PERFORMING SINGLE FILE RECOVERY FROM NETAPP SNAPSHOTS OF NFS DATASTORES

Version 1.3
January 2011
Michael Arndt
Michael.Arndt@netapp.com

Abstract
This paper documents the configuration and tools that can be used to perform recovery of single files from NetApp SnapShots of VMware virtual machines within NFS datastores.

TABLE OF CONTENTS

1 INTRODUCTION

This paper documents the configuration and tools that can be used to perform recovery of single files from NetApp SnapShots of VMware virtual machines within NFS datastores. A variety of configurations, tools, and techniques can be used to accomplish these tasks for both Microsoft Windows and Linux Guest Operating Systems. This paper will focus on tools that are freely available, and can be enhanced with custom scripting, to help automate the processes around single file recovery in a VMware environment that uses NFS datastores. Note that the techniques given in this document will not work directly with VMFS datastores. The diagram below will illustrate the connectivity required for some common single file recovery scenarios.

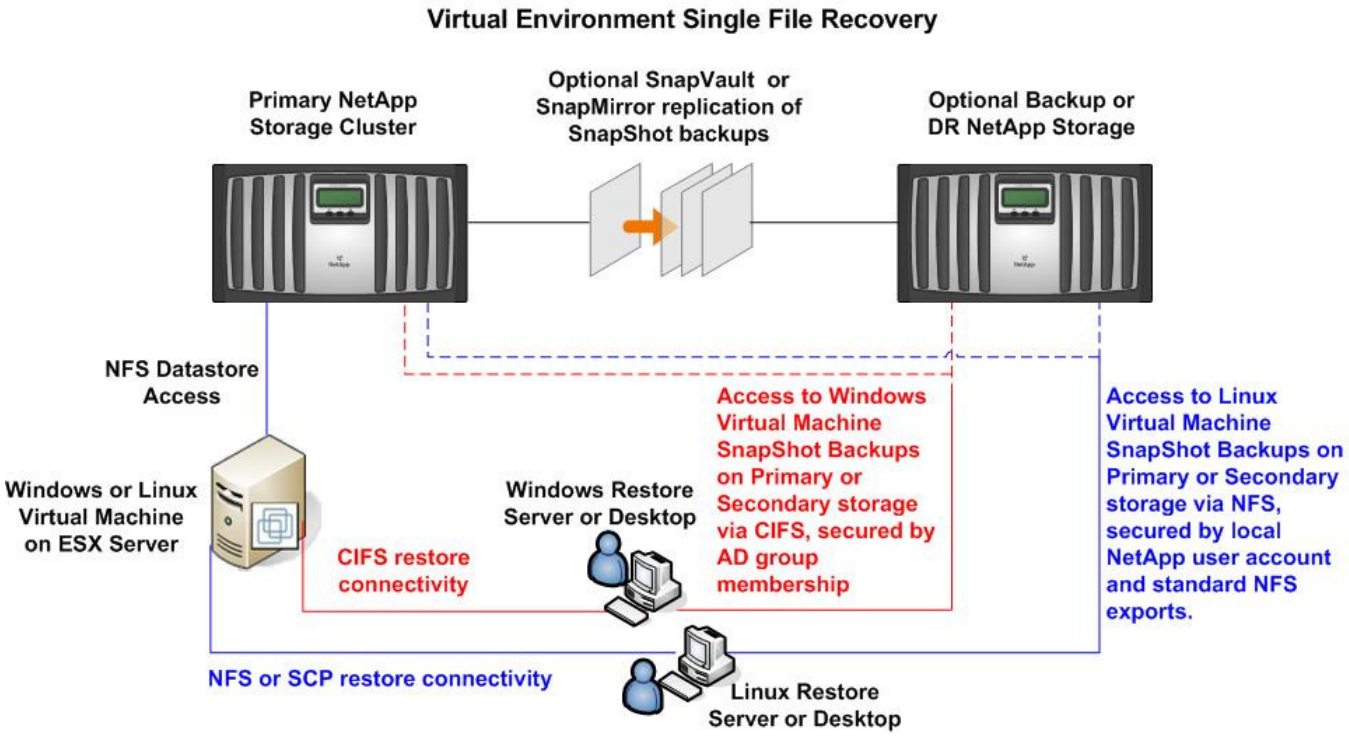

2 PERMISSIONS AND AUTHORIZATION ON NFS DATASTORES

In this section, we will review the methods available for allowing access to the NetApp SnapShots of the VMware virtual disks (VMDKs) that contain the data on which single file recovery is to be performed. This section should be reviewed carefully in order to allow access to the virtual machine disk files so that single file recovery can be performed as described in the subsequent sections of this document.

