

Using EMC NetWorker with Solaris SUNWndmp

David Stes
Molenstraat 5
2018 Antwerp, Flanders, Belgium
email: stes@pandora.be

November 11, 2009

Abstract

This document describes setting up EMC NetWorker with the Solaris SUNWndmp package for NDMP backups. The integration between EMC NetWorker and Solaris SUNWndmp offers the following benefits. There's no EMC NetWorker agent necessary for Solaris. The Solaris SUNWndmp daemon supports Solaris ZFS snapshots. The Solaris SUNWndmp package supports both LAN and LAN-free backups. Also, EMC NetWorker supports both browsable, i.e. individual file, and saveset restores for Solaris SUNWndmp. There's only one TCP port for Solaris SUNWndmp, the NDMP port 10000, so Solaris SUNWndmp is easy to firewall. The Solaris SUNWndmp package can be used with tape drives and autochangers (tape robots) to perform LAN-free backups. Also, we show how the Solaris SUNWndmp package supports NDMP version 4, so it is compatible with EMC NetWorker's new path-to-tape cloning feature, where non-NDMP savesets can be cloned via NDMP. Solaris SUNWndmp works with both EMC NetWorker 7.2.2 and EMC NetWorker 7.5.1 so it is not necessary to upgrade the EMC NetWorker server for Solaris SUNWndmp support. In fact, Solaris SUNWndmp provides a way to backup Solaris ZFS with EMC NetWorker 7.2.2.

1 EMC NetWorker backup server

As mentioned in the abstract, EMC NetWorker server supports NDMP, the Network Data Management Protocol, already for a long time. It provides a way to make a backup of a system without having to install an EMC NetWorker client on that system. If one wishes to make backups of Solaris ZFS using non-NDMP, regular NetWorker backups, it should be noted that the Solaris ZFS support in EMC NetWorker requires a client that is more recent than EMC NetWorker 7.2.2.

One interesting feature of EMC NetWorker is the DSA (Direct Save Access) support, which supports encapsulation of NDMP savesets in regular EMC NetWorker savesets. DSA NDMP is supported in EMC NetWorker 7.2 or higher. Since EMC NetWorker 7.2, EMC NetWorker has an integrated NDMP tapeserver, so NDMP backups can be sent simply over the LAN towards the backup server, and there is no need for a locally attached tape drive, attached to the NDMP server. With DSA NDMP, it is possible to send NDMP backups to disk, or to tape drives attached to the backup server, or to multiplex non-NDMP backups with NDMP backups.

Note that it is by no means necessary to run a Solaris backup server. In fact, we can also test with a Microsoft Windows or Linux backup server running EMC NetWorker 7.2.2, and this works fine as well.

For our testing purposes, we used an EMC NetWorker server, called *darkstar*. This server is a small Solaris 11 system. Here is the version information of *darkstar* :

```
# uname -a
SunOS darkstar 5.11 snv_125 i86pc i386 i86pc
```

On this Solaris server, we have a directly attached, physical tape unit :

```
# cat /etc/release
                                OpenSolaris Development snv_125 X86
Copyright 2009 Sun Microsystems, Inc.  All Rights Reserved.
Use is subject to license terms.
                                Assembled 12 October 2009

# cfgadm -avl c5
Ap_Id Receptacle Occupant Condition Information
When Type Busy Phys_Id
c5 connected configured unknown
unavailable scsi-bus n /devices/pci@0,0/pci1022,7450@a/pci9005,44@1:scsi
c5::rmt/0 connected configured unknown IBM ULTRIUM-TD3
unavailable tape n /devices/pci@0,0/pci1022,7450@a/pci9005,44@1:scsi::rmt/0
```

The EMC NetWorker packages, manpages, client, server software, are installed as usual :

```
# strings /usr/sbin/nsrd | grep '(#)'
@(#) Product:      NetWorker
@(#) Build number: 269
@(#) Build arch.:  sol10amd64
@(#) Release:     7.5.1.Build.269
@(#) Build date:  Fri Mar 20 23:05:02 PDT 2009
@(#) Build info:  DBG=0,OPT=-O2 -fno-strict-aliasing
```

2 Installing the Solaris SUNWndmp client

Normally, the Solaris SUNWndmp machine is a different server, called the NDMP server, but in our case, we will install the Solaris SUNWndmp packages on the same system *darkstar*.

It is by no means evident that installing EMC NetWorker server and Solaris SUNWndmp on the same machine will work, because EMC NetWorker uses resources called *NSR storage node* resources. These resources are created in the database of EMC NetWorker, and in the database of EMC NetWorker, it is recorded whether a machine is a SCSI storage node or an NDMP storage node (which implies that they are not both NDMP storage node and non-NDMP storage node).

Officially, for EMC NetWorker, a storage node cannot be a regular (non-NDMP) storage node and an NDMP storage node at the same time. So in order to avoid this problem, it is better to use different machines for the EMC NetWorker backup server and the Solaris SUNWndmp client, but the packages can be co-hosted on the same machine and it can be made to work.

The Solaris SUNWndmp packages are installed as follows:

```

# pkg install SUNWndmp
DOWNLOAD                                PKGS      FILES      XFER (MB)
Completed                                1/1       13/13      0.3/0.3

PHASE                                    ACTIONS
Install Phase                            40/40

```

There are two packages : the NDMP root components and the NDMP user components.

