


Disaster Recovery Using NetWorker and Data Domain



Michael Dick

Principal Solutions Architect

Dell EMC

michael.dick@dell.com

Table of Contents

Introduction	3
Disaster Recovery	4
Requirements	5
Overall Methodology	6
<i>Creating the Recovery Environment</i>	<i>7</i>
<i>Install DR NetWorker server.....</i>	<i>8</i>
<i>Create DDBoost Storage Unit</i>	<i>8</i>
<i>Fast Copy Data between Storage Units</i>	<i>9</i>
<i>Create the NetWorker Devices.....</i>	<i>10</i>
<i>Finding the client ID and other necessary information.....</i>	<i>11</i>
<i>Create the Client Definition to be recovered</i>	<i>13</i>
<i>Rebuild the CFI Entries.....</i>	<i>16</i>
Recovering data in the DR data zone.....	17
<i>Recovering of savesets and specific Files</i>	<i>17</i>
Client Migrations between NetWorker Servers	18

Disclaimer: The views, processes or methodologies published in this article are those of the author. They do not necessarily reflect Dell EMC's views, processes or methodologies.

Introduction

This Knowledge Sharing paper provides guidance on how to achieve disaster recovery of a client using NetWorker™ software and Data Domain™ on an existing or new NetWorker server without the need to execute a full bootstrap recovery of this server. The benefits of this approach are mainly shortened time to client recovery and protection of the actual backup data. Another benefit of this approach is that an already running NetWorker server can be used to recover a client's data without compromising already existing clients on this NetWorker server. While NetWorker 9 is used throughout the paper, the concepts shown hold true for pre-NetWorker 9 systems as well.

Disaster Recovery

Before we go into Disaster Recovery (DR) itself, we must define what constitutes a disaster for the purpose of this paper. In this context, a disaster is a situation where the primary NetWorker server is unavailable to execute recoveries. Reasons for this unavailability do not matter. It does not matter if the NetWorker server is down merely because of a hardware fault, or if a complete data center has gone down because of some natural or man-made disaster. It may just be that the network between two data centers is unavailable, and a restore has to be executed in the remote (from the NetWorker server's point of view) data center. If such a failure occurs, we have several ways of effecting a restore given that backups, including bootstrap backups, have been cloned to the remote data center.

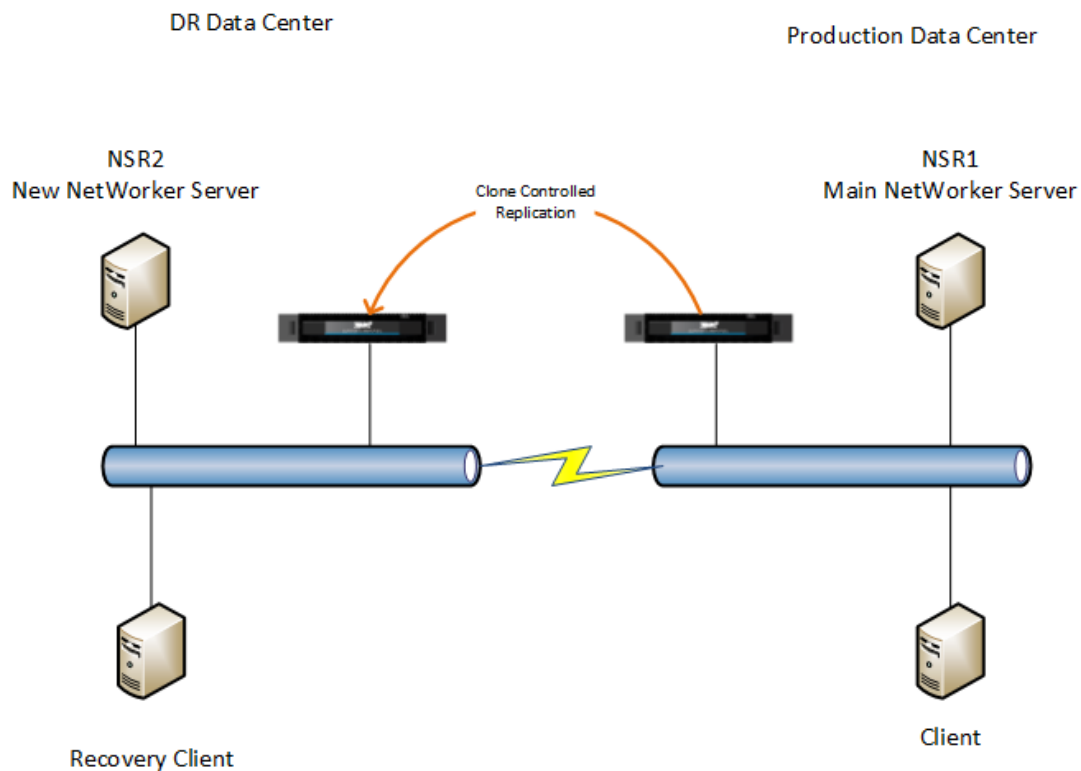
We can restore the complete NetWorker server using a bootstrap backup (if the bootstrap backup is available), or we can just restore a single (or more clients) by scanning in the cloned save sets to a new or existing NetWorker server. The method of DR depends on the cause of the failure. If the main data center will be down for a long time, it makes sense to bootstrap the NetWorker server in the remote location. If the main data center is only unreachable for a short amount of time (hours or days), it may be more efficient to only restore certain clients or save sets if needed. This paper will introduce methodologies to recover client data in a DR data center without having to completely rebuild the NetWorker server. We will discuss recovering single files, complete save sets, and application data (SQL, Oracle, etc.) to clients residing in a DR data center. Other uses for this methodology include:

- Client migrations between NetWorker servers without losing the ability to execute file level recoveries
- Restore/DR Tests in isolated environments

Requirements

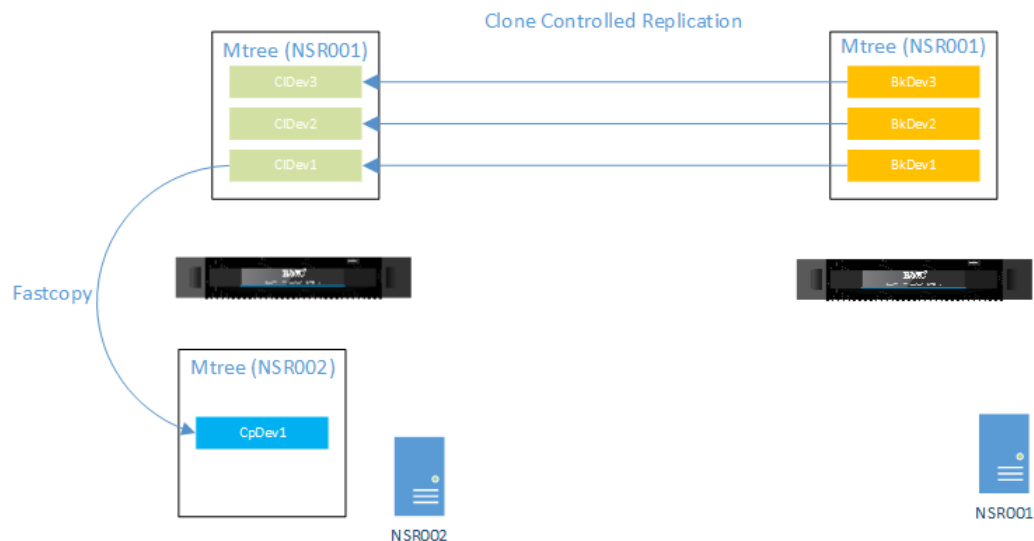
The requirements for a Disaster Recovery of client in a remote datacenter are as follows:

- Two Data Domain systems. One in the main data center and one in the DR center. Backups to be recovered must exist on the DR Data Domain system
- NetWorker is already installed or there is a VM (or physical server) readily available to install NetWorker
- A client is available for the recovery and is known to NetWorker
- If OS and/or application backups are not available for restore, the application must be installed and correctly configured to execute application data recovery



Overall Methodology

The overall methodology for these recoveries is the same regardless of type of recovery (single file, save set, or application data). The overarching idea here is to protect the actual backup data (backup or clone) at all cost. One could imagine that importing the actual clone device could lead to data loss if we were to label this device by accident. By using fastcopy to copy the data to a new mtree, we isolate the data to be restored from the backed up data. We then can import these backups to a new NetWorker server without risking corruption to the cloned data.



