization vendors' documentation and architecture guides for details relevant to your particular environment.

Let's set a baseline by assuming the high-level storage replication architecture shown below. Note: The network storage could be network attached storage (NAS) or either a Fibre Channel (FC) or iSCSI storage array.



All major network storage vendors offer tools for replicating data on an array from one site to another. Most of them use asynchronous replication for site-to-site network storage synchronization, since the WAN network throughput or distance between sites is usually inadequate for synchronous replication. With asynchronous replication, writes are committed to primary storage, then replicated based on the replication policy set by the storage administrators.

Vendor support

Although most storage array vendors offer some form of asynchronous replication, the choice of array vendor nevertheless usually matters. When evaluating storage options, vendor support is a key criteria. A storage array should be supported on products from your environments virtualization vendor and OS vendor. Support should also be considered for enterprise application vendors that name supported storage platforms. Storage platforms that leave a portion of your infrastructure unsupported constitute a risk.

You should also look at your backup vendor's list of supported storage platforms. Many enterprise backup products are capable of managing snapshots on most popular network storage platforms. A storage platform that integrates with your existing data protection software should be given more consideration than one that that does not.

Storage architecture

The way in which storage is architected to support virtualization can have a dramatic effect on replication performance, and thus DR response. Fault-tolerant capabilities via RAID support are required, as any storage array should be deployed as RAID level 5 at a minimum.

In terms of DR response, you need to look at how each VM's virtual disk storage is allocated, as well as how temporary file locations are configured in each VM's guest operating system. When a storage array is configured to support virtualization, you should set aside a volume set for transient or temporary data. How you deal with transient data should be determined by the service level requirements of the VMs you support. For VM data that is synchronously mirrored over dark fiber between two locations, certain application- or service-centric temporary files may be critical and will need to be replicated too. However, for VMs that are asynchronously replicated to a DR site, in most cases replicating temp files would be a waste of bandwidth and storage space.

Getting back to the storage configuration details, assume you've set aside enough volume space (e.g. storage LUN, NFS mount, etc.) for your virtual infrastructure's temporary data. Once the storage for transient data has been allocated, you should configure the virtual infrastructure so that the following files are stored on the transient data volumes:

- Hypervisor swap files;
- Virtual machine guest OS:
 - o Swap file
 - o Pagefile
 - o OS and application temp folders
 - o User temp directories

For individual VMs, you'll need to create a separate virtual hard disk just for transient data, which in turn would be stored on the "transient" volume space of your network storage device. While this may seem like a lot of work, it can result in substantial savings in storage requirements for your DR site, since you won't have to replicate any of the transient data to your DR facility. A VM's pagefile will generally require a high degree of storage I/O, so you may want to use a dedicated virtual hard disk just for the VM's paging file or swap file to gain better control of pagefile quality of service (QoS).

Replication options

Each application's service level requirements should drive the replication requirements of any storage platform. Platforms that offer synchronous and asynchronous replication features, along with block level incremental replication and granular snapshot features, are more likely to meet all of your storage replication requirements. The bottom line should always be the storage solution's ability to leverage replication in order to meet your recovery time objectives (RTOs) and recovery point objectives (RPOs).

Deduplication or single instance storage support

A high number of VMs with identical OSes, applications or services will often reside on the same storage array. Storage nodes with built-in data deduplication or single instance storage support will offer significant storage savings by eliminating data redundancy on storage blocks. Note: To realize these storage savings, the storage array should also support thin provisioning. Otherwise a virtual hard disk file (for example, a .vmdk file on a VMFS volume) would consume all of its allocated space at the time it is provisioned. Thin provisioning would allow the virtual hard disk to consume its assigned storage as the virtual hard disk grows in size. With ESX server, thin provisioning is supported by thin formatting VMDKs.

One of the key benefits to deduplicated storage is that the amount of data to be replicated to the DR site will be significantly reduced, by as much as 60%. You could optimize WAN throughput with a WAN accelerator device, but this won't reduce storage costs. Deduplicated storage will not only reduce the WAN bandwidth needed to replicate storage but will also reduce the total amount of storage needed for a given virtual infrastructure. By reducing the amount of storage you need to replicate, you'll also be able to replicate storage more frequently and thus reduce your RTO.

Recovery options

Many storage arrays only provide volume-level recovery for virtual machines. While volume-level recovery is usually what you need for DR, you should look at storage platforms that offer granular file level recovery for files residing in virtual hard disks. Platforms that offer you the ability to recover previous volumes or previous versions of single files from snapshots allow you to leverage the storage solution for both DR and day-to-day file recovery operations. Such solutions would save on the required storage space for data protection operations, as most file-level backups would be unnecessary since file recovery could come from previous volume-level snapshots.

About the author: *Chris Wolf is a senior analyst for Burton Group and author of several IT books. Check out a chapter on backup from Wolf's book, Virtualization: From the Desktop to the Enterprise.*

Resources from VMware



[VMware White Paper: Transforming Disaster Recovery—VMware Infrastructure for rapid, reliable and cost-effective Disaster Recovery](#)

[Yankee Group: "Disaster Strikes! Is Your Business Ready? Disaster Preparedness for Mid-Sized Firms"](#)

About VMware:

VMware (NYSE: VMW) is the global leader in virtualization solutions from the desktop to the datacenter. Customers of all sizes rely on VMware to reduce capital and operating expenses, ensure business continuity, strengthen security and go green. With 2007 revenues of \$1.3 billion, more than 100,000 customers and nearly 14,000 partners, VMware is one of the fastest growing public software companies. Based in Palo Alto, California, VMware is majority-owned by EMC Corporation (NYSE: EMC) and on the web at <http://www.vmware.com>.