The virtual disks for virtual machines running Microsoft Windows based operating systems will typically be formatted with the NTFS filesystem. Accessing the virtual disk files via a Microsoft Windows client (either physical or virtual) is preferred, as a Microsoft Windows client will understand both the NTFS filesystem, as well as the NTFS ACLs that reside on the virtual disk. As such, we will present methods for accessing the virtual disks of Microsoft Windows based virtual machines via the CIFS protocol.

The virtual disks for virtual machines running Linux based operating systems will typically be formatted with either the ext3 or ReiserFS filesystem. Whatever filesystem is on the virtual disk, using a Linux client (either physical or virtual) is ideal for performing single file recovery, as most Linux distributions will be able to understand both the filesystem that resides on the virtual disk, as well as the permission bits that accompany the data. As such, we will present methods for accessing the virtual disks of Linux based virtual machines via the NFS protocol.

2.1 Environments using UNIX security style on NetApp NFS datastores

The use of UNIX security style on NetApp storage is the most popular in VMware environments that use NFS datastores, as this is the native security model that the NFS protocol supports. When using UNIX security style, VMware ESX server creates the virtual disk files (VMDK's) with permission mode bits of 600. The implications of this are as follows:

- From a NFS perspective, this means that only the root account on NFS clients, and only on those clients that are given root privileges to the NFS export containing the datastore, can access the virtual disk files in order to perform single file recovery.
- From a CIFS perspective, this means that any account accessing the datastore via the CIFS protocol must map to the root account on the NetApp storage system.
- There is no need to write to the virtual disk file, and therefore no need for a FlexClone of the NetApp Snapshot containing the datastore, as the tools we will present for accessing the virtual disk can do so without committing any writes to the virtual disk filesystem.

2.1.1 Access to SnapShots of Microsoft Windows virtual machines

On a Microsoft Windows client, we can access the virtual disk files directly in the NetApp SnapShot via the CIFS protocol. The requirements in this case are as follows:

- You must create a CIFS share of the NetApp NFS datastore. Optionally, this CIFS share can point directly to the .snapshot directory of the datastore.
- The "create_ucode" and "convert_ucode" volume options must be turned on for the NetApp volume containing the datastore, prior to taking snapshots that are to be accessed for the purpose of single file recovery. See the following knowledge base article for more information on these settings: https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb23656
- You must perform user mapping on the NetApp storage system to map a Microsoft Windows AD account or accounts to the root account. This can be accomplished in one of two ways:
    - You can directly setup a one-for-one mapping in /etc/usermap.cfg, with an entry such as "<AD-DOMAIN>\<AD-USERNAME> => root".
    - You can create a group of AD accounts that should be given this privilege and add that group to the Local Administrators group on the NetApp storage system. After doing this, you can run "options wafl.nt_admin_priv_map_to_root on" in order to automatically map any administrative account in the AD domain to the root user on the NetApp storage system.
    - Once your mapping is setup correctly, validate it using the "wcc" command on your NetApp storage system. For example, run "wcc –s DOMAIN\account" to validate that the given AD account is being mapped to root. If you recently modified the way that mapping was done, you may need to flush the credential cache on the NetApp storage system by running "wcc –x".