The diagram above shows the basic premise. A NetWorker server (NSR001) has backups in the production data center and clones in the DR data center. The clone devices can be fast copied to the mtree owned by the DR NetWorker server (NSR002). NSR002 can now scan the fast copied device and make use of the data on this device. If we have separation of index and backup data then we need to fast copy both the data clone and the index clone device.

Creating the Recovery Environment

This section will describe the creation of the recovery environment. It is the same for all examples below. Any differences will be pointed out in the examples themselves. The creation of the recovery environment can be broken down into 6-six steps

1. Install new NetWorker server and configure a Device on the DR Data Domain (this will create the default NetWorker DDBoost storage unit).
2. Fast copy device(s) from the originating NetWorker server mtree to the new NetWorker server mtree. The Devices are folders beneath the mtree.
3. Create the NetWorker device(s) using GUI or command line.
4. Create the client to be recovered on the NetWorker DR server. This client must use the same client ID as the client on the originating server.
5. Scan the device using scanner `-i -c <clientname> <devicename>`. This will scan the save sets of this device into the media database and recreate the client indexes on the DR NetWorker server.
6. Create the client definition to be recovered to if this is a directed recovery.

Please note: If the methodology described is used to affect test restores in an isolated environment, it is recommended to clone the savesets needed to a separate clone device and to fast copy this clone device. This will speed up the scanning process needed. An example clone command to clone all backups since yesterday to a separate device is below:

```
Nsrclone -b <clonePool> -S -t yesterday -c <clientname>
```

This command will copy all savesets for a client from yesterday to a specific pool. This pool can then be fast copied to the DR storage unit.

Install DR NetWorker server

The NetWorker server used for the DR should be installed on the same OS and with the same version as the originating NetWorker server. This NetWorker server can be managed by its own NetWorker management Console (NMC) or by an already existing NMC server. After installation, the NetWorker server is fully licensed for 30 days. Please consult the NetWorker installation guide for information about installing NetWorker on different Operating Systems. The screenshot below shows the packages for NetWorker 9 that need to be installed

```
=====
Installing:
lgtoauthc          x86_64          9.0.1.0-1          /lgtoauthc-9.0.1.0-1.x86_64
lgtoctrl           x86_64          9.0.1.0-1          /lgtoctrl-9.0.1.0-1.x86_64
lgtoadmin          x86_64          9.0.1.0-1          /lgtoadmin-9.0.1.0-1.x86_64
lgtonmc            x86_64          9.0.1.0-1          /lgtonmc-9.0.1.0-1.x86_64
lgtonode           x86_64          9.0.1.0-1          /lgtonode-9.0.1.0-1.x86_64
lgtooserv          x86_64          9.0.1.0-1          /lgtooserv-9.0.1.0-1.x86_64
lgtoxtctrl         x86_64          9.0.1.0-1          /lgtoxtctrl-9.0.1.0-1.x86_64
```

In this case, the new NetWorker server is installed as a completely self-contained unit which can be deleted after the recoveries are complete. After installing the packages, complete the configuration of authc and NMC before starting and using NetWorker.

Create DDBoost Storage Unit

As the DR Data Domain is already configured for DDBoost, we just need to create the DDBoost storage unit for the DR NetWorker server. The DDBoost storage unit can be created manually or we can just create a Data Domain Device on the DR NetWorker server to create a DDBoost storage unit with default values (recommended). If you create a device, ensure that the device name will be unique.

Fast Copy Data between Storage Units

We then can copy either the complete originating storage unit to this new mtree, or we can copy only the devices we need.

To copy the complete storage unit, use the following command:

```
Filesys fastcopy source /data/col1/<srcname> destination /data/col1/<destname>
```

Where <srcname> and <destname> are the source and destination mtree respectively.