In addition there exists also a test program, called SUNWstc-ndmpcopy, but this is not used here (as we use EMC NetWorker to initiate the NDMP backups).

```

# pkginfo SUNWndmpr
system      SUNWndmpr Network Data Management Protocol Service (Root)

# pkginfo -l SUNWndmpu
  PKGINST:  SUNWndmpu
    NAME:   Network Data Management Protocol Service (User)
  CATEGORY: system
    ARCH:   i386
  VERSION:  11.11,REV=2009.09.30.21.42
    VENDOR: Sun Microsystems, Inc.
    DESC:   Network Data Management Protocol Service (user components)
  HOTLINE:  Please contact your local service provider
  STATUS:   completely installed

```

There exists older versions of these packages. For example, EMC NetWorker officially supports the SUN NAS 7000 series, and we believe that there are older versions of Solaris SUNWndmp shipped with those NAS systems.

```

# pkginfo -l | grep ndmp
PKGINST: SUNWndmpr
PKGINST: SUNWndmpu

# uname -a
SunOS sun7000 5.11 snv_111b i86pc i386 i86pc

# pkginfo -l SUNWndmpu | grep VERSION
VERSION: 11.11,REV=2009.04.08.17.26

```

3 Configuring the Solaris SUNWndmp client

Once the Solaris SUNWndmp packages are installed, we have to enable the NDMP service.

```
# svcadm enable ndmpd
```

After enabling the service, we can check that the daemon runs :

```
# ps -ef | grep ndmpd
root 3651      1    0 06:29:25 ?                0:00 /usr/lib/ndmp/ndmpd
```

This process listens (usually) on TCP port 10000 :

```
# netstat -an | grep 10000
*.10000          *.*              0          0 49152        0 LISTEN
```

There are some utilities included, such as *ndmpadm* and *ndmpstat*. For example, with *ndmpadm* we can see that the logfile is in */var/ndmp* :

```
# ndmpadm get
debug-path=/var/ndmp
dump-pathnode=no
tar-pathnode=no
ignore-ctime=no
token-maxseq=9
version=4
dar-support=no
tcp-port=10000
backup-quarantine=no
restore-quarantine=no
overwrite-quarantine=no
```

We can also see at the Solaris SUNWndmp level, that we have a tape unit to work with :

```
# ndmpadm show-devices
Tape drive(s):
Name       : /dev/rmt/0n
LUN #     : 0
SCSI ID #  : 0
Vendor    : IBM
Product   : ULTRIUM-TD3
Revision  : 64D0
Serial    : 1230146020
WWN       : IBM          ULTRIUM-TD3      1210
```

In order to use EMC NetWorker with Solaris SUNWndmp, we have to set a password.

If we try, for example, to use the EMC NetWorker command *inquire* without setting a password, we get an error :

```
# inquire -N localhost
Enter NDMP user name: ? ndmp
Communicating to devices on NDMP Server 'localhost', this may take a while...
42597:inquire: connect auth: connection has not been authorized
42610:inquire: The NDMP connection is not successfully authorized on host 'localhost'.
```

The username (ndmp) and password (ndmp) is set as follows :

```
# ndmpadm enable -a cram-md5 -u ndmp
Enter new password: ndmp
Re-enter password: ndmp
```

Once this is done, we can test with EMC NetWorker that the command *inquire* works :

```
# inquire -N localhost
Enter NDMP user name: ? ndmp
Enter 'ndmp' password on NDMP host 'localhost' (characters will not be echoed):

Communicating to devices on NDMP Server 'localhost', this may take a while...
scsidev@0.0.0:IBM      ULTRIUM-TD3      64D0|Tape, /dev/rmt/0n
                               S/N: 1230146020
                               ATNN=IBM      ULTRIUM-TD3      1210346020
```

Note that it is not required to have a locally attached tape unit, because EMC NetWorker supports DSA NDMP backups where the data can be sent via the LAN to the backup server.

4 Configuring an EMC NetWorker NDMP device

Because EMC NetWorker supports DSA NDMP, we do not really need an NDMP tape device in order to perform a backup. We could do the NDMP backup simply to disk or to a tape unit that supports both NDMP and non-NDMP backups.

In this case, we will do a LAN-free backup and set up a small disk device and the tape unit as NDMP device.

For the disk device, we label it in a pool called *INDEX*,

```
# zfs create -o mountpoint=/b2d rpool/b2d

# nsradmin
NetWorker administration program.
Use the "help" command for help, "visual" for full-screen mode.
nsradmin> create type: NSR device; name: /b2d; media type: file
                    type: NSR device;
                    name: /b2d;
                    media type: file;

# nsrmm -l -b INDEX -f /b2d DISK
```

We need this device for index and bootstrap backups, but normally, this is already dealt with if the backup server and the NDMP server are different hosts.

For EMC NetWorker 7.5.1, the name of the NDMP device must be in a specific format, including the word NDMP in parentheses. This was not the case in EMC NetWorker 7.2.2 where the device name could simply be something like */dev/rmt/0n*.

```

# nsradmin
NetWorker administration program.
Use the "help" command for help, "visual" for full-screen mode.
nsradmin> create type: NSR device; name: "rd=darkstar:/dev/rmt/0n (NDMP)";
  media type: LTO Ultrium-3; ndmp: Yes; target sessions: 1;
  remote user: ndmp; password: ndmp

```

Check that there exists a storage node resource called *darkstar* and check that the NDMP password is correctly set, not only at the device level, but also at the storage node level :

```

# nsradmin
NetWorker administration program.
Use the "help" command for help, "visual" for full-screen mode.
nsradmin> . type: NSR storage node; remote user: ndmp
nsradmin> update password: ndmp
                    password: ndmp;
Update? y

```

Once the storage node resource has the correct ndmp password set, label and mount the NDMP tape device as follows :

```

# nsrmm -l -f "rd=darkstar:/dev/rmt/0n (NDMP)" TAPE
Volume 'TAPE' is not recyclable - are you sure you want to over-write it with a new label?
# nsrmm -m -f "rd=darkstar:/dev/rmt/0n (NDMP)" TAPE
LTO Ultrium-3 tape TAPE mounted on rd=darkstar:/dev/rmt/0n (NDMP), write enabled

```

It is not possible to update the NDMP password at the device level :

```

# nsradmin
NetWorker administration program.
Use the "help" command for help, "visual" for full-screen mode.
nsradmin> . type: NSR device; remote user: ndmp
Current query set
nsradmin> update password: ndmp
Update? y
update failed: The password for NDMP device cannot be updated from NSR device resource.
Use NSR storage node resource to update this information.

