

## Chapter 19: Virtual Machine Backup

### What's New

- **VMware Data Recovery Appliance (vDR)** – vDR is a virtual machine that facilitates backup and comes with vSphere Client plug-in. Works with ESX “Classic” and ESXi to backup VMs – which allows for the backup of the VMs files. Support for advanced features such as compression of de-duplicated data. No network hit when used with shared storage. In the beta release there is no support for backing up individual files within the VM – and this feature will be only “experimentally” supported at the time of the vDR’s general availability. However, after the first backup – all subsequent backups are merely the delta changes within the virtual disk – using special change block tracking functionality available only in VMs with Hardware Level 7.

### How Do Hot Backups Work?

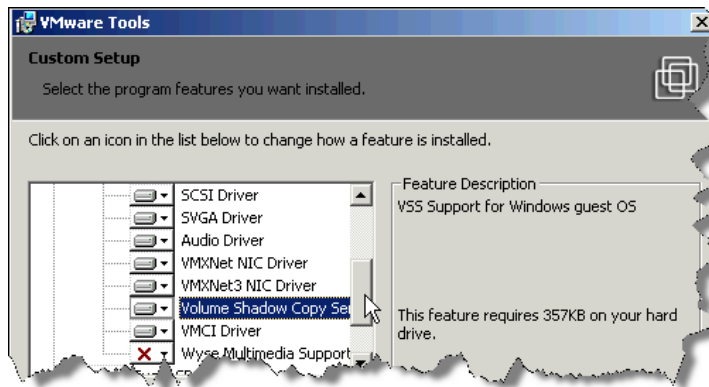
It is possible to backup the files that make up a virtual machine while it is powered on and running – in other words a hot backup. This is achieved by either VMware or the third party backup vendor leveraging VMware’s Snapshot feature. In fact with a few simple commands (vmware-cmd and vmkfstools or VMware’s Power CLI) you could actually cobble together your own “home-brewed” backup solution. I wouldn’t recommend that you do – because others have got there before you and have been doing it longer. As you might know, when a snapshot is applied to a VM the files that make up the VM become unlocked in the file system. This “locking” process not unlike the way an application like Microsoft Word “locks” the files you are using – and I imagine more than once in life you have had the experience of trying delete files that are “in use”. The locking process is intended to protect you basic operator errors such as deleting a powered on VM.

However, when a VM has snapshot – this behavior is modified such that the file become unlocked. When a file is unlocked the possibility of backup becomes available. It’s important to know that ALL vendors use VMware Snapshots in this way. The quality of the backup vendor’s software comes in the validation process. Validating that the snapshot was successfully created prior to the backup starting, and critically that the snapshot is deleted if or when the backup completes. If you recall the section about snapshots you might remember that one of the most common reason for VMware Snapshots to become excessively large is bad backup programs that do not successful validate the snapshot process or alert the administrator when something goes wrong.

By now you might be pretty worried. After that entire paragraph above wasn’t especially reassuring was it? Well, buddy welcome to the real world! What I want to do is make you aware of the realities of hot backups of VMs because forewarned is forearmed. This is not to say that they are not good thing, but to remind you that bad things can happen and you should be looking out for them.

Firstly, let’s take the quality of VMware Snapshots in the context of backup. VMware’s Snapshots in Vi3.5 introduced a File System Sync driver installed to the VM during the install of VMware Tools. The job of this File System Sync driver is to flush the file system cache, forcing a write-to-disk prior to the backup process. The intention is that you get the most complete backup you can have with live data still being memory – that is to say a consistent backup, rather than a crash-consistent backup.

This File System Sync driver from VMware is pretty good, but in the early days of Vi3 it was found by customers to be limited when it came to specific systems that store their data in modes beyond just being ordinary text or graphics files. If you want some specific examples the File System Sync driver was found to be somewhat lacking when it came to backing up systems such as Active Directory, Microsoft SQL, Oracle and so on. Essentially, the driver lacked the level of file system awareness required to backup these candidates reliably. The response to this was implementing call outs or hooks into Microsoft's Volume Shadow-Copy Service. This came into the Vi3.5 product at around the Update 1 or 2 period, and it has been maintained within vSphere4. Interestingly, VMware appeared to be quite slow in taking up Microsoft VSS, and there were a number of third-party backup vendors such as VizionCore who were offering this level integration with Microsoft much earlier.



The second issue associated with VMware Snapshots are the reasons why the sometimes the snapshot is left behind after the backup process completes. As you recall this could be dangerous if left unchecked as the VMware Snapshot has its heart a "delta" file which grows incremental in blocks of 16MB. One reason why VMware Snapshots can be left behind after the backup has completed is simply sloppy coding. I've come across a number of home-brewed "solutions" where because of poor scripting the VM's snapshot has not been deleted at the end of process. I've often been asked why this happens – and unfortunately the answer of "poor scripting" is not one that goes down especially well, I can tell you! However, putting this issue aside the main reason can be the fragility of your network. Most backup systems either need communication to the vCenter to function or to ESX. It is quite possible that the backup system might trigger the backup process without error, but mid-way through the backup job loose communication to either the ESX host or vCenter. This could be due to something VMware responsible for such as a failure of the vCenter service or something as simple as some idiot in your environment creating a IP conflict. It could be both. The net result of this lost communications is backup which is successful, but fails to delete the snapshot at the end of the process. Personally, I think we can accept that no network is perfect and whilst we strive to protect ourselves from mass network outages by teaming devices – when it comes to software or configuration it becomes increasingly harder to offer the same level of redundancy. If we accept this as reality – then really the benchmark then becomes how well the given backup solution then detects, repairs and logs this event and how well it alerts the administrator to a potential problem.

## Backup Strategies

There are a million ways of getting copying files to be the same in two different storage locations – which is as brief a description of backup as think of – such that if the file in one location is lost, damaged or deleted – the original data can be retrieved. As a reader you will be grateful that I will absolutely bore you with lengthy descriptions of normal, differential, incremental or synthetic backups or extol the merits of disk-to-disk de-duplication of data. In my current role as an instructor I'm often asked for a survey of backup options or strategies together with overview of different methods. These can be as crude or sophisticated as you like. As individual my backup concerns are very modest – I backup my laptop data – and I “backup” the 9 critical virtual machines I have that make my life workable. None of these 9 VMs are important to me – they could be rebuilt from scratch without any data loss. The only reason I backup them is because I know what a royal pain in the butt it would be to that rebuild. You might be curious to know how I backup such a modest environment. Well, once a week, quarter or year (depending on how my mood takes me) I power them down, and clone them as templates in the compressed format to alternative removable drive. In my mind a copy of a VM held in different location is a backup by any other name. I do recognize that the field of backup is immensely more complicated than what I have just described, and I am not by any means promoting or endorsing such an approach – but it gives you an idea that if you are prepared to think – anything is possible.