Devices are merely folders beneath the mtree. To copy a device, use the command below.

```
filesystem fastcopy source /data/col1/<srcname>/<folder> destination  
/data/col1/<destname>/<folder>
```

Where <srcname> and <destname> are the source and destination folders and <folder> is the device folder beneath the mtree.

The screenshot below shows the command to create a fast copy of CloneDevice1 in the DR NetWorker storage unit

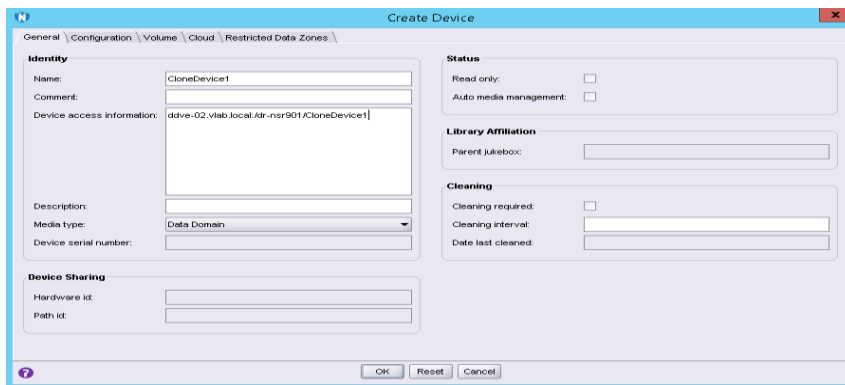
```
sysadmin@ddve-02# filesystem fastcopy source /data/col1/nsr901/CloneDevice1 destination /data/col1/dr-nsr901/CloneDevice1  
(00:04) Waiting for fastcopy to complete...  
Fastcopy status: fastcopy /data/col1/nsr901/CloneDevice1 to /data/col1/dr-nsr901/CloneDevice1: copied 567 files, 654 directories in 6.08 seconds  
sysadmin@ddve-02#
```

It is recommended to use the original device name when copying device folders only. Only the device(s) holding the actual client backups need to be scanned in. The save sets on this device will be used to recreate the CFI for the client to be recovered.

Create the NetWorker Devices

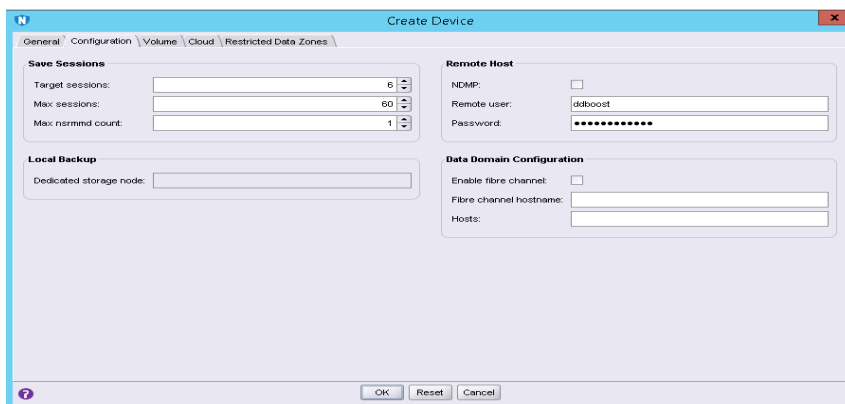
One can use the device wizard or device properties screen on the GUI or command line to create the necessary device definition. While the device definition can be created without the pool to be available, we need to create the appropriate pool (same name as the originating server) before we can mount the device for data access. Also, if creating the device using the wizard, one must ensure that the device to be created is not getting labeled in the process as this would delete the data on the device. The screenshots below show how to create the device using the device property window.

Open the device properties window and add the device name and device access information in the General tab (it is recommended to keep the device name the same as the original device)



The screenshot shows the 'Create Device' dialog box with the 'General' tab selected. The 'Identity' section contains fields for Name (CloneDevice1), Comment, Device access information (ddve-02.vlab.local:ddr-nstr901/CloneDevice1), Description, Media type (Data Domain), and Device serial number. The 'Status' section has checkboxes for Read only and Auto media management. The 'Library Affiliation' section has a Parent jukebox field. The 'Cleaning' section has checkboxes for Cleaning required, Cleaning interval, and Date last cleaned. The 'Device Sharing' section has fields for Hardware id and Path id. At the bottom are OK, Reset, and Cancel buttons.