```

During the label operation, which is requesting to write a label to the tape, via the Solaris SUNWndmp package, you may get an error in the NDMP logfile :

```

# tail -1 /var/ndmp/ndmplog.0
11/11 7:49:59 ERROR: [12][ndmpd_tape_read_v3:935]:Tape read error: I/O error.

```

However, after the label operation *nsrwatch* should show something like this :

Server: darkstar

Wed Nov 11 7:02:27 2009

```
Up since: Wed Nov 11 6:12:07 2009 Version: NetWorker 7.5.1.Build.269 Eval
Saves: 0 session(s) Recovers: 0 session(s)
Device          type          volume
/b2d            file          DISK          mounted file disk DISK
/rmt/0n (NDMP) LTO Ultrium-3 TAPE          mounted LTO Ultrium-3 tape TAPE
```

Also, we can inspect the EMC NetWorker processes and check that there are *nsrmmmd* processes for the two devices:

```
# ps -ef | grep nsr
root 3057 3051 0 16:12:12 ? 0:00 /usr/sbin/nsrjobd
root 1856 1 0 12:30:56 ? 0:01 nsrexecd
root 3056 3051 0 16:12:10 ? 0:00 /usr/sbin/nsrindexd
root 3082 3051 0 16:14:11 ? 0:00 /usr/sbin/nsrmmmd -n 1
root 3051 1 0 16:12:07 ? 0:00 nsrd
root 3139 3051 0 16:21:46 ? 0:00 /usr/sbin/nsrmmmd -n 2 -r darkstar
root 3055 3051 0 16:12:08 ? 0:00 /usr/sbin/nsrmmdbd
```

5 Configuring an EMC NetWorker NDMP jukebox

In the case of an autochanger (a tape library) containing a tape drive, there are multiple options, such as dynamic drive sharing between NDMP and non-NDMP machines, or control of the robot arm via NDMP or via ACSLS or SCSI directly attached to a EMC NetWorker server.

In the case of an NDMP jukebox, we attach the robot arm to a machine running Solaris SUNWndmp and we control the robot arm via NDMP commands.

Note that it is not necessary to use a NDMP jukebox in order to use a NDMP device; the NDMP tape drive can be in a non-NDMP jukebox.

Here's a jbconfig session for an Solaris SUNWndmp controlled jukebox :

```
# jbconfig
```

```
Jbconfig is running on host darkstar (SunOS 5.11),  
and is using darkstar as the NetWorker server.
```

- 1) Configure an AlphaStor Library.
- 2) Configure an Autodetected SCSI Jukebox.
- 3) Configure an Autodetected NDMP SCSI Jukebox.
- 4) Configure an SJI Jukebox.
- 5) Configure an STL Silo.

```
What kind of Jukebox are you configuring? [1] 3
```

```
Enter NDMP Tape Server name: ? darkstar
```

```
Enter NDMP user name: ? ndmp
```

```
Enter on NDMP password (characters will not be echoed):
```

```
Communicating to devices on NDMP Server 'darkstar', this may take a while...
```

```
14484:jbconfig: Scanning SCSI buses; this may take a while ...
```

```
Installing 'Standard SCSI Jukebox' jukebox - scsidev@0.0.0.
```

```
What name do you want to assign to this jukebox device? ndmpjb
```

```
15814:jbconfig: Attempting to detect serial numbers on the jukebox and drives ...
```

```
15815:jbconfig: Will try to use SCSI information returned by jukebox to configure drives
```

```
Turn NetWorker auto-cleaning on (yes / no) [yes]? no
```

```
The following drive(s) can be auto-configured in this jukebox:
```

- 1> LTO Ultrium-2 @ 0.0.0 ==> /dev/rmt/0n (NDMP)

```
These are all the drives that this jukebox has reported.
```

```
To change the drive model(s) or configure them as shared or NDMP drives,  
you need to bypass auto-configure. Bypass auto-configure? (yes / no) [no] no
```

```
Jukebox has been added successfully
```

```
The following configuration options have been set:
```

- > Jukebox description to the control port and model.
- > Autochanger control port to the port at which we found it.
- > Autocleaning off.
- > Barcode reading to on.
- > Volume labels that match the barcodes.

```
You can review and change the characteristics of the autochanger and its  
associated devices using the NetWorker Management Console.
```

```
Would you like to configure another jukebox? (yes/no) [no]no
```

```
The jbconfig has configured an NDMP tape drive in the NDMP jukebox :
```

```
# nsrmm
(nothing) mounted on LTO Ultrium-2 tape rd=darkstar:/dev/rmt/0n (NDMP)
```

We can label a tape as follows :

```
# nsrjb -L -S 1 -b Default
1:          ndmpjb [enabled]
There is only one enabled and configured jukebox: ndmpjb
Label 'LTO999' is a valid NetWorker label. Overwrite it with a new label? y
Info: Operation 'Verify label' in progress on device 'rd=darkstar:/dev/rmt/0n (NDMP)'
Info: Expected volume '-' in slot '1'. The actual volume is 'LTO999'.
```

The following daemons will run on the backup server :

```
# ps -eo args | grep nsr
/usr/sbin/nsrmmdbd
/usr/sbin/nsrindexd
/usr/sbin/nsrmmmd -n 1 -r darkstar
/usr/sbin/nsrexecd
/usr/sbin/nsrjobd
/usr/sbin/nsrd
/usr/sbin/nsrmmgd
/usr/sbin/nsrlcpd -s darkstar -N 1 -n 1
```

Note that the *nsrlcpd* keeps an NDMP session permanently open.

```
# ndmpadm show-sessions | tail -1
1 active sessions.
```

So it is normal to see an active NDMP session that remains active as long as the *nsrlcpd* daemon is running, even when no NDMP backups are running.

If the NDMP jukebox is disabled then the *nsrlcpd* process will go away, and the permanently open NDMP session will go away as well.