The Holy Grail of backup in the world of virtualization is this, to be able to backup a VM directly from its storage location (SAN, iSCSI, NAS) without having to install agents into the guest operating system. This backup should offer all the features of conventional backups (normal, differential, incremental). The backup should be quick, and even more importantly so should the restore process. It should be simple and easy to use – and not take a degree in scripting languages to setup or use. You might be interested to know that I personally think that the virtualization sector has taken some time to mature – in such a desperately important part of IT. It's for this reason that installing backup agent inside a VM is still very popular because often the virtualization eco-system has failed to deliver a solution which has approached anything near universal adoption by the VMware Community. So put simply there is no one way of backing up, and when that happens what you will see is many competing methods which all come with advantages and disadvantages. Of course one method of getting the best of both all strategies – is to use more than one method. However, even this strategy has the disadvantage of significantly increasing the cost of storing that backup data in multiple locations. Essentially these competing strategies can be broken down into 4 approaches:

- Install Backup Agent inside the VM
- Use “Free” backup utilities
- Buy a third-party “virtualization-only” solution
- Use VMware Technology – The Data Recovery Appliance and VMware Consolidated Backup

### Strategy 1: Backup Agents inside the VM

It's entirely possible to backup individual files (.doc, .xls and so on) inside the virtual disks but it still remains popular to install backup agents inside the VM and backup up across the network. There are some distinct advantages such as:

- **Mixed Environments** - VMs and physical machines are backup in the same consistent way

- **Very mature** – this backup method will reliably backup troublesome candidates like Microsoft AD, SQL, Exchange, and Oracle making sure your data is restored in a consistent state
- **Easy Adoption** - No need to re-rehearse your restore strategy or educate staff about new procedures and tasks
- **Improved by virtualization** – If you lost an entire VM you can very quickly deploy a new VM containing a backup agent pre-installed – and start the restore process very easily

However, using backup agents is not without disadvantages either

- **Cost** – Most backup vendors charge a license fee for each backup agent. Some don't like HP DataProtector, and some vendors like Symantec have shifted the licensing policy to cancel out this disadvantage
- **Throttled by Network** – By definition with an agent installed inside the VM, you will see that backup and restore will be constrained by the bandwidth from the backup storage location
- **Performance Hit** – During backup there will increased I/O on the ESX host – to absorb the backup load
- **Encapsulation; Encapsulation; Encapsulation** – One of the major benefits of virtualization is if you virtual disks, everything is already “just a file” anyway. In away, using backup agents represents a “if it ain't broke don't fix it” approach IT which some would say significant reduces one the advantages of virtualization. When you use a backup agent inside a VM you affectively treating the VM as if it was physical. Some would reply – What's wrong with that?

### Strategy 2: Free Virtualization Backup Utilities

In the early days of virtualization (ESX1 and ESX2) there really was no backup solution beyond installing backup agents inside the VM. However, as stated before as the VM is really just a bunch of files, there were command-line utilities that would allow you to snapshot a VM, and then export the VMs virtual disk to another location. Essentially, this represented a normal backup every time – and so the storage penalties were quite high. They were also significantly dependent on the ESX Service Console as a scripting environment – and so would be unsuitable in a ESXi environment. Over time certain individuals developed these scripts in such a way that they almost became products – with the critical difference being that they were still free and only support by the community. You might be surprised to know – that despite the increased regulatory burden we live under where corporate compliance is king – that these methods were very popular in the VMware Community. This was mainly because there were very good, and there was no alternative.

Two free backup utilities are currently worth mentioning – but whether they will be continued to be improved and supported in vSphere4 is that this stage unclear – they can be currently found on these locations. One free utility have passed into private hands, and is longer freely available. This leaves the grand daddy of them all – the vmbk.pl script available here:

<http://www.vmts.net>

### Strategy 3: Third-Party "Virtualization-only" Solutions

As you can see with backup story I provided earlier, there is plenty of scope for third parties to develop their own unique backup solution. There are basically three players currently in this virtualization-only space they are:

- PhD Technologies – esXpress product
- Veeam Backup (which also includes a replication engine included)
- VizionCore - vRanger Product (owned by Quest Software)

I'm probably not saying anything contentious by saying that VizionCore's vRanger (formerly esxRanger) is probably the most popular product in terms of units sold – and the product that has been around for some time. Although popularity and longevity are solid attributes when compared with the likes of Symantec or Legato – these organizations are still very much the new kids on the block. As with installing backup agent inside the VM, these vendors come with their own advantages including:

- **Delivering the Holy Grail** – all these vendors claim to deliver the Holy Grail I described at the beginning – of an agent-less backup of the VMs. Whether they do deliver will largely depend on their quality of coding
- **Unique Features** – The often deliver so called bleeding-edge features much quicker than VMware. For example these vendors delivered a delta backup of the entire VM (.nvram, vmx, and so on) well before VMware's Consolidate Backup. With these delta backups – you backup all the files of the VM just once as with a normal copy, and then from that point onwards you are merely backing up the blocks that change of those files
- **Simple to Use** – Very often these vendors have delivered very easy to use point-and-click management tools which make virtualization backup and restore incredibly easy – and accessible to even the smallest of SMBs with own modest data recovery skills. They also come with powerful scripting engines for the more able administrator to automate the backup process.

Likewise there are disadvantages as well:

- **IT Politics** – The decision to introduce a brand new propriety backup solution in addition to an existing backup solution that might have enterprise license agreement decided in different continent is often a reason why these third party vendors don't get a look in. Basically, the politics of data recovery is a prickly one, and may be beyond your control
- **Quality Regional Support** – Some third party backup vendors may have quite modest support teams – they may not operate in multiple time zones and in multiple languages
- **The Start-Up Mentality** – There has been stories (admittedly in the early days) that third party backup vendors issuing customers with new versions and bug fixes on a daily basis. This can be very annoying, irritating and undermine the confidence in the technology. You could argue that that very large ISV fair no better – software is software after all
- **Cost** – It would be silly to over-exaggerate the cost of these third-party vendors – in fact they are extremely competitive products for what they deliver. But some might see acquiring them as yet another additional cost to the business. After all you already have a backup solution – why can't you just use that?

- **Vendor Run Around** – By definition all of these vendors are to some degree dependent on the APIs and management platforms of VMware. So if they don't work reliably – who is to blame – the third-party vendor or VMware? Critically, when you raise an SR – who is responsible for fixing that issue? There's nothing new about this vendors run around – and some would say this is reality we just live within the world of IT support

#### Strategy 4: Use VMware Technology

As time has gone by VMware have developed their own backup APIs to which the conventional third-party vendors can “hook” in to. VMware's strategy in the past was something called VMware Consolidate Backup. Oddly enough VCB wasn't actually a backup solution per se but rather a collection of command-line tools, scripts and driver that would allow an existing backup vendor to use their backup solution for Windows to access the VMFS volumes and the files within for backup purposes. New to vSphere4 is the Data Recovery Appliance – which is downloadable virtual appliances which assists in the backup process. Of course even VMware Technologies come with advantages, and disadvantages

- **Reduced Vendor Run Around** – I say reduced rather removed, as if you are using VCB you will find the vendors have merely switched places – from being a third-party “virtualization-only” vendor, to a more agnostic traditional backup vendor
- **Improve Backup Performance** – Performance is always point of argument between vendors – and one which I would like to side step here. But with VCB the backup load is removed from the network and the ESX hosts the means of dedicated physical server with connection to your SAN, iSCSI or NAS based systems.
- **Leverage Existing Backup Vendor** – With VCB there's no additional purchase or education cost beyond becoming familiar with the VCB Framework. You can carry on using the same tools for backing up physical machine and virtual machines.