Click on the configuration tab and add the ddbboost user name and password.



The screenshot shows the 'Create Device' dialog box with the 'Configuration' tab selected. The 'Save Sessions' section has fields for Target sessions (6), Max sessions (60), and Max nsrmmid count (1). The 'Local Backup' section has a Dedicated storage node field. The 'Remote Host' section has checkboxes for NDMP, Remote user (ddbboost), and Password (masked with dots). The 'Data Domain Configuration' section has checkboxes for Enable fibre channel, Fibre channel hostname, and Hosts. At the bottom are OK, Reset, and Cancel buttons.

Click OK to create the device.

Finding the client ID and other necessary information

We must create the clients to be recovered on the DR NetWorker server with the same client ID as on the originating NetWorker server. We must also determine the originating pool names and label templates in order to create the appropriate pools to mount the DR devices for recovery. If the originating server is accessible, we can get this information by running the commands below on the original server without the need to scan in the device first. If the originating server is not accessible, we must scan the device we just created using `scanner -m <devicename>`. Scanner `-m` scans savesets into the media database. Depending on the size of the volume, this command can take some time to populate the media management database. If this is just a DR test to be executed in an isolated environment, then it would be beneficial to take the time and create another clone of the needed data on a separate clone device to shorten the scan time. Once the command is done, we can query the media management database to gather the necessary information to create the appropriate client and pool definitions. To find the client ID, use:

```
mminfo -q "client=<clientname>" -r "clientid" | tail -1
```

The example below shows the command to find the client id for a client with the name `linclient1.vlab.local`

```
[root@dr-nsr901 linux_x86_64]#  
[root@dr-nsr901 linux_x86_64]# mminfo -q "client=linclient1.vlab.local" -r "clientid" | tail -1  
8f4358c2-00000004-57aca84b-57aca8d7-00095000-691c0156  
[root@dr-nsr901 linux_x86_64]#
```

The example below shows the command to find volume and pool information

```
[root@dr-nsr901 linux_x86_64]# mminfo -q "client=linclient1.vlab.local" -r "clientid,volume,pool" | tail -1  
8f4358c2-00000004-57aca84b-57aca8d7-00095000-691c0156 ClonePool1.001 ClonePool1  
[root@dr-nsr901 linux_x86_64]#
```

We now can use this information to determine how many clients we have on this volume and the client IDs for these clients

```
[root@dr-nsr901 linux_x86_64]# mminfo -q "pool=ClonePool1" -r "client,clientid" | sort -u  
client      client ID  
linclient1.vlab.local 8f4358c2-00000004-57aca84b-57aca8d7-00095000-691c0156  
oracle-01.vlab.local b0a7486c-00000004-57aca84c-57aca8d7-000a5000-691c0156  
win2012-01 d9607099-00000004-57aca84a-57aca849-00085000-691c0156  
[root@dr-nsr901 linux_x86_64]#
```

Finding the needed information easily requires some knowledge of the source data zone. If this information is not known, we can use `mminfo -avot` to find the volume information and then go from there.

```
[root@dr-nsr901 linux_x86_64]# mminfo -avot | awk '{ print $1 }' | sort -u | grep -v volume  
ClonePool11.001  
[root@dr-nsr901 linux_x86_64]#
```