6 NDMP client resource

In order to perform a backup now, we have to setup an NDMP client resource. There are many options, but let us assume that we have created a group called *NDMP* and setup a client with save set *All* as a client in that group:

```

        group: Default    [NDMP];
        save set: All;
Backup renamed directories: Enabled    [Disabled];
        priority: 500;
File inactivity threshold: 0;
File inactivity alert threshold: 0;
        remote access: ;
        remote user: ndmp;
        password: ndmp;
        backup command: nsrndmp_save -T dump -c darkstar;
application information: HIST=n,DIRECT=n,EXCLUDE=/b2d,UPDATE=n;
        ndmp: [Yes]    No ;
NDMP array name: darkstar;

```

The client is called *darkstar* and we have configured it as a NDMP client resource (*ndmp* is set to *Yes*). Also there is a property called *NDMP array name* that relates to NDMP licensing that we have set to *darkstar*.

The backup command is:

```
backup command: nsrndmp_save -T dump -c darkstar;
```

In the case of a DSA NDMP backup, which is EMC NetWorker's powerful technology to do a LAN-based NDMP backup, the backup command is:

```
backup command: nsrndmp_save -M -T dump -c darkstar;
```

7 Performing a Solaris SUNWndmp probe

With the above settings, we can *probe* the Solaris SUNWndmp daemon and discover what we can backup.

```

# savegrp -pvI -l full NDMP

32451:savegrp: darkstar:All                                level=full
7236:savegrp: Group will not limit job parallelism
32493:savegrp: darkstar:probe                            started
savefs -s darkstar -c darkstar -g NDMP -p -l full -R -m -v
64690:savegrp: savegrp:NDMP * darkstar:All See the file /nsr/tmp/sg/NDMP/sso.darkstar.r
7340:savegrp: darkstar:probe succeeded.
7076:savegrp: --- Probe Summary ---

darkstar:All                                level=full, dn=-1, mx=0, vers=ssbrowse, p=12
darkstar:All                                level=full, pool=NDMP, save as of 11 November 2009 05:09:53 PM
darkstar:/                                  level=full, dn=0, mx=1, vers=ssbrowse, p=12
darkstar:/                                  level=full, pool=NDMP, save as of 11 November 2009 05:09:53 PM
darkstar:/rpool                              level=full, dn=4, mx=1, vers=ssbrowse, p=12
darkstar:/rpool                              level=full, pool=NDMP, save as of 11 November 2009 05:09:53 PM
darkstar:/b2d                                level=full, dn=5, mx=1, vers=ssbrowse, p=12
darkstar:/b2d                                level=full, pool=NDMP, save as of 11 November 2009 05:09:53 PM
7241:savegrp: nsrim run recently, skipping

```

See the file /nsr/tmp/sg/NDMP/sso.darkstar.rYaibh for output of save command, indicates that we have to look in that file to see what EMC NetWorker has found out by asking the Solaris SUNWndmp package.

```
7212:savefs: ***** FS INFO[1] *****
7213:savefs: fstype = zfs ; logical device = / ; physical device =
7215:savefs: total-size 73.88 GB
7217:savefs: used-size 4.478 GB
7219:savefs: Avail size 69.40 GB
7221:savefs: Total Indoos 145691402
7223:savefs: Used Indoos 144474
7224:savefs: LOCAL y
7224:savefs: TYPE zfs
7224:savefs: AVAILABLE_BACKUP tar,dump
7224:savefs: AVAILABLE_RECOVERY tar,dump
7225:savefs: Status of the file system =
7212:savefs: ***** FS INFO[2] *****
7213:savefs: fstype = zfs ; logical device = /rpool ; physical device =
7215:savefs: total-size 69.40 GB
7217:savefs: used-size 81.00 KB
7219:savefs: Avail size 69.40 GB
7221:savefs: Total Indoos 145546940
7223:savefs: Used Indoos 12
7224:savefs: LOCAL y
7224:savefs: TYPE zfs
7224:savefs: AVAILABLE_BACKUP tar,dump
7224:savefs: AVAILABLE_RECOVERY tar,dump
7225:savefs: Status of the file system =
7212:savefs: ***** FS INFO[3] *****
7213:savefs: fstype = zfs ; logical device = /b2d ; physical device =
7215:savefs: total-size 69.40 GB
7217:savefs: used-size 86.50 KB
7219:savefs: Avail size 69.40 GB
7221:savefs: Total Indoos 145546934
7223:savefs: Used Indoos 6
7224:savefs: LOCAL y
7224:savefs: TYPE zfs
7224:savefs: AVAILABLE_BACKUP tar,dump
7224:savefs: AVAILABLE_RECOVERY tar,dump
7225:savefs: Status of the file system =
7226:savefs: *****
savefs darkstar: succeeded.
```

The probe output tells us what the valid save set names are, and it also indicates that Solaris SUNWndmp supports both *dump* and *tar* as backup applications. However, these applications are builtins for the Solaris SUNWndmp package, and no actual dump or tar process is seen on the host during backups.

We will now make a *dump* of the save set root (/), which according to the probe is about 4.478 GB.

8 Performing a local Solaris SUNWndmp backup with HIST=n

Prior to the backup we check that there are no ZFS snapshots. The Solaris SUNWndmp package creates snapshots on the fly, as we immediately see.

```
# zfs list -t snapshot
no datasets available
```

Next, we launch the backup using the following command:

```
# savegrp -vI -l full NDMP
```

This will then print the following command-line :

```
64690:savegrp: savegrp:NDMP * darkstar:/ See the file /nsr/tmp/sg/NDMP/sso.darkstar.66a
32494:savegrp: darkstar:/ started
nsrndmp_save -T dump -c darkstar -s darkstar -g NDMP -l full -LL -W 78 -N / /
```

It is therefore possible to check (using *tail -f*) the progress of the NDMP backup by inspecting the `/nsr/tmp/sg/NDMP` files.

Alternatively, we can use the Solaris SUNWndmp tools to monitor the backup.

First of all, during the backup, Solaris SUNWndmp automatically creates snapshots.

```
/nsr/tmp/sg/NDMP# zfs list -t snapshot
NAME                               USED  AVAIL  REFER  MOUNTPOINT
rpool/ROOT/opensolaris@bk-ndmp.2  1.89M      -    4.48G  -
```

At the end of the backup, these snapshots are destroyed.