#### What is the VMware Data Recovery Appliance (vDR)?

The VMware Data Recovery appliance is a virtual machine, in other words a virtual appliance that assists in backup process of virtual machines. It backup up the entire virtual machine files (.nvram, .vmx, .log, .vmdk including RDMs in a virtual compatibility mode). Experimentally, it can also backup files within the VM such as Word or Excel documents. In the early beta the vDR utilized a “helper” VM which allows SCSI based hot add backups which are more efficient than network based backups. This functionality was later integrated into the vDR. You have the choice of backing up your VMs to virtual disk or experimentally RDMs. It also supports the mounting of Windows Shares (CIFS/SMB) for an over-the-network back up process. Currently, the vDR is not backwards compatible with Vi3.5, and you will require vCenter 4.x to running. As with VMware HA you will need vCenter and the vDR appliance to be running to for backups to occur, but you do not need vCenter to carry out a restore process. If something goes wrong with the vDR you merely re-import the .ovf file, and point it to your backup destination. VMware intends to re-release the appliance rather than issue patch up dates. The vDR requires the new “quiesced snapshot” option, and therefore only VMs that have been upgraded or created using Hardware Level 7 qualify as valid target VMs. As with VMware Consolidated Backup, the vDR does not backup up end-user created snapshots. Theoretically, there is no reason why this isn't possible – but appears VMware wish to avoid the complexity of having to backup snapshots when the very method of backup uses snapshots.

VMware recommend that the vDR appliance runs on an ESX host. Critically, it must have less than 80% CPU activity – and there must be 5GB free for each VM you backup. The maximum number of VMs you can backup is 500 with a single vDR appliance – however in practice there is no limitation on the number of vDR appliances you run simultaneously. As for concurrency backups VMware support the maximum of 8 backups and 8 restores as the same time.

To reduce disk storage cost of backing up VMs, VMware have developed their own vendor independent data de-duplication. De-duplication is a process that makes sure that you only backup genuinely unique data. In conventional backup, even with normal, differential, and incremental strategies, many organizations backing the same data more than once – because it is duplicated in many locations. Simple examples of this might be such documents as corporate logos that appear on multiple internal and external web-servers. De-duplication in the backup and archive arena has been developed in such a way that the block-by-block comparison can be made of data – so that backup seems can spot these duplicates regardless of storage location, file name or data stamps. VMware's de-duplication process operates by analyzing the VM to be backed up and breaking it up into smaller variable blocks size chunks which are anywhere between the range of 2KB to 64KB. This is use of a variable blocks size is important for performance. Each of these variable blocks and SHA1 hash is generated and referenced to a block inside the virtual disk. This process eventually generates a list of SHA1 hashes which collectively reference the entire data file. It's by analyzing these SHA1 values that vDR can recognize if a block has been backed up before and if it has been modified. This de-duplication process is free and has been developed in house by VMware engineers. The de-duplication process cannot be disabled (because the disk costs would make vDR unfeasible), the data is also encrypted to prevent malicious interception of the data. Whilst it is possible to use your storage vendors de-duplication process in addition to the VMware de-duplication the additional workload of de-duplicating data, which has already been de-duplication is in most case not going to improve the utilization of the storage. If you forgive the pun – it would be an unnecessary duplication of the de-duplication process! What does pay dividends is backing up the *same type* of VM to the same destination – because by definition this increases the number of blocks that overlap – which increases the benefits of the de-duplication process. The last thing you should know is that de-duplication data is stored as 1GB files in the VMwareDataRecovery folder in the virtual disk.

All these features make vDR very attractive – however, there a limitations. Critically, I think the adoption of vDR will be in the early days hampered by the mere “experimental” support for the backup of individual files contained inside the VM. A similar issue has inhibited the wide spread adoption of VCB in the Vi3.5 release. Now, don't me wrong, VCB can backup individual files inside the virtual disk, and did from its first release. What hampered VCB was its complexity of setup, and the lack of a really efficient method of restoring individual files. As such a system that lacks full support for the backup of individual files (not to mention what the restore process might be like!) is likely to have similar affected. In fact simply stating that there only exists “experimental” support for backup of individual files inside the VM - will be enough for some organization to discount vDR for the time being until it has had the opportunity to mature.

Another weakness of both VCB and vDR is their inability of either technology to backup the end-user generated snapshots. In my mind these files are as important to me as any other file that makes up the VM. By definition we all looking for the most complete of backups possible – by definition if neither VCB or vDR backup the snapshot files – for many this will be a bitter pill to swallow. What's

somewhat annoying about both the file level backup and snapshot limitations – is that they appear to stem from a lack of time and QA resources for them to be included within the time frame of the vSphere4 release. No doubt this position will change as we progress to perhaps a vSphere4.5 release.

The final and last limitation is that whilst the vDR can utilize tape you will need a 3<sup>rd</sup> party solution for this. Whilst many businesses have moved over to disk-to-disk backups for performance and reliability, for large archive sets tape is still important for lots of organizations – especially those who must retain that data because of audit or compliance reasons.

My personal belief is that vDR will be *immensely* popular in the SMB market where virtualization is new, and backup requirements a relatively modest. It provides a very point-and-click interface that even someone with modest IT skills would find easy to use.

### Using the VMware Data Recovery Appliance – vDR (Beta 2)

As stated before the vDR ships as virtual appliance in the .OVF file and vSphere Client plug-in. I've documented this process a number of times, so I don't intend to waste valuable time with a step-by-step guide to this process. The vDR is built around the same APIs within vStorage that allows VMware Consolidated Backup to function. VMware do not intend (at this stage) that the vDR will replace VCB – the two will co-exist. I've chosen to start with the vDR because I think it's significantly easier to get up and running with it, and because it is new. And I like new, don't you? Once you have successfully imported the vDR, you can *optionally* add a second virtual disk which will act as your destination drive for backups. It is possible to connect the vDR to CIFS (Windows) shares for backup across the network.


### Configure the vDR and Install the VMware Data Recovery

1. **Right-Click the vDR** and choose **Edit Settings**
2. Click the **Add** button
3. Choose **Hard Disk** from the device list
4. Choose **Create new virtual disk**
5. **Specify a size and datastore location**
6. **Power on the vDR** for the first time.

**Note:**

Use the Networking wizard to set a static IP address.

7. **Install the vDR plug-in**, and **confirm it is installed and enabled in the Plug-in Manager** window



This should produce a “VMware Data Recovery” icon in the **Home >> Solutions and Applications >>** location:

[chapter19-the-vmware-data-recovery-plug-in-icon.jpg]

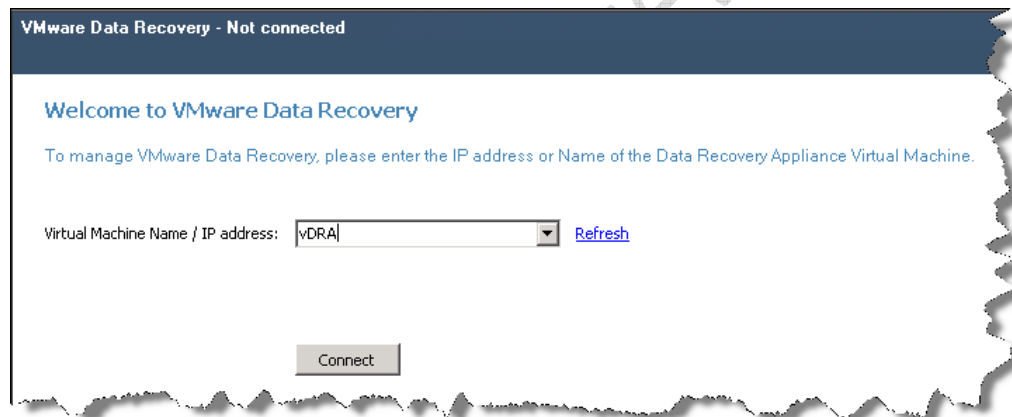


### Connecting vDR to vCenter


The next step is to couple the VDR and vCenter systems together. This will allow the vDR to enumerate the vCenter environment and allow you to select VMs for backup

1. Navigate to **Home >> Solutions and Applications** and click the **VMware Data Recovery** icon
2. **In the edit box type in the VM name of the vDR or its IP address**

