
VMWARE BACKUP, DISASTER RECOVERY (DR) AND STORAGE
Alternatives to VCB for VMware backup
By George Crump, Founder, Storage Switzerland

Alternatives to VCB for VMware backup

Protecting the VMware environment has its own unique set of data protection challenges that have yet to be fully experienced in data centers. There are three common practices available for protecting VMware: the guest OS method, the console backup method and the VMware Consolidated Backup (VCB) method. While each backup practice has its own advantages and shortcomings, you can compensate for their faults with alternative methods such as block-level incremental (BLI) backups.

VMware Backup Methods

The first practice, and probably the oldest, is the guest OS practice. This practice essentially treats each virtual machine as a standalone server; a software agent is installed on a virtual server, which is then protected as if it were a physical server.

This method of VMware data protection has its advantages. First and foremost, the backup administrator is already familiar with it. It requires essentially no change in process or procedures. Support for database applications, and other open file situations, is well defined through the use of modules from the backup application.

The second practice is the console backup practice, in which virtualization administrators back up the VMware ESX Server with no regard of the underlying virtual machines in the ESX environment.

The final and the newest practice is VCB. The VCB Backup Practice requires both VMware Infrastructure 3 (VI3), and shared storage as the foundation of the architecture, such as Fibre Channel or iSCSI SAN. With VI3 and shared storage in place, you then can add a dedicated Windows Server 2003 to act as the backup proxy. You then install the VCB software on the Windows Server and provide access to the same SAN Logical Unit Number (LUN) used for the VMware Virtual Disk Files. You must also install the VCB agent on the same Windows Server (one must exist) from your backup application software vendor, and then perform the integration steps provided by VMware and the backup application provider.

The backup application launches the backup of a specific virtual machine and then communicates with VMware VirtualCenter and the appropriate ESX Server to prepare the virtual machine for backup. The virtual machine then flushes pending I/O operations, and a redo log is created for each virtual disk to be protected. Once created, the virtual machines can be returned to their normal state. The backup proxy then mounts the virtual machines' disk files and allows the backup agent to send a copy across the network to the backup server, which then writes the stream to the backup target.

VCB Has Its Shortcomings

While VCB sounds like the perfect backup method, many customers do not use it. The biggest obstacle is the requirement for shared storage and need for VMware 3.1. It also still requires scripts in order to execute properly.

VCB is also fairly complex to set up and manage. Many customers do not have the excess space on their shared SAN environment to handle the large storage requirements of the VCB created data set. The backup of the VCB area is still a file-by-file backup that is both time consuming and prone to error. Lastly, VCB does not provide specific support for databases and, as a result, there is concern about getting clean snapshots of the environment. Restores are a two step process; first to the backup proxy, then to the primary storage area.

Because of the drawbacks to VCB, most customers are most likely protecting their systems using the guest OS practice. In other words, they are installing a standard backup agent on the virtual machine (guest OS) and backing it up as if it were a standalone physical machine.

While the major backup software providers have all now delivered VCB modules to manage the backup of the VCB snapshot, few have addressed the guest OS backup issue. This is in spite of the guest OS method being the customer-preferred method for the foreseeable future.

The challenges for guest OS backup and the reason VCB was created are significant. For example, installing a backup agent on each guest OS and then triggering a backup could, on a large VMware server with many guest OSes, consume so much processing power that the guest OSes would crash or the backups would run so slow as a collective unit that they never complete. To get around this, backup administrators need to develop a very complex scheduling process to prevent too many guest OS backups from occurring simultaneously on the same physical ESX server.

A Lighter Alternative To Guest OS Backup Agents

The solution to using guest OS backup agents is to use a backup application that utilizes very little of the backup client's CPU. One such solution is block-level incremental backups. BLI backups communicate with the guest OS at a lower level, so they do not require the complex file system walks of traditional backup agents. Changed blocks can be identified quickly and with less overall CPU impact than looking for changed files. Once the changed blocks are identified, only those changed blocks are transferred and thus the transfer of data to the backup server occurs much more quickly.

The initial BLI backup backs up each virtual server block by block, recreating the virtual server block by block on a standalone backup disk target outside of the virtual environment. Prior to the second and all subsequent backups, a snapshot of the backup disk is taken, preserving file versioning information. After this snapshot is taken, changed blocks are sent only to the backup disk target. For most servers, if not all, this results in a very small backup load and takes less than five minutes to complete. With some backup vendors, the image on the backup target is a live, browsable file system and is immediately usable.

One benefit of this method is that it enables virtual to physical server movement. Since the backup target is a browsable file system, a standalone physical server can iSCSI mount the backup target while the application can be working with this data. This can be valuable not only for testing purposes, but in the instance that an application develops problems while working in a virtual state you have the ability to move quickly to a physical state. Database applications are also often supported with BLI backup software, something that VCB lacks.

Additional Guest OS Backup Challenges

The guest OS backup method does not capture the state of the ESX environment. There are utilities provided by VMware that do this easily, but they require a separate step outside of the backup process to capture and protect this information. You can also use an agent from your backup software provider to protect your environment by using the aforementioned console backup practice.

The console backup practice is the inverse of the guest OS practice. The service console is used to back up the ESX server, as the OS and virtual machines are just "files" under that OS. The service console is a special virtual machine based on Red Hat Linux that runs on each ESX server host. With this method, a Linux client is installed on the VMware service console. The service console is visible to VMFS. In this file system are the virtual disk files, the various virtual machine and ESX configuration files and the redo log files. The agent then sends the backup stream across the network to the backup server, which writes the backup image to the appropriate backup target. Then, the backup software manages this data like it would any other data.

The VCB backup method has its advantages in providing a single backup point that is offloaded from the VMware ESX host, but its impact on resources may be too large of a burden to bear. A hybrid approach of guest OS and console backups may provide a more reasonable solution. The console backup practice,

when used in combination with the guest OS backup practice, needs only to capture the virtual machine and ESX configuration files, which means that the virtual disk files can now be ignored. A guest OS backup practice that leverages BLI reduces the CPU load that causes scheduling challenges and reduces the amount of data that needs to be transferred to the backup server and the backup disk target. That same BLI technology can be used on the Linux VMware Service Console for backup of the ESX environment files, resulting in a complete solution.

About The Author:

George Crump is President and Founder of An IT analyst firm focused on the storage and virtualization segments. With 25 years of experience designing storage solutions for data centers across the US, he has seen the birth of such technologies as RAID, NAS and SAN. Prior to founding Storage Switzerland, he was CTO at one the nation's largest storage integrators where he was in charge of technology testing, integration and product selection.