During the backup, we can see that there are NDMP sessions active:

```
# ndmpadm show-sessions | head
  Session Id:      10
  Protocol version: 4
  Authenticated:   Yes
  EOF:             No
  Client address:  172.16.0.8:44731
  scsi.open:       No
  scsi.adapter:    ""
  scsi.valid target: No
  tape.fd:         12
  tape.record count: 1504
```

Also, the Solaris SUNWndmp package contains a tool to monitor the backup :

```
# ndmpstat 1 2
wthr ops      file      disk      tape      bytes      perf      prcnt
r w  bk rs  rd  wr  rd  wr  rd  wr  rd  wr  bk rs  dsk tpe idl
0 0  0  0   0   0   0   0   0   0   0   0   0  0   0  0  100
1 0  1  0   80   0 26746  0   0 26880 26813  0  26  0   49  50  1
```

Note that during the backup, there is no actual dump process :

```
# ps -ef | grep ndmp
root    411      1    0 17:46:39 ?                0:27 /usr/lib/ndmp/ndmpd
root    859     585    0 17:51:04 ?                0:00 /usr/sbin/nsrndmp_save -T dump -c dar
root    919     851    0 17:54:27 pts/3           0:00 grep ndmp
```

```
# ps -ef | grep 411
root    411      1    0 17:46:39 ?                0:28 /usr/lib/ndmp/ndmpd
root    923     851    0 17:54:37 pts/3           0:00 grep 411
```

The prstat tool shows activity of the ndmpd daemon :

```
# prstat
PID USERNAME  SIZE  RSS STATE  PRI NICE      TIME  CPU PROCESS/NLWP
411 root        12M 9280K sleep  59  0    0:01:10  11% ndmpd/9
```

After a while we see that the snapshot is deleted, and when we inspect the EMC NetWorker logfile (under /nsr/tmp/sg/NDMP) we find the logfile of *nsrndmp_save* :

```
42785:nsrndmp_save: Performing backup to NDMP type of device
42897:nsrndmp_save: Opened the tape device : /dev/rmt/0n
42910:nsrndmp_save: Performing Non-DAR Backup..
42617:nsrndmp_save: NDMP Service Log: Direct Access Restore information is supported
42952:nsrndmp_save: darkstar:/ NDMP save running on 'darkstar'
42617:nsrndmp_save: NDMP Service Log: Backing up "/".

42617:nsrndmp_save: NDMP Service Log: Tape record size: 61440.

42617:nsrndmp_save: NDMP Service Log: File history: N.

42617:nsrndmp_save: NDMP Service Log: Date of the last level '0': the epoch.

42617:nsrndmp_save: NDMP Service Log: Date of this level '0': Tue Nov 11 18:07:25 2009.

42617:nsrndmp_save: NDMP Service Log: Update: FALSE.

42617:nsrndmp_save: NDMP Service Log: Runtime [/] 4828200960 bytes (4828200960): 1032 se

The NDMP backup is complete.
darkstar: / level=full, 4715 MB 00:17:27      0 file
completed savetime=1257959243
42913:nsrndmp_save: Save session closed with NW server successfully
```

We can check using the command `mminfo` that the save set is in the media database of EMC NetWorker :

```
# mminfo -av
volume      client      date      time      size ssid      fl  lvl name
DISK        darkstar   11/11/2009 06:24:53 PM 256 KB 4260995557 cr full bootstrap
TAPE        darkstar   11/11/2009 06:07:23 PM 4715 MB 4277771724 crN full /
```

The flag 'N' indicates that this is an NDMP backup. The flag 'Ns' stands for DSA NDMP backup (which also works fine, and which is very interesting, because it makes it unnecessary to use a local tape drive).

In this case , the average throughput was only 4 MB/s , so using a dedicated LTO-3 tape drive for this NDMP backup does not make any sense. A DSA NDMP backup (via the LAN), for example using 100 mbit or gigabit LAN, makes much more sense in this particular case.

In other cases, the internal disk storage of the Solaris SUNWndmp server can be very fast, and in such cases a LAN-free backup to a regular NDMP tape device makes sense.

9 Performing a local Solaris SUNWndmp backup with HIST=y

The EMC NetWorker index database can be updated if we change the application information in the client resource, and set HIST=y instead of HIST=n :

```
application information: HIST=y,DIRECT=n,EXCLUDE=/b2d,UPDATE=n;
```

When running a DSA NDMP backup with such settings, we obtain a *browsable* flag for the save set in the media database :

```
DISK      darkstar 11/11/09 15:21:02      7561 KB 16003278 cbNs full /testvol
```

The save set flags are : c (complete) b (browsable) and Ns (DSA NDMP).

Other modifications are possible as well. For example, we can use -T tar instead of -T dump.

```
backup command: nsrndmp_save -M -T tar
```

So far we have only considered full backups (-l full). We can also do level-backups of a different level (level 0 is a full backup). These level backups are and should be supported by the Solaris SUNWndmp package, as EMC NetWorker is not actually doing any backup (there is not even a EMC NetWorker agent active on the NDMP server, all the work is done by the Solaris SUNWndmp package).

Note that Solaris SUNWndmp also supports the UPDATE=y variable, but some configuration changes to the Solaris SUNWndmp are necessary such as:

```
# ndmpadm set -p dar-support=yes
# ndmpadm set -p tar-pathnode=yes
# ndmpadm set -p dump-pathnode=yes
```

For reference, we now include here example output of such a browsable Solaris SUNWndmp backup. This backup is used later on for the recover and cloning examples.