[chapter19-settings-the-vDR-name.jpg]

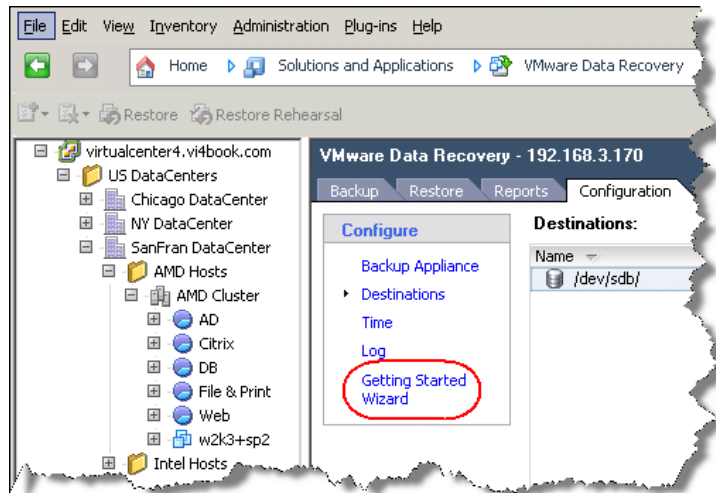


3. Click Connect, and after a short while complete the login dialog box



#### Note:

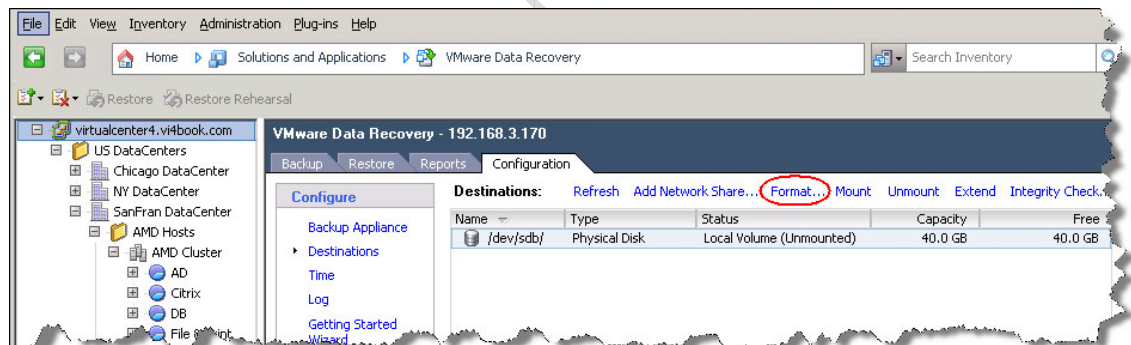
After completing the dialog box a “Getting Started Wizard” will run to guide you through the process of preparing your backup destinations. If you cancel this wizard by mistake (as I did), it can be triggered again from the “Getting Started Wizard” link available in the Configuration tab



### Prepare vDR Format and Mount a vDR Backup Disk

Before powering on the vDR I added a virtual disk, this disk now needs formatting and mounting. This can be done from the “Getting Started Wizard” if you wish. I prefer to do it from the main options in the vCenter window


1. Navigate to **Home >> Solutions and Applications** and click the **VMware Data Recovery** icon
2. Select the **Configuration** Tab
3. **Select the virtual disk** that needs to be formatted, and click the **Format...** option



**Note:**

In the screen grab above you can see I have one backup destination drive (/dev/sdb) which is not mounted (Local Volume (Unmounted)) it is 40GB in size and is blank because it has 40GB Free.

4. Confirm that you are happy for this format to be carried out




**Note:**

At the end of the process the disk will be formatted and mounted for you. You can mount and unmount volumes as will using the Configuration tabs mount and unmount links. You may need to use these options, if you want to add additionally storage to the vDR after its initial configuration.

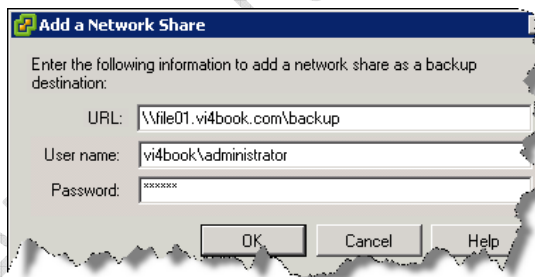
**Adding a Network Backup Location to vDR**

In addition to using an internal mounted virtual disk, you can also mount network shares using Microsoft Windows sharing, commonly referred to SMB/CIFS shares.

1. In the Data Recovery **Configuration** tab
2. Click **Add Network Share**


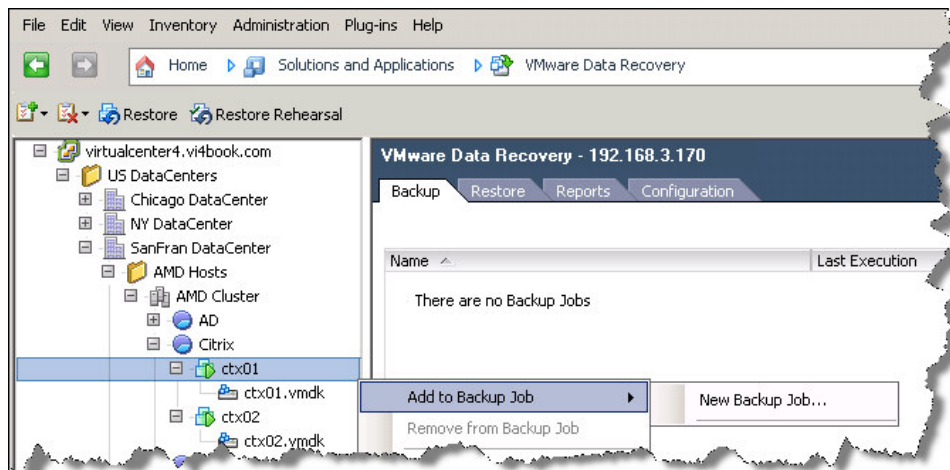


3. In the **Add a Network Share dialog box**, supply the **UNC Path, Username and Password** settings

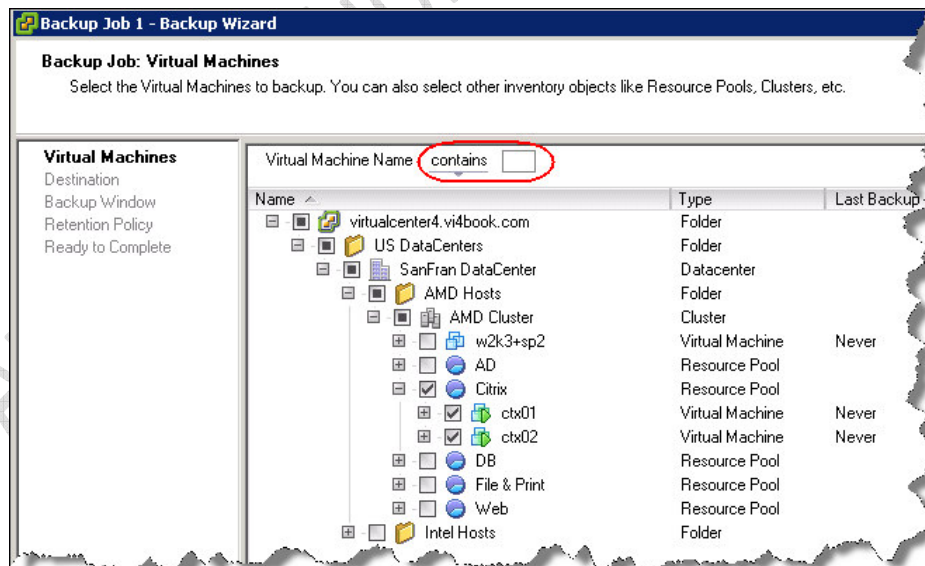


**Backing Up a Virtual Machine**

There is two ways to backup a virtual machine. You can either right-click the VM in the Data Recovery inventory, or else create a backup job from the Backup Tab which allows you to multiple select VMs in the tree.