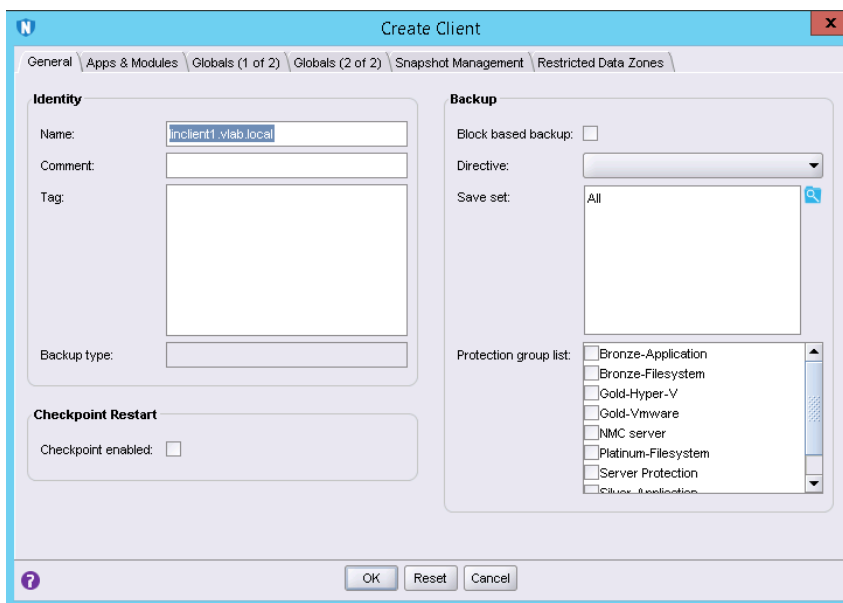
Once we have this information we can use `mminfo` to acquire the client names, client ID's, and pool names by querying the volume instead of the pool.

Create the Client Definition to be recovered

Before we can scan in the savesets and indexes for the client to be recovered, we must recreate the client to be recovered on the new NetWorker server. This client MUST be recreated with its original client ID!

Also, if the client to be created is not resolvable/reachable from the DR NetWorker server, the client must be created using nsradmin in offline mode. This is because NetWorker will attempt to communicate with the client, and the client creation will fail if communication cannot be established.

If the client is reachable, we can use the client properties window to recreate the client definition on the DR NetWorker server. Ensure that the GUI is in Diagnostic Mode. Go to Protection and right click on Clients. Choose New, which will open the Create Client window. Add the client name in the General section.



Click on Globals 1 of 2 and add the discovered Client ID for the client to be created.

Create Client

General | Apps & Modules | Globals (1 of 2) | Globals (2 of 2) | Snapshot Management | Restricted Data Zones | Legacy

Identity

Client id: 4b-57ac68d7-00095000-691c0156

Aliases:

Optimizations

Parallelism: 4

Server network interface:

Priority: 500

Physical client parallelism: ☐

Parallel save streams per save set: ☐

Save session distribution: max sessions

Add an *@* into the Remote Access field and click on OK to create the client.

Create Client

General | Apps & Modules | Globals (1 of 2) | Globals (2 of 2) | Snapshot Management | Restricted Data Zones | Legacy

Configuration

Remote access: @

Index path:

Executable path:

Backup target disks:

Owner notification:

Centralized logs collection: ☒

Windows Client

Hard links: ☐

Short filenames: ☐

Storage Nodes

Autoselect storage node: ☐

Storage Nodes:

Recover storage nodes:

Archive Management

Archive services: ☐

It must be noted that adding *@* into the remote access field allows any user from any client to access the backups of this client. This is NOT a best practice and should NOT be done on a production client because of data security reasons. It is done here because it is assumed that we are in a DR situation where ease of recovery overrides data security.

If the client to be created is not reachable from the DR NetWorker server, or if there are no DNS entries for this client in the DR data center, the client definition must be created

using nsradmin in offline mode. It must be noted that manipulating the NetWorker data base (nsrdb) in offline mode is dangerous, as there will be no checks. Doing this MAY result in a corrupt NetWorker data base, and as such, editing the NetWorker data base in offline mode should never be done on a production NetWorker server. To add a client to NetWorker in offline mode, shut down NetWorker and start nsradmin with the `-d` option to connect directly to the NetWorker database (nsrdb)

```
[root@dr-nsr901 res]#  
[root@dr-nsr901 res]# nsradmin -d ./nsrdb  
NetWorker administration program.  
Use the "help" command for help, "visual" for full-screen mode.  
nsradmin> █
```

Run the create type: command to create the client definition offline. The example below creates the client with name win2012-01.vlab.local and client id of d9607099-00000004-57aca84a-57aca849-00085000-691c0156

```
nsradmin> create type: NSR Client; name: win2012-01.vlab.local; client id: d9607099-00000004-57aca84a-57aca849-00085000-691c0156  
          type: NSR Client;  
          client id: \  
d9607099-00000004-57aca84a-57aca849-00085000-691c0156;  
          name: win2012-01.vlab.local;  
Create? y  
created resource id 1.0.210.42.0.0.0.104.109.180.87.192.168.1.190(1)  
nsradmin> █
```

After the client is created, we can leave nsradmin and restart NetWorker.