The backup is initiated via the *savegrp* command :

```
# savegrp -l full -m -I -v NDMP
32451:savegrp: darkstar:/TEST level=full
7236:savegrp: Group will not limit job parallelism
32493:savegrp: darkstar:probe started
savefs -s darkstar -c darkstar -g NDMP -p -l full -R -v -F /TEST
darkstar:/TEST level=full, dn=0, mx=1, vers=ssbrowse, p=12
64690:savegrp: savegrp:NDMP * darkstar:Probe See the file /nsr/tmp/sg/NDMP/sso.darkstar
7340:savegrp: darkstar:probe succeeded.
32494:savegrp: darkstar:/TEST started
nsrndmp_save -c darkstar -T dump -s darkstar -g NDMP -l full -LL -W 78 -N /TEST /TEST
64690:savegrp: savegrp:NDMP * darkstar:/TEST See the file /nsr/tmp/sg/NDMP/sso.darkstar
7341:savegrp: darkstar:/TEST succeeded.
32494:savegrp: darkstar:bootstrap started
save -s darkstar -S -g NDMP -LL -f - -m darkstar -V -l 0 -LL -W 78 -N bootstrap -x /nsr/
64690:savegrp: savegrp:NDMP * darkstar:bootstrap See the file /nsr/tmp/sg/NDMP/sso.dark
7341:savegrp: darkstar:bootstrap succeeded.
7241:savegrp: nsrim run recently, skipping
```

The output of the *nsrndmp_save* command is in the *sg* file :

```

42785:nsrndmp_save: Performing backup to NDMP type of device
42897:nsrndmp_save: Opened the tape device : /dev/rmt/0n
42617:nsrndmp_save: NDMP Service Log: Direct Access Restore information is supported
42952:nsrndmp_save: darkstar:/TEST NDMP save running on 'darkstar'
42617:nsrndmp_save: NDMP Service Log: Backing up "/TEST".

42617:nsrndmp_save: NDMP Service Log: Tape record size: 65536.

42617:nsrndmp_save: NDMP Service Log: File history: Y.

42617:nsrndmp_save: NDMP Service Log: Date of the last level '0': the epoch.

42617:nsrndmp_save: NDMP Service Log: Date of this level '0': Wed Nov 11 15:19:19 2009.

42617:nsrndmp_save: NDMP Service Log: Update: TRUE.

42617:nsrndmp_save: NDMP Service Log: Runtime [/TEST] 1096286208 bytes (1096286208): 70 seconds

The NDMP backup is complete.
darkstar: /TEST level=full, 1070 MB 00:02:15 1448 files
completed savetime=1257949108
42913:nsrndmp_save: Save session closed with NW server successfully

42914:nsrndmp_save: Sorting File History....
42916:nsrndmp_save: Sorting File History completed Successfully in 00:00:00 Hours

42917:nsrndmp_save: Processing NDMP File History...
42918:nsrndmp_save: darkstar:/TEST Processing NDMP File History completed Successfully on 'darkstar' in 00:00:00 Hours

42920:nsrndmp_save: browsable savetime=1257949108

```

10 Browsable Solaris SUNWndmp restores

Assume that we have a browsable Solaris SUNWndmp backup of a filesystem TEST :

```

# mminfo -av
LTO999      darkstar      11/11/2009 03:18:28 PM 1070 MB 16435138  cbN full /TEST

```

We can use the EMC NetWorker command *recover* as follows :

```

# recover -c darkstar
cd /TEST
add test
relocate darkstar::/RECOV
recover

```

As an example, here's the output, using Solaris SUNWndmp, of such a recover session :

```
# recover
Current working directory is /TEST/
recover> add test
/TEST
1 file(s) marked for recovery
recover> relocate /RECOV
recover> list
/TEST/test @ 11 November 2009 03:18:28 PM CET
1 file(s) marked for recovery
recover> recover
Recovering 1 file from /TEST/ into /RECOV
Volumes needed (all on-line):
    LTO999 at rd=darkstar:/dev/rmt/0n (NDMP)
Total estimated disk space needed for recover is 1048 MB.
Requesting 1 file(s), this may take a while...
42787:nsrndmp_recover: Performing recover from NDMP type of device
42897:nsrndmp_recover: Opened the tape device : /dev/rmt/0n
42690:nsrndmp_recover: Performing non-DAR Recovery..
42617:nsrndmp_recover: NDMP Service Log: DAR is disabled. Running Restore without DAR
42617:nsrndmp_recover: NDMP Service Log: Restoring to "/RECOV/test".

42617:nsrndmp_recover: NDMP Service Log: Tape server: local.

42617:nsrndmp_recover: NDMP Service Log: Tape record size: 65536.

42706:nsrndmp_recover: Retrieved file '/TEST/test' successfully.
42927:nsrndmp_recover: Successfully done
recover> quit
```

11 Saveset Solaris SUNWndmp restores

For non-browsable restores, we can restore the entire save set.

First we have to figure out the *save set id* or *ssid* of the save set:

```
# mminfo -av
LTO999    darkstar      11/11/2009 03:18:28 PM 1070 MB 16435138  cbN full /TEST
```

The save set id is 16435138 in the above case.

Next, start the Solaris SUNWndmp restore with *-v* off i.e., index verification disabled, as follows:

```
# nsrndmp_recover -c darkstar -m /RECOV -S 16435138 -v off /TEST
42787:nsrndmp_recover: Performing recover from NDMP type of device
42897:nsrndmp_recover: Opened the tape device : /dev/rmt/0n
42690:nsrndmp_recover: Performing non-DAR Recovery..
42617:nsrndmp_recover: NDMP Service Log: DAR is disabled. Running Restore without DAR
42617:nsrndmp_recover: NDMP Service Log: Restoring to "/RECOV/".

42617:nsrndmp_recover: NDMP Service Log: Tape server: local.

42617:nsrndmp_recover: NDMP Service Log: Tape record size: 65536.

42706:nsrndmp_recover: Retrieved file '/TEST' successfully.
42927:nsrndmp_recover: Successfully done
```

12 NDMP Cloning

Here's an example of regular NDMP cloning with EMC NetWorker and Solaris SUNWndmp.

Label a tape in the *Default Clone* pool :

```
# nsrmm -l -b "Default Clone" -f "rd=darkstar:/dev/rmt/1n (NDMP)" CLONE
```

We will clone the following saveset :

```
#mminfo -av
LTO999    darkstar    11/11/2009 03:18:28 PM 1070 MB 16435138  cbN full /TEST
```

For DSA NDMP savesets, we have to use *nsrclone*.