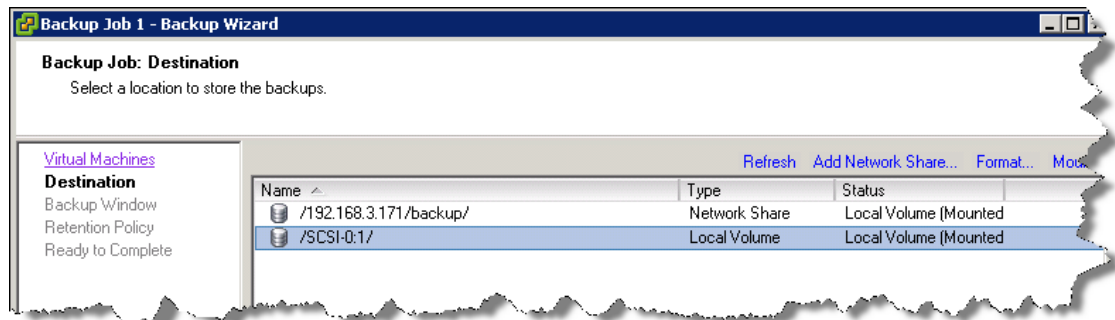
1. Click the **Backup** Tab, and select the **New** option
2. **Select the VMs you wish to backup** from the inventory



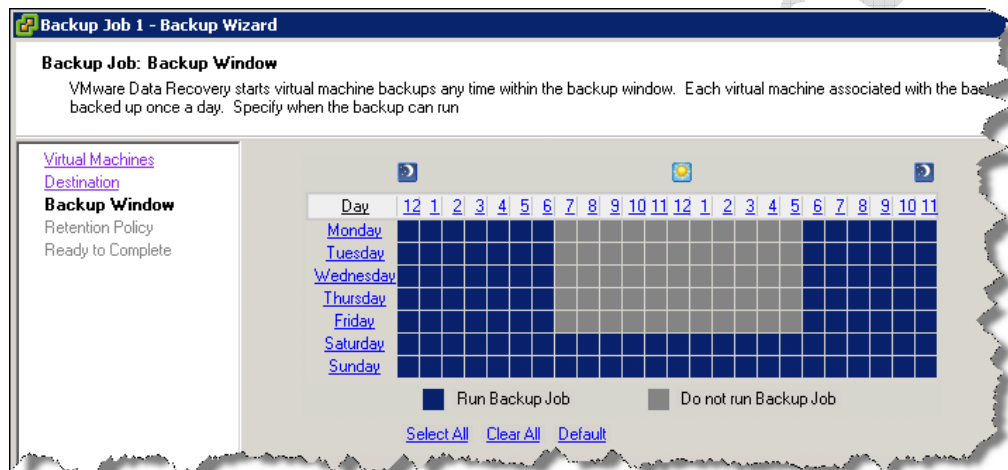
**Note:**

Here I select ctX01 and ctX02 by selecting the Citrix Resource Pool. Note how you can search for VMs by string using the “contains:” option.

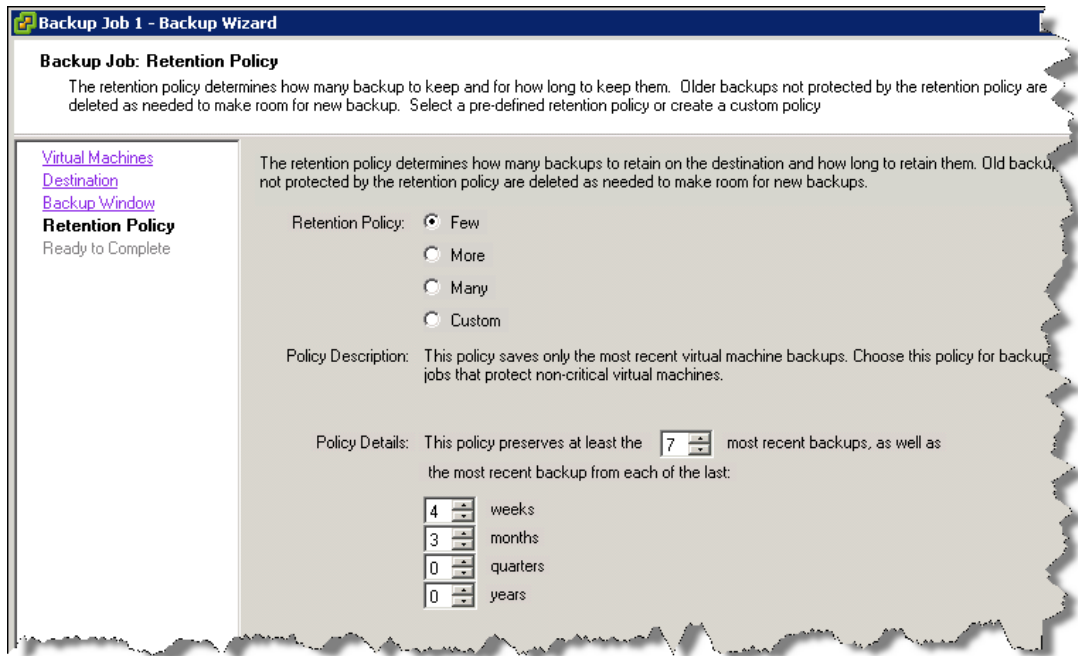
- Next select your backup destination from the list



- Select a Backup Window which controls when backup job is schedule to run. The default is out of standard business hours, once each day



- Next set your Retention Policy

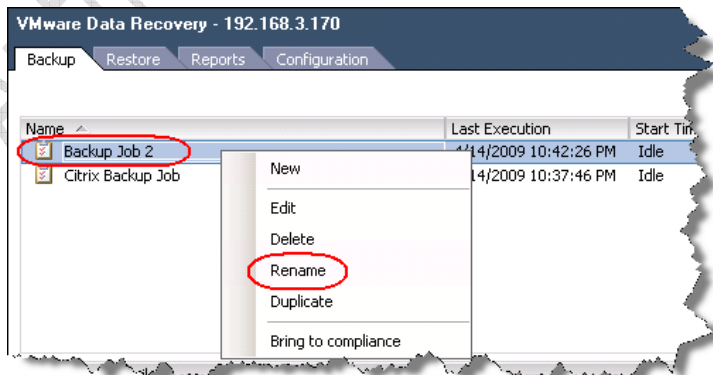


**Note:**

As you might expect the retention policy controls how long your previous backups are kept. The longer the retention backup – the more backups you will have – and the more options you will have to restore the VM. Despite the power of de-duplication these multiple backups are not free from a storage perspective. So as ever your decisions is a compromise of the power of the recovery against the “wasted” space of the backup.

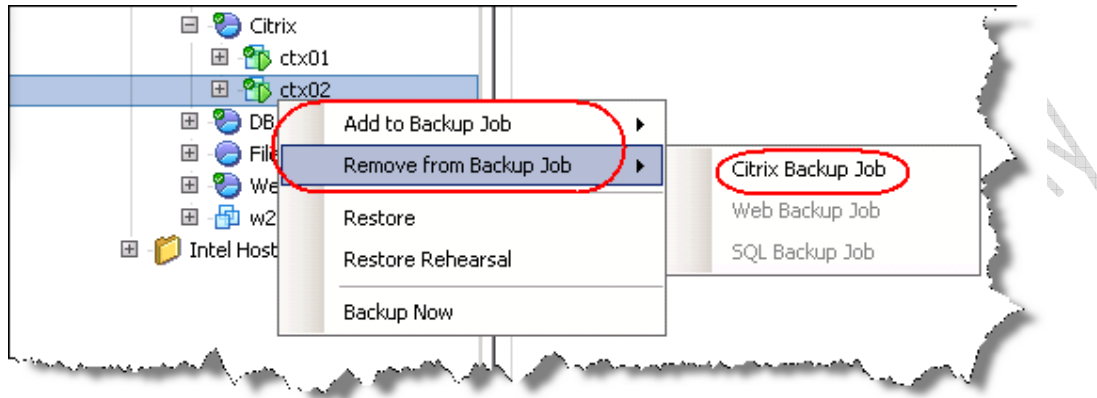
Three built-in Retention Policies exist (Few, More and Many) together with a custom policy – each progressive retain the backups for a longer and longer period. The custom retention policy allows you additional state how many backups from each period.

After finishing the wizard a backup job will be created. In this release you cannot name your backup jobs when you create them in the wizard but you may rename them after they have been created.



### Adding/Removing a VM to existing Backup Job

If your backup job was created using a VM container such as DataCenter, Cluster or Resource Pool the mere act of creating a VM in that container should automatically add it to the existing backup job. If this does not happen you can manually add a VM to an existing job or alternatively, if a VM has been added to a backup job accidentally – then you can remove it from a backup job with a simple right click of the virtual machine.




### Backing up Individual Files within a Virtual Disk

The beta release of vDR does not support backup of individual files contained in the VM.

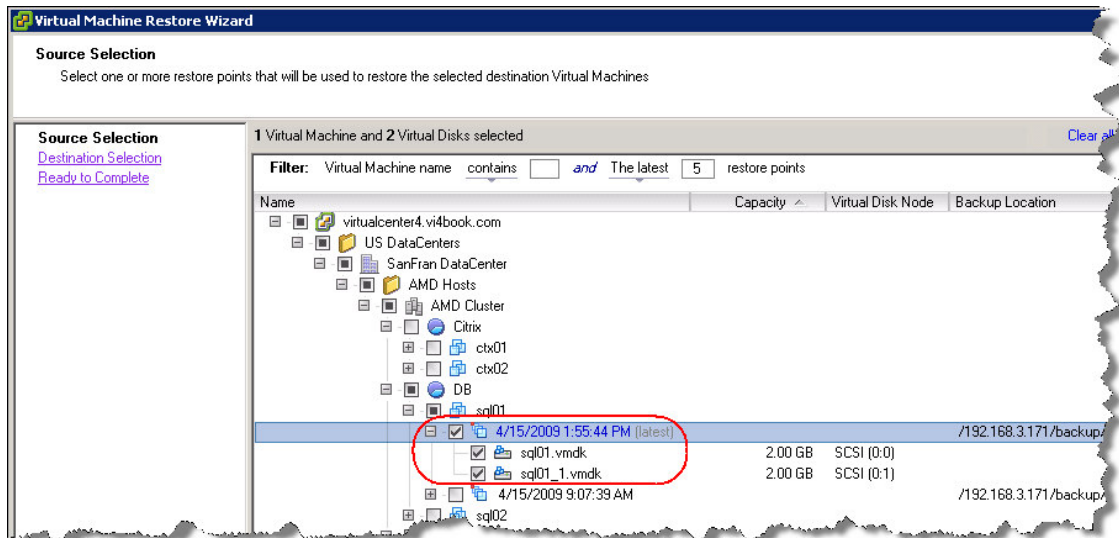
### Restoring a Virtual Machine to Alternative Location

Clearly at some stage you will want to carry out a test restore of a VM you have backed up to alternative location, as a way of validating the backup process. Prior to attempting a restore I created a Resource Pool called "Backup" and VM Folder called "Backup" as well. These will serve as the alternative inventory locations for backup. I will use local storage as the restore destination as I am running a little low on Fibre-Channel SAN space.

1. Select the **Restore** tab, and click the **Restore** icon



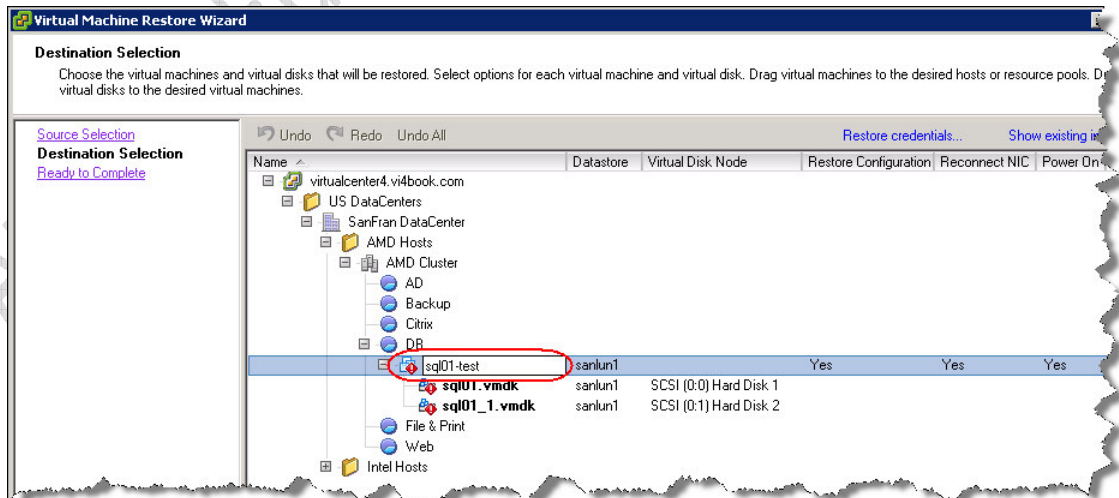
2. In the **Restore Wizard Window** select the VM(s) you wish to restore



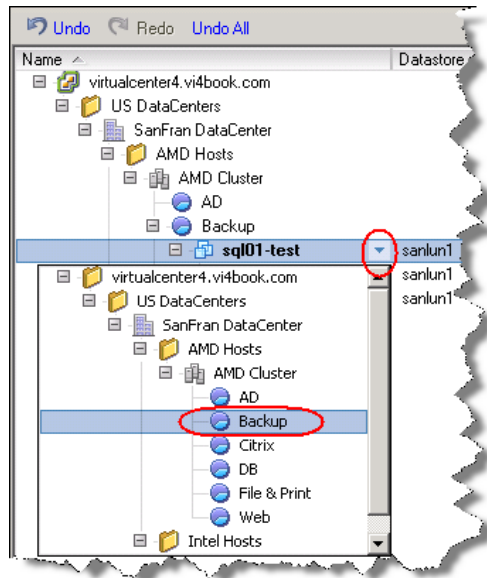
**Note:**

In this case I have chosen to restore both disks that make up the VM called SQL01, it possible to merely select individual disks in this process. You can see that this is the latest backup taken on the 4/15/2009 compared to the older backup taken on the 4/15/2009. Notice also that the UI defaults to showing the last 5 restore points – whether these appear depend greatly on the frequency of the backup and your retention policy.

3. **Next we modify the restored VM so it does not conflict with the existing VM called SQL01. To do this you must first rename the restore VM.** This can be done back clicking into edit box where the name of the VM is displayed. **A red alert icon on the VM indicated that if you preceded with the restore you would overwrite the original VM.** Only by renaming, and relocating it to can you avoid this from happening.

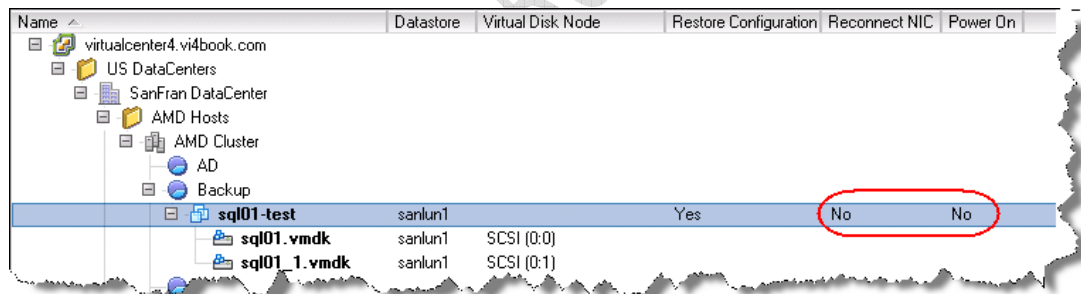


Once you have renamed the VM to be restored it is possible to **relocate the VM to alternative resource pool** like so:



I would also recommend setting **Reconnect NIC** and **Power On** to be set to **No**. This will allow you to make sure you restored backup does not generate any IP conflicts or Microsoft “Duplicate name exists on this network” errors.

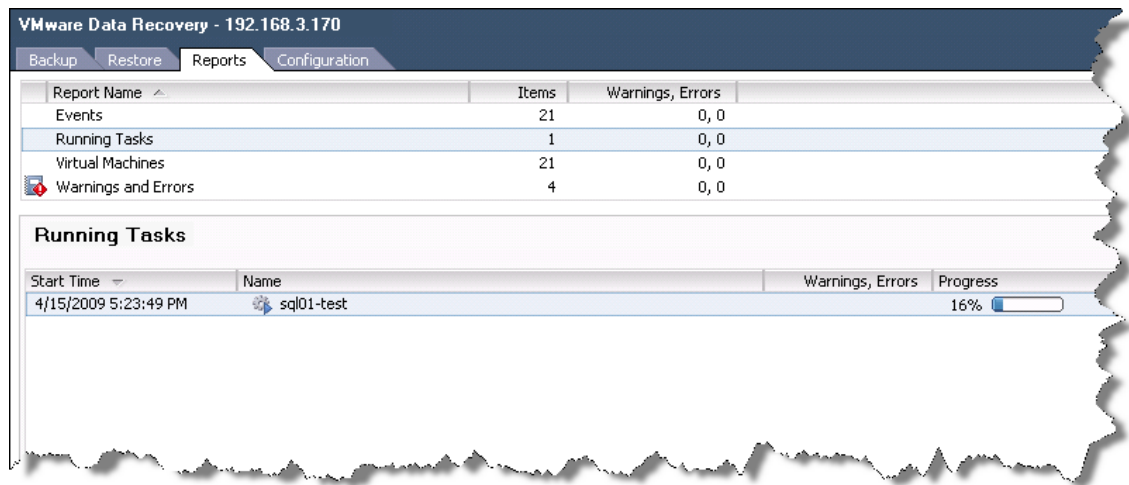
[chapter19-disconnect-network-and-dont-power-on.jpg]



- When you are finished, click **Next** and then **Restore** to trigger the restore process.

**Note:**

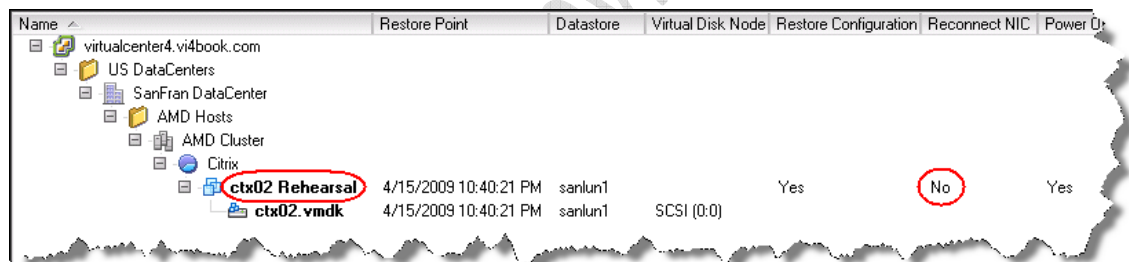
After click the button you will be shelled out the wizard into the Reports Tab, where you may track the progress of the restore



### Restore Rehearsal

A simple way to test your restore process is using the restore rehearsal feature. This reduces the amount of questions asked compared to a manual restore which was covered just a moment ago.

1. Right-click the VM
2. Choose **Restore Rehearsal**



#### Note:

Notice how with the quick restore rehearsal the wizard automatically renames the VM as vm-name Rehearsal, additionally the Reconnect NIC option is automatically set to NO. This should prevent the administrator overwriting the existing VM accidentally or creating IP and NetBIOS name conflicts.

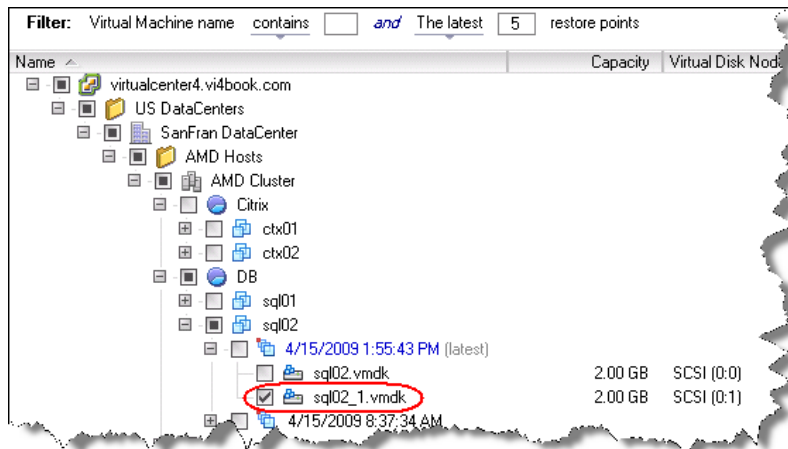
### Restoring an Individual Virtual Disk

It is possible to restore an individual virtual disk this is possible, and the process is very similar to the procedure outline above. For example you could restore an individual virtual disk, if you have lost a significant amount of data in a virtual disk, but your guest operating system is in a fair shape. Conversely, perhaps your data is fine, but your operating system is damaged beyond simple repair you could restore just the boot or system disk.