Rebuild the CFI Entries

After the savesets are scanned in and the required client entries are created, we can do saveset restores. To perform file level restores, we must recreate the indices for the client. To recreate the index for a particular client, we can use `scanner -i -c <clientname> <devicename>`.

The example below shows the rebuilding of the index for `linclient1.vlab.local` from device `CloneDevice1`.

```
[root@dr-nsr901 res]#  
[root@dr-nsr901 res]# scanner -i -c linclient1.vlab.local CloneDevice1  
8909:scanner: using 'CloneDevice1' as the device name  
8936:scanner: scanning Data Domain disk ClonePool1.001 on CloneDevice1  
8939:scanner: Data Domain disk ClonePool1.001 already exists in the media index
```

Using `nsrck -L7` would only be possible if the DR NetWorker server uses the same hostname and client ID as the originating NetWorker server.

Recovering data in the DR data zone

We now have prerequisites in place to be able to recover data from a particular client to a client in the DR data zone. Before we can recover, we must create a pool to mount the device used for recovery. How to determine the originating pool name has been described earlier in this paper. Once the pool is created, we can mount and access the device.

Recovering of savesets and specific Files

To recover savesets, we specify the server to recover from, the saveset ID / clone ID, and the recover destination. A saveset recovery can be executed without the need to rebuild the client file index. Note that the client must be reachable to execute the recovery from the NetWorker Administration console. The example below shows the steps to execute a directed saveset recover of an imported client. The recovery is executed on the target machine. The screenshot below is an example of finding the saveset and clone ID.

```
[root@DR-LinClient1 recover]# mminfo -s dr-nsr901.vlab.local -q "client=linclient1.vlab.local,level=full" -r "ssid,cloneid,name,totalsize,savetime"
ssid      clone id name      total  date
4188843351 1470939486 /      12166258160 08/11/2016
1320224263 1471220290 /      12171420292 08/14/2016
4172066139 1470939486 /boot  24794652 08/11/2016
2997729087 1471003208 /boot  24794648 08/12/2016
1303447052 1471220290 /boot  24794504 08/14/2016
[root@DR-LinClient1 recover]#
```

Once we have the saveset and clone ID, we can execute the recovery.

```
[root@DR-LinClient1 recover]# recover -s dr-nsr901.vlab.local -d /recover -S 1303447052/1471220290
Recovering a subset of 34 files within / into /recover
Recover start time: Wed 17 Aug 2016 12:29:05 PM EDT
Requesting 1 recover session(s) from server.
129290:recover: Successfully established direct file retrieve session for save-set ID '1303447052' with Data Domain volume 'ClonePool11.001'.
```

Once the client file index has been created, we can recover specific files. Please be aware that command line must be used to recover files.

Client Migrations between NetWorker Servers

Another use case of this methodology is the migration of clients from one NetWorker data zone to another without losing the ability to perform file level recoveries. Since index rebuilds take a long time, it is recommended to use a migration clone device and to clone only the savesets which are absolutely necessary. A migration strategy could look like the following:

- Find all long term retention backups for a client.
- Clone these backups to a specific clone device. This device must have a unique name. The group name this devices belongs to must also be unique.
- Fast copy this clone device to the new data zone.
- Create the client using the same client ID on the new data zone.
- Scan the fast copied clone device into the new data zone.
- Rebuild the indices.

We can now set up the client to do backups on the new data zone. Any shorter term retentions can be recovered from the old data zone, while all new backups and long term retentions can be recovered from the new data zone.

Dell EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL EMC MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell EMC software described in this publication requires an applicable software license.

Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries.