

## Backup Strategy – Continuous Backup Needed?

Posted By [Texiwill](#) on February 11, 2010

I had an interesting conversation with Vizioncore yesterday about how backup is not as much a decision about what software to use but what process to use. In addition, this process needs to be considered and thought about from the very beginning of your virtualization architecture design process. With the quantity of virtual machines being used today by the SMB and Enterprise customers, the backup window has grown to nearly an all day event. What you say? An all day event! My backups happen with the window I set.

However, if your virtual machine quantity is sufficiently high you may indeed take all day to make a backup of all your VMs. VM sprawl makes this worse, but that is another subject. Virtualization Backups happen using three basic technologies, but each have their own issues with respect to the architecture and design to be created for your virtual environment. These technologies are:

1. Backup from Within the VM using traditional backup agents and tools
2. Backup from Without the VM using network protocols to transfer the entire disk image to some other location
3. Backup from Without the VM using storage protocols to transfer the entire disk image to some other location

The last two sound like the same thing as storage protocols are often network protocols, but storage protocols are those implemented within the storage devices to mirror LUNs and transfer data within the storage device(s) instead of using a network cable connected to the virtualization host.

There are several factors that affect your backup strategy and therefore the design and architecture of your backup process.

- Backup is disk IO intensive. You should have a IO requirement for any tool you choose to use.
- Backup can be network intensive. You may need more network that you currently have within your virtual environment
- Not all storage devices support protocols for LUN backup and mirroring within the hardware. You may need to spend more money to achieve this level of backup
- Not everything can be backed up without first handling data integrity issues such as required for databases. You may need to investigate the requirements of your in use Applications.
- Off-site storage of backups maybe required, is this via Tape, disk, etc.
- Existing backup tools (tape libraries, software, etc.) may already be in use within the organization.

Given all these possible issues, it is wise to start your backup design and architecture when you start your virtual environment design and architecture. Do not get tied to anyone product but out of your backup design and architecture will come requirements for such software. You may find that your existing software may be sufficient or may need to be augmented to perform backups. Unfortunately, choosing new software often contains a political component. Let the requirements speak for themselves.

One of the major issues to consider when you are looking at architecting a virtual environment backup solution is how much data will need to be backed up for each of the methods listed above. You may find you need to do continuous backups. An example of this would be the time it takes to initially backup several TBs of data. Can your existing network and storage handle the IO requirements or do you need to switch to a more continuous backup of sustained bandwidth requirements to get your backup load completed.

Continuous backup is not continuous data protection, but more about sustaining backup throughput throughout the day instead of during a single backup window. If it takes 2 hours to backup a single VM how long would it take to backup 20 VMs? This generally exceeds any backup window that exists these days as systems are required to be on-line 24/7. So now we need to consider minimally the following:

- when to backup each VM,
- whether to do incremental backups,
- how many incrementals before we do another full backup,
- how to perform data integrity work prior to each backup (VSS is one solution perhaps)
- how to keep 'Applications' running while data is being backed up
- etc.

The list is fairly endless, and differs from organization to organization. Backup may require serious changes to hardware (introduction of more networking for example), may require changes to storage, and definitely will require forethought before implementation. Like everything else with respect to the virtual environment, Backup is not something you should leave to the last minute. Ensure it is part of your initial architecture and design.

### **About the author**

Edward L. Haletky is the author of [VMware vSphere\(TM\) and Virtual Infrastructure Security: Securing the Virtual Environment](#) as well as [VMWare ESX Server in the Enterprise: Planning and Securing Virtualization Servers](#). Edward owns [AstroArch Consulting, Inc.](#), providing virtualization, security, network consulting and development. Edward is also a guru and moderator for the VMware Communities Forums, providing answers to security and configuration questions, prolific blogger, and is working on new books on Virtualization.