For regular NDMP savesets, we have to use *nsrndmp_clone* :

```

# nsrndmp_clone -v -S 16435138
5815:nsrndmp_clone: NDMP Cloning the following save sets (ssid/cloneid) :
-----
valid = 16433684
-----
ssid/cloneid = 16435138/1257949108
-----
5816:nsrndmp_clone:
Starting NDMP Tape cloning operation...
42798:nsrndmp_clone: Cloning the following save sets from the volume(valid) 16433684...
32379:nsrndmp_clone: Cloning save set /TEST (ssid/cloneid 16435138/1257949108)
72985:nsrndmp_clone: Source volume 'LTO999' ready for cloning
72986:nsrndmp_clone: Destination volume 'CLONE' ready for cloning
5783:nsrndmp_clone: Tape-to-tape duplication of size 1070 MB is in progress...
67140:nsrndmp_clone: Open the source tape device : /dev/rmt/0n
69700:nsrndmp_clone: record size is set to 64 KB
67141:nsrndmp_clone: Open the destination tape device : /dev/rmt/1n
69700:nsrndmp_clone: record size is set to 64 KB
42617:nsrndmp_clone: NDMP Service Log: End of tape reached. Load next tape
42744:nsrndmp_clone: Tape server paused: reached the end of file
Source Mover reached EOF
Total bytes cloned so far 1070 MB
5781:nsrndmp_clone: Cloning done for save set : /TEST (ssid 16435138)
67791:nsrndmp_clone: Size cloned for save set: /TEST (ssid 16435138) is 1070 MB
darkstar:/TEST size 1070 MB, 1448 file(s), took 3 min 56 sec
72979:nsrndmp_clone: Total size cloned so far: 1070 MB.
67790:nsrndmp_clone: Total size cloned: 1070 MB.
5820:nsrndmp_clone: Verifying Cloned Save sets with Media Database.....
5884:nsrndmp_clone: Successfully cloned all requested save sets
5886:nsrndmp_clone: Clones were written to the following volume(s):
CLONE

```

We obtain like this two clones of this NDMP saveset :

```

#mminfo -avr ssid,cloneid,sumflags,savetime,volume -qssid=16435138
  ssid      clone id fl  date    volume
16435138   1257950526 cbN 11/11/2009 CLONE
16435138   1257949108 cbN 11/11/2009 LTO999

```

13 Tape Copy

Note that NDMP cloning, or EMC NetWorker cloning in general, is not to be confused with tape copy or tape mirroring.

Here's a sample output of a *tcopy* session where we use the Solaris command to make an identical copy of the source tape instead of an EMC NetWorker clone.

This means that the EMC NetWorker label of the tape is copied as well (the two tapes will have the same volume identifier *valid*, which will confuse EMC NetWorker if we mount the tapes at the same time).

```
#tcopy /dev/rmt/0cbn /dev/rmt/1cbn
file 1: record 1: size 32768
file 1: eof after 1 records: 32768 bytes
file 2: record 1: size 32768
file 2: eof after 1 records: 32768 bytes
file 3: record 1: size 65536
file 3: eof after 1 records: 65536 bytes
file 4: records 1 to 344: size 65536
file 4: eof after 344 records: 22544384 bytes
file 5: record 1: size 65536
file 5: eof after 1 records: 65536 bytes
file 6: record 1: size 65536
file 6: eof after 1 records: 65536 bytes
file 7: records 1 to 16728: size 65536
file 7: eof after 16728 records: 1096286208 bytes
file 8: record 1: size 65536
file 8: eof after 1 records: 65536 bytes
eot
total length: 1119158272 bytes
```

The tape that is obtained by this command, can be mounted via EMC NetWorker, but it is not a EMC NetWorker clone. Specifically, it will not be in a *clone pool*, it will just be impossible to distinguish between the original and it will be a volume from a regular backup pool.

```
#nsrmm -m -f "rd=darkstar:/dev/rmt/1n (NDMP)"
LTO Ultrium-3 tape LTO999 mounted on rd=darkstar:/dev/rmt/1n (NDMP), write enabled
```

14 Path-to-tape cloning

EMC NetWorker 7.5 and higher supports NDMP *path-to-tape* cloning of non-NDMP savesets.

This functionality is enabled by the *-p* flag of the *nsrndmp_clone* command.

First, we make a normal backup of the file */etc/motd*, to a normal (non-NDMP) tape drive.

```
# mminfo -av
LTO999 darkstar 11/11/2009 05:50:52 PM 6 KB 4278807612 cb manual /etc/motd
```

This backup has level set to *manual* because it was a manual (client-initiated) backup. The volume LTO999 was mounted on a local tape drive that was configured as a normal EMC NetWorker SCSI-attached tape drive.

Next, we reconfigure the machine to use NDMP tape drives (without deleting the volume LTO999 that contains our non-NDMP backup).

```
(nothing) mounted on LTO Ultrium-3 tape rd=darkstar:/dev/rmt/1n (NDMP)
(nothing) mounted on LTO Ultrium-2 tape rd=darkstar:/dev/rmt/0n (NDMP)
```