2.1.2 Access to SnapShots of Linux virtual machines

Since we are using a Linux client, we can easily access the virtual disk files via NFS, as long as we are given access to the NFS export containing NetApp SnapShots of the datastores on which our virtual machines reside. The requirements in this case are as follows:

- You must have the ability to create a FlexClone of the NetApp SnapShot containing the NFS datastore. This is required because access to the filesystems on the virtual disk files will typically involve some log replay on the ext3 or ReiserFS filesystem. The log replay will require a small amount of data to be written to the virtual disk, which can not be performed if you are accessing the virtual disk in a NetApp read only SnapShot. A FlexClone is essentially a writable version of a NetApp SnapShot. The FlexClone feature is a licensed option on a NetApp storage system.
- You must given read/write and root access to the NFS client that is accessing the cloned version of the datastore.

2.2 Environments using NTFS security style on NetApp NFS datastores

The use of NTFS security style on NetApp storage is also supported for VMware environments that use NFS datastores. NetApp storage systems are unique, in that they allow the administrator to configure access to NTFS security style data directly from NFS clients (as well as allowing access to UNIX security style data from CIFS clients). When using NTFS security style on a NetApp NFS datastore, special care must be taken to give proper access to the root user on the VMware ESX servers. There are two options for doing this:

- Edit the /etc/usermap.cfg configuration file on the NetApp storage system and add a line such as "Administrator <= root" in order to map the root account from NFS clients to the Administrator account.
- Run "options cifs.nfs.root_ignore_acl on" on the NetApp storage system so that the root user from NFS clients is allowed to bypass the NTFS security on the NetApp volume used for the datastore.

2.2.1 Access to SnapShots of Microsoft Windows virtual machines

On a Microsoft Windows client, we can access the virtual disk files directly in the NetApp SnapShot via the CIFS protocol. The requirements in this case are as follows:

- Since the virtual disk files will be created from a NFS client that is specifying that only the file owner has privileges to read the file, this will result in a NTFS ACL that gives only full control to the Administrator account. As such, the AD account you are using on your CIFS client must either be a Domain Administrator, or they must be in an AD group that is a member of the Local Administrators group on the NetApp storage system.
- The "create_ucode" and "convert_ucode" volume options must be turned on for the NetApp volume containing the datastore, prior to taking snapshots that are to be accessed for the purpose of single file recovery. See the following knowledge base article for more information on these settings: https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb23656

2.2.2 Access to SnapShots of Linux virtual machines

Since we are using a Linux client, we can easily access the virtual disk files via NFS, as long as we are given access to the NFS export containing NetApp SnapShots of the datastores on which our virtual machines reside. The requirements in this case are as follows:

- You must have the ability to create a FlexClone of the NetApp SnapShot containing the NFS datastore. This is required because access to the filesystems on the virtual disk files will typically involve some log replay on the ext3 or ReiserFS filesystem. The log replay will require a small amount of data to be written to the virtual disk, which can not be performed if you are accessing the virtual disk in a NetApp read only SnapShot. A FlexClone is essentially a writable version of a NetApp SnapShot. The FlexClone feature is a licensed option on a NetApp storage system.
- You must given read/write and root access to the NFS client that is accessing the cloned version of the datastore.

3 ACCESSING VIRTUAL DISKS FOR SINGLE FILE RECOVERY

The following sections give examples of how to access the actual data within a virtual disk file once you have obtained access to the file using one of the configurations described in section 2.

3.1 Procedures for Microsoft Windows virtual machines

The VMware Virtual Disk Development Kit comes with a utility called vmware-mount.exe. This utility can be used to access the contents of a NTFS formatted virtual disk, directly via a CIFS share of the NetApp SnapShot of the NFS datastore. Here is an example of using the vmware-mount.exe utility to access the contents of a virtual disk file:

C:\> vmware-mount.exe z: \\<Hostname>\<Share>\.snapshot\<SnapShot>\<VM>\<VM>.vmdk

The above command would connect the contents of the given virtual disk to the Z:\ drive on the Microsoft Windows system from which the command was run. Once this is done, tools such as Microsoft Windows Explorer or Robocopy can be used to copy data from the Z:\ drive back to the original virtual machine or an alternate location. When the single file recoveries are complete, use the following command to disconnect the virtual disk file from the drive letter on which it was mounted:

C:\> vmware-mount.exe z: /d

3.1.1 Using the VMDKMounter utility from NetApp

The VMDKMounter utility from NetApp is a GUI front end wrapper around the vmware-mount.exe utility from VMware. Once VMDKMounter and the Virtual Disk Development Kit are installed, you will be able to browse to the VMDK file using the standard Microsoft Windows Explorer utility, right click on the VMDK, and select **VMDK Mount** from the drop down list:

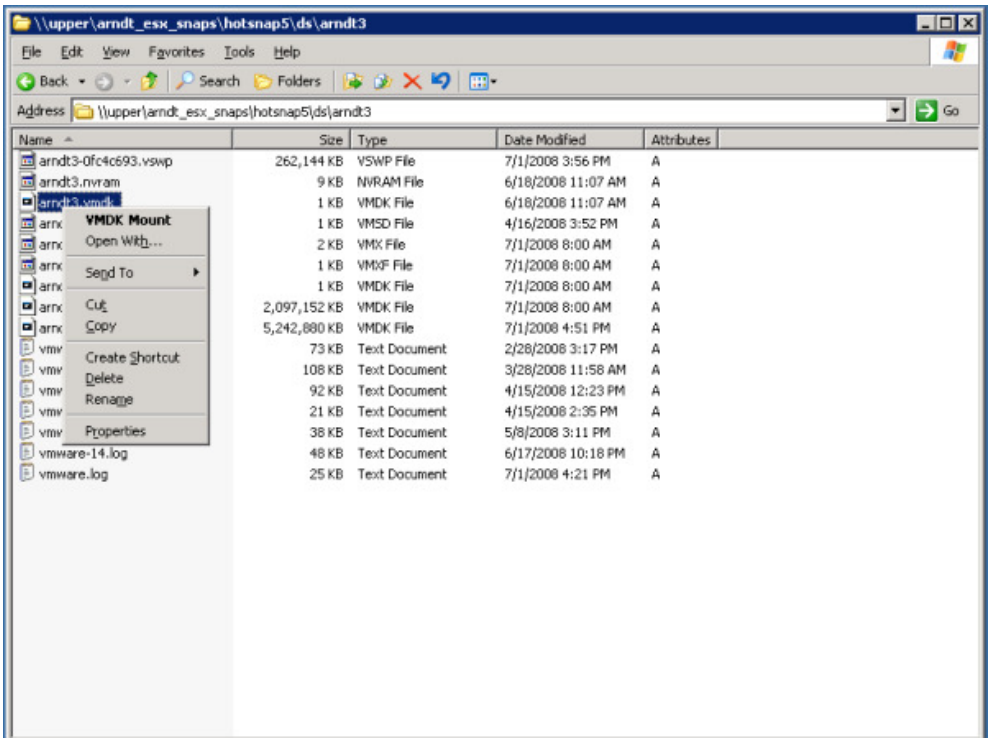

After selecting the **VMDK Mount** option, VMDKMounter will give you a simple interface which you can use to map the volumes (partitions) on a virtual disk to the drive letters on which you would like to connect them:

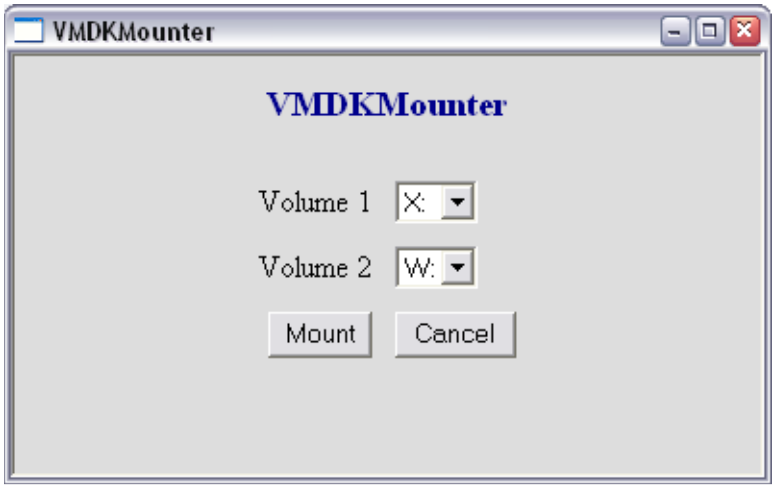

After completing your single file recovery, right click on the drive letters where the VMDK is mounted and select **VMDK UnMount**:

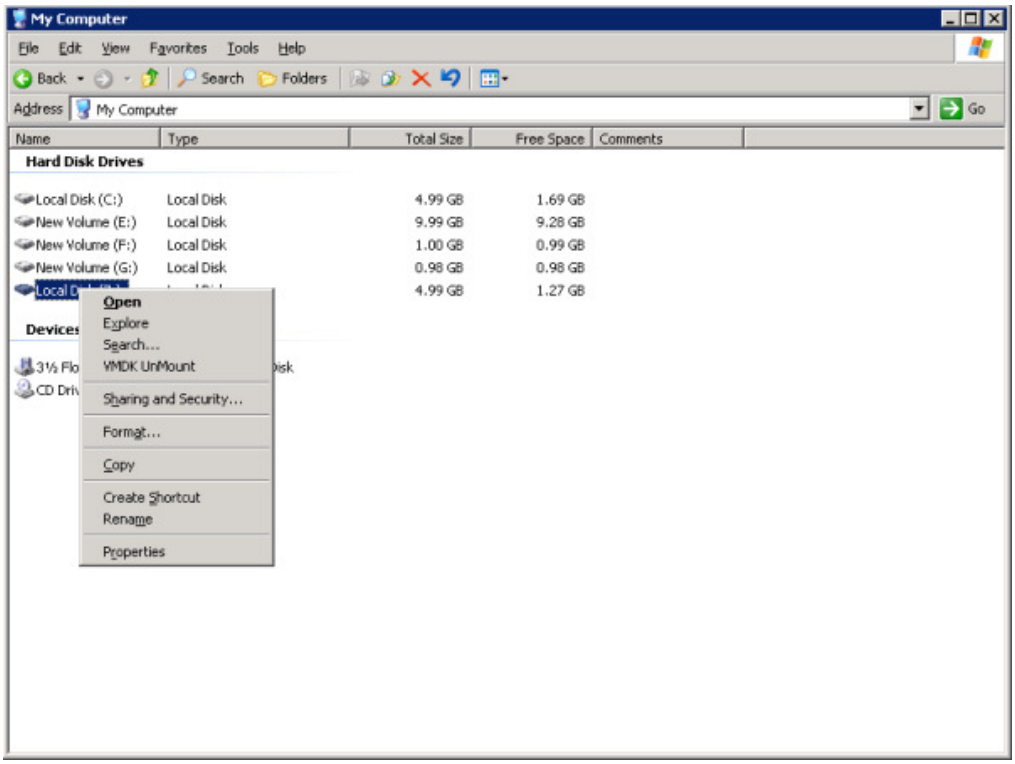

The **VMDK UnMount** operation will disconnect the partition from the drive letter which you chose to unmount.

3.2 Procedures for Linux virtual machines

Linux systems have ability to connect directly to virtual disk files via a loopback device, which is a feature built in to the Linux operating system. As previously mentioned, this will require write access to the virtual disk file in order to replay the virtual disk filesystem logs when mounting via the loopback device. In order to obtain write access to a NetApp SnapShot, we can use the NetApp FlexClone feature. Since there are a number of steps required to connect to Linux VMDKs in NetApp SnapShots, we will show how this is scripted in the next section.

3.2.1 Using the vmdkmounter.pl script from NetApp

The vmdkmounter.pl script was created to automate the many steps required in order to connect to Linux VMDKs in a NetApp SnapShot. This Perl script will automatically detect the number and type of partitions on a VMDK, including any Logical Volumes, and connect them via a loopback device in order to allow access to the filesystems. The requirements for using vmdkmounter.pl are as follows:

1. An account with the appropriate role must be created on the NetApp storage system or vFiler.
2. The vfiler.vol_clone_zapi_allow option must be enabled on the physical filer if you are using vmdkmounter.pl against a vFiler
3. The NetApp system must have a FlexClone license.
4. If you use LVM, you must not use vmdkmounter.pl on a client that has identically named volume groups or logical volumes as exist on the VMDKs being restored from.

The following commands show how to create an account named "vmware_sfr" with the appropriate role on the NetApp storage system or vFiler:

```
> useradmin role add vmware_sfr_role -a login-http-admin,\
  api-system-get-version,api-snapshot-list-info,api-volume-*,\
  api-file-list-*,api-nfs-exportfs-append-rules
> useradmin group add vmware_sfr_group -r vmware_sfr_role
> useradmin user add vmware_sfr -g vmware_sfr_group
```

The following examples show the basic syntax for using vmdkmounter.pl. To connect to a Linux VMDK in a NetApp SnapShot:

**linux# vmdkmounter.pl -o connect -h storagehost -u storageuser -m mountpoint -f vmfolderpath **
-s snapshot [-v]

NOTE: Starting with version 2.3, you must pass in the path to the folder/directory of the virtual machine that you are restoring from. Previous versions required the full path to the actual "-flat.vmdk" file, but that is not allowed starting with version 2.3. This was done to support logical volumes that span multiple VMDK files.

Once the partitions and any logical volumes on the VMDK are mounted, tools such as scp or rsync can be used to copy data from mountpoints back to the original virtual machine or to an alternate location. When the restores are completed, use the following syntax to disconnect a Linux VMDK:
linux# vmdkmounter.pl -o disconnect -h storagehost -u storageuser -m mountpoint [-v]

To illustrate the usage of vmdkmounter.pl, here is an example of connecting to a Linux VMDK:

```
linux#  ./vmdkmounter.pl -o connect -h arndtvf -u vmware_sfr -m /sfr -f \
        /vol/arndt_esx/ds/arndt_sles10_lvm -s sv_nightly.0
Enter password for vmware_sfr@arndtvf:
Wed Mar 11 15:41:10 2009: Checking connectivity for vmware_sfr@arndtvf.
Wed Mar 11 15:41:10 2009: Connectivity OK - NetApp Release 7.3.1.
Wed Mar 11 15:41:10 2009: Checking storage system configuration.
Wed Mar 11 15:41:10 2009: Connecting to clone on storage system.
Wed Mar 11 15:41:12 2009: Mounting partition 1 under /sfr/partitions/arndt_sles10_lvm/1.
Wed Mar 11 15:41:13 2009: Mounting logical volume /dev/system/root under /sfr/LV/system/root.
Wed Mar 11 15:41:13 2009: Mounting logical volume /dev/system/testlv under /sfr/LV/system/testlv.
```

Also, the corresponding example of disconnecting the VMDK that was connected in the example above:

```
linux# ./vmdkmounter.pl -o disconnect -h upper -u root -m /sfr
Enter password for vmware_sfr@arndtvf:
Wed Mar 11 15:41:52 2009: Checking connectivity for vmware_sfr@arndtvf.
Wed Mar 11 15:41:52 2009: Connectivity OK - NetApp Release 7.3.1.
Wed Mar 11 15:41:52 2009: Unmounting /sfr/LV/system/root.
Wed Mar 11 15:41:52 2009: Unmounting /sfr/LV/system/testlv.
Wed Mar 11 15:41:53 2009: Unmounting /sfr/partitions/arndt_sles10_lvm/1.
Wed Mar 11 15:41:53 2009: Disconnecting from storage system.
```

In order to see the exact commands on the Linux client and the NetApp storage system used to accomplish the steps during the connect and disconnect operations, simply append a **"-v"** switch to the vmdkmounter.pl command line.

4 RESOURCES

In order to obtain the VMDKMounter application for performing single file restores from Windows virtual machines, or the vmdkmounter.pl script for performing single file restores from Linux virtual machines, contact the author of this document at Michael.Arndt@netapp.com.

Further resources pertaining to the concepts discussed in this document are as follows:

- VMware Virtual Disk Development Kit:
  - http://www.vmware.com/support/developer/vddk/
- Multiprotocol Permission on NetApp systems:
  - http://media.netapp.com/documents/wp_3014.pdf
- Configuring SSH PublicKey Authentication on a NetApp storage system:
  - http://now.netapp.com/NOW/knowledge/docs/ontap/rel724/html/ontap/sysadmin/9secadm9.htm