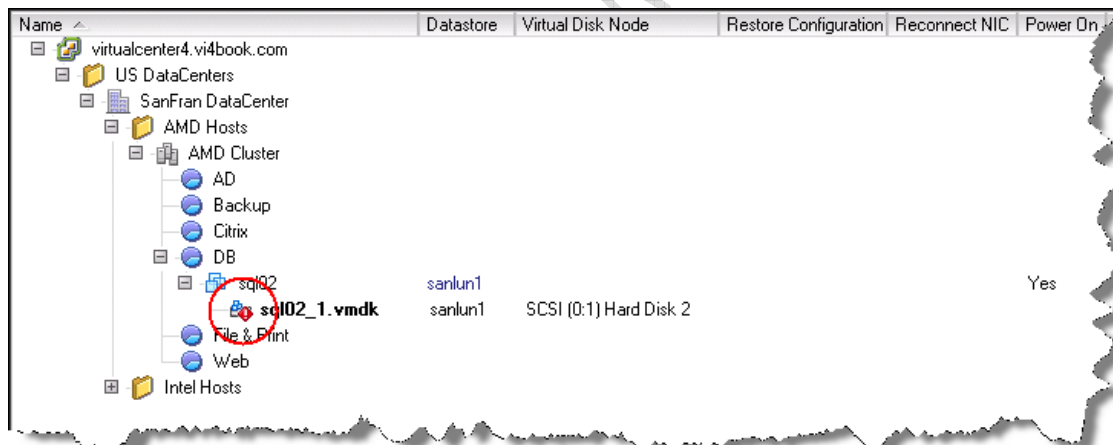
vDR does support the use the restore process to overwrite an existing VM or virtual disk – but care must be taken to avoid an scenario where your VM is in *worse* state after the restore process. That rather defeats the object of restore process. But you would be surprised to know how often this happens. I've seen it in business I've worked in. In one notable case when I was a junior contactor and the bad restore was actually carried out by senior manager of IT! Unfortunately, neither I nor

anyone else was around to stop this “test” restore wiping out end-users data. I won’t mention any company names. But remember a bad restore can cause as much damage as no backup strategy at all or infrequent backups. In this case I’m going to restore the virtual disk and in doing so – overwrite the original virtual disk.

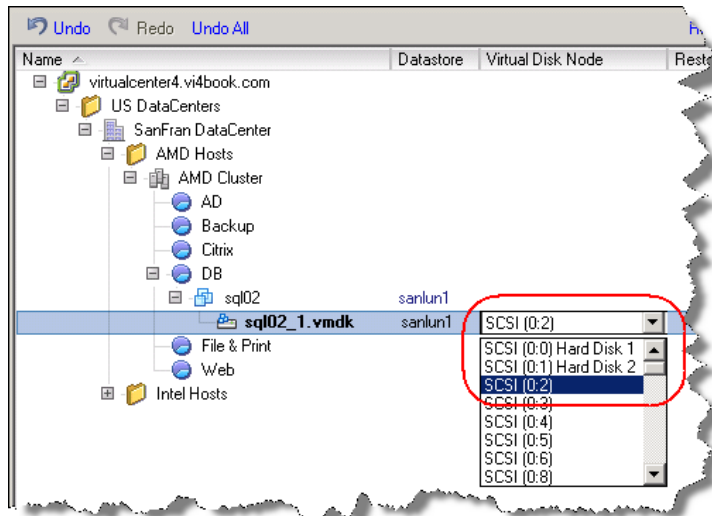
To restore an individual virtual disk start the restore wizard and select the disk that needs be restored like so:



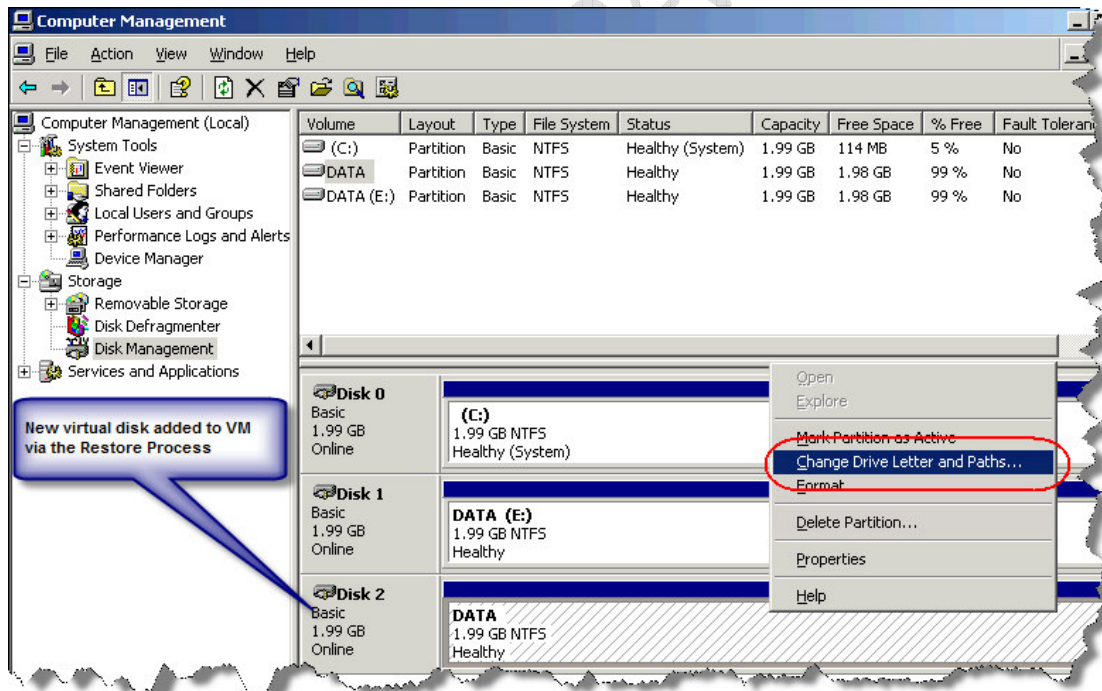
In the next part of the wizard notice the red ! exclamation mark.



By continuing you are agreeing to have the original virtual disk overwritten. Also notice how the option “Restore Configuration” and “Reconnect NIC” is not available – because we are only restoring an individual virtual disk. Continuing in this manner is clearly quite dangerous for reason I hope I made clear earlier. It is possible to restore the virtual disk to different SCSI ID. This would modify the VM to have 3 virtual disks – the boot, the original data drive, and the restore data drive. This configuration can be achieved by clicking the pull-down arrow next to SCSI Disk reference in the Virtual Disk Node as shown below:



As you can see once I select a different SCSI ID (SCSI 0:2) the red ! exclamation mark is cleared on the VM. Once you click next and Restore, the virtual machine will power off, and then restore process will begin. Currently, the vDR does not power on the VM after the restore process has completed. Once you manually power on the VM you will need to use Microsoft's Disk Management – to get the restored virtual disk (in my case SCSI0:2) to be visible, by assigning a drive letter to it.



**WARNING:**

IF you do the restore process are overwrite the original disk, Windows will see the restore virtual disk as brand new device – as such drive letter mappings are not restored. The Windows defaults allocating the next available drive letter to the restored disk.

## Restoring Individual Files within a Virtual Disk

File level restores are not currently supported in the beta release.

## VMware Data Recovery Appliance - Operational Maintenance

### Bring a VM to Compliance

No information at this point

### Locking a Restore Job


No information at this point

### Starting a Integrity Check

It is possible to select a either disk or mounted network location and check it to see if the metadata and de-duplication data contains is in a good state.

1. Select the **Configuration Tab**
2. In the **Configure Pane**, select the **Destinations** link
3. **Select the datastore**, and then click the **Integrity Check...** link

[chapter19-integrity-check.jpg]



### Extending a Backup Destination

If the size of the backup destination is increased, then you can use the extend option to force the vDR to extend the file system of the disk to take up the new space

1. Select the **Configuration Tab**
2. In the **Configure Pane**, select the **Destinations** link
3. **Select the datastore**, and then click the **Extend...** link

#### Note:

This can take some time complete and then refresh - so please be patient.