Now we clone the non-NDMP saveset by using the Solaris SUNWndmp package :

```

# nsrndmp_clone -p -J darkstar -v -S 4278807612
5815:nsrndmp_clone: NDMP Cloning the following save sets (ssid/cloneid) :
-----
volid = 617237
-----
ssid/cloneid = 4278807612/1258908652
-----
5816:nsrndmp_clone:
Starting NDMP Tape cloning operation...
42798:nsrndmp_clone: Cloning the following save sets from the volume(volid) 6172
37...
32379:nsrndmp_clone: Cloning save set /etc/motd (ssid/cloneid 4278807612/1258908
652)
72985:nsrndmp_clone: Source volume 'LTO999' ready for cloning
69699:nsrndmp_clone: Cloning from one media type to another with different block
sizes.
72986:nsrndmp_clone: Destination volume 'TAPE' ready for cloning
5783:nsrndmp_clone: Tape-to-tape duplication of size 6 KB is in progress...
67140:nsrndmp_clone: Open the source tape device : /dev/rmt/0n
68527:nsrndmp_clone: tape file size is set to 2097 MB
69700:nsrndmp_clone: record size is set to 64 KB
67141:nsrndmp_clone: Open the destination tape device : /dev/rmt/1n
68527:nsrndmp_clone: tape file size is set to 2097 MB
69700:nsrndmp_clone: record size is set to 64 KB
42617:nsrndmp_clone: NDMP Service Log: End of tape reached. Load next tape
42744:nsrndmp_clone: Tape server paused: reached the end of file
Source Mover reached EOF
No more fragments to clone
Total bytes cloned so far 128 KB
5781:nsrndmp_clone: Cloning done for save set : /etc/motd (ssid 4278807612)
67791:nsrndmp_clone: Size cloned for save set: /etc/motd (ssid 4278807612) is 12
8 KB
darkstar:/etc/motd size 128 KB, 3 file(s), took 0 min 13 sec
72979:nsrndmp_clone: Total size cloned so far: 128 KB.
67790:nsrndmp_clone: Total size cloned: 128 KB.
5820:nsrndmp_clone: Verifying Cloned Save sets with Media Database.....
5884:nsrndmp_clone: Successfully cloned all requested save sets
5886:nsrndmp_clone: Clones were written to the following volume(s):
TAPE

```

The clone volume can be mounted in a non-NDMP tape drive and it is possible to restore from the clone, using a regular non-NDMP restore :

```

# recover
Recovering 1 file into its original location
Requesting 1 file(s), this may take a while...
Received 1 file(s) from NSR server 'darkstar'
Recover completion time: 11 November 2009 06:15:09 PM CET
Recovering 1 file into its original location
Total estimated disk space needed for recover is 1 KB.
Requesting 1 file(s), this may take a while...
Requesting 1 recover session(s) from server.
./motd
Renaming ./motd to ./motd.R
Received 1 file(s) from NSR server 'darkstar'
Recover completion time: 11 November 2009 06:15:16 PM CET

```

To test the path-to-tape functionality the NDMP and non-NDMP drives can be shared (using dynamic-drive-sharing or DDS) between multiple Solaris machines that access the tape drives via a SAN switch.

In our case, we just one machine, and we reconfigured the tape drives, first using them as non-NDMP drives, then next, for the cloning, as NDMP drives, and finally for the restore test again as a non-NDMP device.

The application of this technology (called *path-to-tape*) is to use NDMP for VTL (virtual tape libraries) where the cloning can be performed by a VTL where the actual volumes are just containing regular, non-NDMP backups.

15 Some issues encountered so far

During our test configuration, we have encountered so far two issues.

When doing NDMP restores, the NDMP logfile ndmplog.0 has printed some information such as

```
Could not set correct file permission for file /home/stes/.local/share/codeina/mozembed/
```

However, the restore finished successfully and Solaris SUNWndmp restored everything (without obviously setting the file permission on that particular lock file).

Inspection of this file shows:

```

# cd ~/.local/share/codeina/mozembed
# ls -l lock
lrwxrwxrwx 1 stes staff 14 2009-07-17 08:46 lock -> 127.0.0.1:+612

```

This type of error message is perfectly normal, and of course one should carefully inspect the ndmplog.0 file for error messages related to Solaris SUNWndmp backups and restores.

This issue is very much a minor issue, in view of the fact that the Solaris SUNWndmp restore was successful.

Another issue, which is unfortunately, more important, is that in our experience so far, the Solaris SUNWndmp package and EMC NetWorker work together, but are not yet very stable.

At this point, with the versions of Solaris SUNWndmp and EMC NetWorker as indicated, we sometimes see the backup work correctly and at the end we get the error :

```
42613:nsrndmp_save: Receive NDMP shutdown message.
42725:nsrndmp_save: Failed to estimate the size of backup data being processed.
42889:nsrndmp_save: NDMP data server host is in the process of shutdown.
42890:nsrndmp_save: NDMP tape server host is in the process of shutdown.
42891:nsrndmp_save: The entire backup is aborted.
42725:nsrndmp_save: Failed to estimate the size of backup data being processed.
42641:nsrndmp_save: NDMP write failed, communication failure
```

The error sometimes happens with Solaris SUNWndmp, and a similar message was reported with the Legato Celestra NDMP module when performing backups of systems using a vxfs managed root disk. Similar bugs are known to happen with NetApp as well, specifically when the NetApp is using snapshots. Information from the NetApp bug (36554) about using Snapmirror and DUMP at the same time, reports that some buffers may be flushed. These might coincide with the use of Snapmirror, as the Snapmirror process aborted around the same time as the DUMP failed.

So this error can be seen in some other NDMP environments as well, and Solaris SUNWndmp is possibly difficult to troubleshoot, and NDMP certainly has its own problems.

Further research is needed, possibly by using Solaris SUNWndmp on SPARC based systems, because our testing so far was on Intel and AMD machines.

16 Conclusion

One straightforward (and good) way to backup Solaris systems, is by installing the EMC NetWorker client and doing non-NDMP (regular) EMC NetWorker backups.

We have seen that there is an alternative, the integration between EMC NetWorker and Solaris SUNWndmp offers the following benefits.

There's no EMC NetWorker agent necessary for Solaris. The Solaris SUNWndmp daemon supports Solaris ZFS snapshots. The Solaris SUNWndmp package supports both LAN and LAN-free backups. Also, EMC NetWorker supports both browsable, i.e. individual file, and saveset restores for Solaris SUNWndmp. There's only one TCP port for Solaris SUNWndmp, the NDMP port 10000, so Solaris SUNWndmp is easy to firewall. Finally, Solaris SUNWndmp works with both EMC NetWorker 7.2.2 and EMC NetWorker 7.5.1 so it is not necessary to upgrade the EMC NetWorker server for Solaris SUNWndmp support.

In fact, Solaris SUNWndmp provides a way to backup Solaris ZFS with EMC NetWorker 7.2.2.