

STORAGE SWITZERLAND

VIRTUALIZATION POWERED RECOVERY



George Crump, Senior Analyst

Backup application vendors are quick to point out that the data protection process is really about recovery, not about backup - sounds very clever. In reality they are partly right. For most legacy backup application vendors recovery is an act of successful data placement, moving data from the backup area to the primary storage area. The job does not end there though, as the application still needs to be returned to service. Recovery is really about a complete restoration of the failed application to users. This is a difficult goal to achieve in the non-virtualized world. Virtualization-powered recovery can lead to a complete restoration of the application but fully exploiting the virtualized environment and achieving these goals may require a data protection application created in the virtual era.

The Challenge of Server Recovery

Recovering even a single server from scratch is challenging. While most backup products now have bare metal restoration capabilities, most customers typically find that these aren't very dependable. For example, a slight change to the video card or ROM Bios in the replacement hardware can cause the restoration to fail. For this reason, most bare metal recovery products also

require that a separate backup job be run to capture the bare metal recovery information.

The typical data center does not have the backup window to be able to support a dual-stage data protection process on a continual basis. As a result this separate bare metal backup job is run sparingly, which leads to a lengthy multi-stage recovery process. In this scenario, the bare metal recovery has to be restored first, and assuming that works, incremental backup jobs need to be recovered to finish restoring the server. If the bare metal recovery fails because of a hardware incompatibility then the backup administrator must manually re-install the operating system as well as go through the normal process of recovering the full backup and the subsequent incremental jobs.

The net result of all of these challenges is an extreme lack of confidence in the full server recovery process. Also these elaborate steps make it hard to justify the time to test the process so that experience in taking these steps can be earned. When a system failure does occur recovery goes slower than it should and mistakes are made because inadequate practice and test time have been invested. In recovery, just like everything else, practice makes perfect.

Recovery Takes a Village

The challenge with a typical recovery is that it rarely involves recovering a single server. When the storage system that supports an application fails or worse, when the site that supports the data center faces an outage, recovery almost always involves several tiers of servers. Beyond just the application server there may be a web server, a database server and infrastructure servers that provide directory services and the like. All of these servers have to be recovered, and often in the correct order.

This makes the above testing scenario even harder since multiple servers need to be recovered as part of the test. As well, sections of the network need to be isolated so as to not conflict with production servers. Once again, with full testing such an infrequent activity (maybe once or twice a year), not enough is really learned to know the recovery process when it must be employed.

This multiserver and multistep recovery process is also too difficult for routine restoration needs like bringing back deleted messages or database records. That difficulty has led to a whole sub-market of products provided by backup software developers, including modules for specific applications, online backups and single message recoveries. These modules add to the cost and complexity of the solution, becoming one more thing to buy and one more thing to learn.

Server Virtualization-powered Recovery

The answer to these challenges lies partly in the server virtualization projects that most data centers already have underway. The missing link is server virtualization-powered backup and recovery products like vPower from [Veeam Software](#) that allows the rapid recovery of not only failed servers but also failed storage systems. The first key ingredient is that the recovery application should leverage, and actually be powered by the server virtualization infrastructure. This is in stark contrast to traditional backup

and recovery products that merely “support” server virtualization. They can protect the components of the environment just like they can any other platform. But few can actually take advantage of that environment to make the backup, and more importantly the recovery process faster and easier for the administrator.

One Backup, Multiple Recovery Options

With virtualization-powered backup the entire server image is protected every night in a single step. By communicating directly with the virtualization hypervisor only the changes or changed blocks of the server image need to be transmitted across the network to the backup device, saving network bandwidth. The backup can go across a standard IP network or can leverage VMware’s vStorage API for Data Protection to transmit the data directly across a high speed fibre connection.

Prior to a new backup being received a snapshot of the image is taken on the backup server before the next backup so that multiple points in time can be held for the server image. From these backup images full servers or individual files within those images can be recovered. Not only is the second backup sweep required by the traditional ‘bare metal’ backup process eliminated, but the initial backup sweep is significantly smaller thanks to only the changed blocks being transferred across the network. In fact, so efficient is this process that many data centers are executing a “full” server backup more than once per day.

Hardware Abstraction

From a single server standpoint, if a virtual machine fails or a host that supports dozens of virtual machines fails, those virtual machines can be restored directly to another host without any concern about hardware compatibility. In the case of virtualization powered backup it most often is merely a disk-to-disk copy and a restart of the virtual machine.

Recovery in Place

The single biggest value of virtualization-powered backup applications vs. traditional backup applications that merely “support” the virtual infrastructure is their ability to provide unprecedented recovery capabilities and to do so in unprecedented recovery windows. Products like Veeam’s vPower have the ability to recover in place, meaning that the recovery process is not delayed by having to copy the virtual machine image to the new host. The new host can mount the backed up image from the target itself and return service to the users, although in a slightly degraded mode since it’s being hosted from the backup device. However, with that service restored, the backup administrator can leverage VMware’s Storage vMotion to move the image in real-time to the higher speed primary storage, without application interruption. This can include a situation where the primary storage device has failed and the service is restarted from the backup device and then vMotioned to a new primary storage device.

Instant Virtual Sandbox

Testing is the key to recovery confidence. It provides knowledge of how a true full-system recovery will work and allows for the verification of the interdependencies between various servers. Ideally, every data center would have a group of servers on standby so that these tests can be performed, often in an infrastructure called a “sandbox”. The challenge for most data centers is the cost to build these sandboxes is out of reach, especially in today’s economic times. In addition, there is the time it takes to configure and reconfigure the sandbox for whatever is to be tested next.

Virtualization-powered backup should provide the ability to create a virtual sandbox that is preconfigured for the specific application, including server interdependencies.

This means that a recovery group can be made of all the servers and recovery can occur within the sandbox. This virtual sandbox also creates a virtual switch that segments these servers from the production network.

For testing, the virtual sandbox can be 100% automated so that after each backup, the group of application servers can be started, tested and shutdown all in a matter of minutes. What once took hours (or even days) to complete in the physical world can now be done for every backup, every virtual machine.

Once the sandbox is up and running, a process that only takes a few minutes, it can be used for testing. But it can also be used as a way to quickly recover individual items within the application, like email messages. Since the application and its dependent servers can all be started in the virtual sandbox, no additional application specific agents (or cost) are needed. This also removed the burden faced by many application owners that they have to wait for their backup vendor to support the latest version of an application before they can upgrade.

The virtual sandbox can also be used to test new patches or software code updates; meaning changes can occur to it without impacting the normal production environment.

Summary

Virtualization-powered backup stands in stark contrast to backup applications that only support the virtual environment. While those legacy applications merely provide protection of the environment, virtualization-powered products like Veeam’s vPower leverage the environment to provide simpler, less disruptive backups, rapid recovery and the ability to extend the backup environment to provide product development testing.

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>

Copyright © 2010 Storage Switzerland, Inc. - All rights reserved