



# NDMP Plugin

**Bacula Systems Documentation**

---

# Contents

1	Scope	2
2	Features	2
3	Backup Strategies	3
4	Installation	5
5	Configuration	8
6	Operations	18
7	Best Practices	20
8	Limitations	20

# Contents

---

This chapter aims at presenting the reader with information about the **Bacula Enterprise Network Data Management Protocol (NDMP) Plugin**. The document briefly defines the scope of its operations, describes the target technology of the plugin, and presents its main features and various techniques and strategies to backup NAS filers using the NDMP or NFS/CIFS protocols with Bacula Enterprise.

## 1 Scope

The current version of the NDMP Plugin supports several NAS vendors, like NetApp, HPE, EMC, and Huawei, among others. With NetApp's DataONTAP, we tested versions 8.x and 9.x.

---

**Note:** Vendor implementations of NDMP differ, so we cannot guarantee that all NDMP implementations will work with our NDMP Plugin.

---

## 2 Features

The following chapter describes the features of NDMP Plugin:

- File to Server mode
- ACL support
- File Restore via the Single Item Restore technique
- NDMP Dump/Tar
- NetApp SMTAPE
- Full, Incremental and Differential support for NDMP Dump/Tar backup

- Full for NDMP SMTAPE backup
- Full, Incremental and Differential support for NDMP SMTAPE backup NetApp
- NetAppCABConfiguration
- Copy and Migration jobs

---

#### Available since Bacula Enterprise 12.4.0

The Bacula Enterprise NDMP Plugin can analyse the native NetApp Dump format, the EMC Unity DUMP format, and the universal (i.e. supported by many different vendors) Tar format to ensure a high deduplication ratio with Bacula Systems' ged. For other NDMP vendors, Aligned Volumes may give good results with the appropriate storage system.

---

## 3 Backup Strategies

This article aims at presenting possible backup strategies for the NDMP Plugin.

### 3.1 Backing Up Using NDMP

NDMP has the main advantage of speed over the NFS/CIFS protocols for backing up NAS filers. Note that NDMP is a control protocol, not a backup format. The NDMP backup format is arbitrary (implementation-specific) and defined individually by each storage platform vendor. Thus in general, an NDMP backup made on one vendor's NAS system cannot be restored to a different vendor's NAS.

NDMP can be used in 3 ways:

- **Filer to Server:** Filer backing up across the LAN to your Bacula server.
- **Filer to Self:** Filer backing up to an attached tape drive.
- **Filer to Filer:** Filer backing up across the network to another filer's tape drive.

At this time, the **Bacula Enterprise** NDMP Plugin supports only the **Filer to Server** mode. In this mode, **Bacula** encapsulates NDMP backup images inside the standard Bacula backup data stream, allowing it to be multiplexed with other backup streams and providing capabilities such as arbitrary (Bacula) storage device use, compression, checksum verification, etc. Inside the saved data stream that contains the NDMP stream, the actual NDMP data remains opaque (vendor specific format, sometimes proprietary).

When using NDMP provided by the NAS vendor, attributes such as ACLs are included in the backup stream and handled correctly by the NAS vendor on backup and restore. Determining which files should be saved is also part of the NAS vendor's job, and Bacula has no way to ensure that all files will be protected with technologies such as our Accurate Mode.

Since Bacula does not implement the NDMP Direct Access Recovery (DAR) protocol (which is reputed to be very slow), doing a single file restore is possible via the Single Item Restore technique (`use_hist` option).

Restoring a single file with the NDMP Plugin without the Single Item Restore is possible by restoring the NDMP datastream to a local machine (same one that has the NDMP Plugin), then opening the resulting file. This requires to know the vendor data format and to have a program that can read it. Since it involves restoring the whole backup and then extracting the desired files, it isn't very efficient.

An alternative, simple, and fast technique is to make use of file system snapshots made on the NAS to keep a history of file system states. These snapshots remain on the NAS box and can thus easily be accessed to restore individual files. In this manner, the NDMP backups are used as a disaster recovery mechanism only, rather than a means of restoring individual files.

In short, the best use of the NDMP Plugin is to protect the NAS from disaster situations, while using snapshots on the NAS box allows to recover files quickly on a daily basis.

---

**Note:** It is also possible to backup a particular snapshot with the NDMP Plugin.

---

## 3.2 Using NetApp SMTAPE Backup Option

The SMTAPE backup option essentially makes a raw binary dump of the selected NAS partitions (volumes). The big advantage of the SMTAPE feature is the speed, especially for volumes with millions of files on them. The standard backup option (DUMP) is much too slow to effectively backup large, multi-terabyte volumes, which are the norm these days. For Full backups, SMTAPE will significantly outperform DUMP on speed. Also, because it includes the entire snapshot history, users can reduce the number of snapshots they keep on their primary volume since the backups include them.

For example, keeping 30 days worth of snapshots on your filer and doing monthly full backups with SMTAPE, and keeping those for years would give you access to your entire snapshot history for the time period without taking up excessive space on your primary volumes. With large capacity tapes, that turns into a pretty granular data protection setup. If you had hourly snapshots, you could potentially get anything back from any hour going back many years just by doing a full backup every month. The recovery path is a bit complex but manageable.

---

### Available since Bacula Enterprise 12.4.0

Recent versions of SMTAPE now support the Incremental backup feature.

---

There are pros and cons to both approaches, but SMTAPE as an NDMP backup method solves a particularly painful problem for many enterprise users. It is a useful feature, and as long as NetApp is supporting its use by backup applications, it should be considered as first choice backup method.

See the section about SMTAPE Incremental/Differential configuration `SMTAPEConfiguration` for more information.

## 3.3 Backing Up Using NFS/CIFS without NDMP

Using NFS or CIFS to mount NAS volumes on a File Daemon machine allows to back up NAS files without the NDMP Plugin the same way it would be done with a local filesystem backup. However, when dealing with millions of small files, directory walking and network latency tend to decrease the backup throughput dramatically. By using techniques such as “Incremental forever” and Virtual Full backups that save only new and modified files, you can, however, achieve reasonable storage space consumption in this situation.

The Virtual Full feature reads the latest Full backup and all the Incrementals and merges them to create a new Full backup that looks like a “real” Full backup, but without having to access the NAS device. This kind of backup is sometimes referred to as a synthetic Full. During a Virtual Full backup, the data will be copied directly from one **Bacula** Volume to another, without accessing the NAS. See `VirtualFullJobs` for more information on this subject.

Most NAS systems provide file system snapshot capabilities, and some of them also provide capabilities to avoid the file system walk to find modified or new files, which is a function typically built on top of the snapshot comparison. If those features are accessible, it is possible to perform a file level backups with both better consistency and less overhead using “Incremental Accelerator”. Bacula Systems offers such functionality for NetApp NAS systems and the `nutanix_hfc` plugins.

Using the NFS or CIFS protocols to do a complete restore of a large NAS file system can be excessively time consuming especially in urgent disaster recovery situations.

In addition, Access Control Lists (ACLs) are not always handled correctly using NFS/CIFS. Unix POSIX ACLs can be handled directly by **Bacula**, but backing up Windows ACLs needs special techniques (see the SavingACLs chapter).

### 3.4 Strategies Comparison

Table 1: Backup Strategies Comparison

Features	NFS Mount	NDMP dump/tar	NDMP SMTAPE
Incremental backup	Yes	Yes	No 3
Accurate support	Yes	Yes 1	Yes 1
ACL	No/Yes 1	Yes	Yes
Exclusion/Inclusion	Yes	Yes	No
Cross-filer recovery	Yes	No 1	No 1

Backup Side	NFS Mount	NDMP dump/tar	NDMP SMTAPE
Network load	Medium	Low	High
Backup size	Small	Small	Large 1
Average backup speed	Slow	Fast	Fast 1
Speed with many files	Slow	Fast	Very Fast

Restore Side	NFS Mount	NDMP dump/tar	NDMP SMTAPE
Network load	Low/Medium	High	High
Single File Restore	Very Fast	Slow	Slow
File Level Restore	Yes	Yes 2	No
Disaster Recovery	No	Yes	Yes
NetApp Cluster Support	Yes	Yes	Yes

1 Platform or configuration dependent

2 With use\_hist option

3 Incremental and Differential backup support with NetApp only. See SMTAPEConfiguration for more information.

## 4 Installation

The NDMP Plugin is installed along the FileDaemon. To avoid network load between the FD and the SD, please install it on the same system of a SD.

It is possible to install the NDMP Plugin on any other host, but doing so is less efficient because it will cause the data to traverse two network links – once from the NAS to the File Daemon, and a second time from the File Daemon to the Storage Daemon.

If you plan to use the SMTAPE Incremental feature, it is required to install the NetAp OnTap API on your system. See NDMPNetAppInstallation for more information.

## 4.1 Prerequisites

The **Plugin Directory** option in the **File Daemon** resource should point to where the `ndmp-fd.so` plugin is installed, which usually is `/opt/bacula/plugins`

```
FileDaemon {  
    Name = bacula-fd  
    Plugin Directory = /opt/bacula/plugins  
    ...  
}
```

The File Daemon should have direct network access to the NAS, which can be verified using `telnet`, for example:

```
# telnet nasbox 10000  
Connected to nasbox.  
Escape character is '^]'.
```

**Note:** NDMP, when doing the actual backup, connects back from the filer to the host initiating the NDMP session. This may require firewalls to be configured accordingly. Using the **data\_port\_range** option, it is possible to control which TCP ports the NAS will attempt to use, allowing for more restricted firewall configuration. When the network connection is not possible, the error `NDMP4_CONNECT_ERR` can be issued.

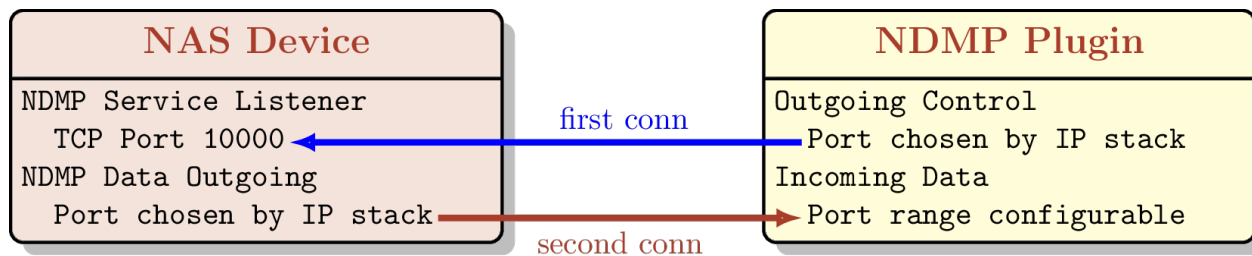


Fig. 1: Network Connections

The job specific configuration of the NDMP Plugin is part of the Fileset that is used and provided as options of the plugin directive.

Depending on NAS setup, the authentication method can be configured to be **md5** or **text**.

## 4.2 NDMP Plugin Installation

This article describes how to install Bacula Enterprise NDMP Plugin.

### NDMP Installation with BIM

In order to install the NDMP Plugin with BIM, install the File Daemon with BIM and choose to install the NDMP Plugin during the FD installation.

[Click here](#) for more details on the plugin installation process with BIM.

### Installation with Package Manager

Installation of the Bacula Enterprise NDMP Plugin is most easily done by adding the repository file suitable for the existing subscription and the distributions package manager configuration.

#### Deb based Linux distribution

An example would be `/etc/apt/sources.list.d/bacula.list` for deb based Linux distributions with the following content:

```
#Bacula Enterprise
[...]
deb https://www.baculasystems.com/dl/@customer-string@/debs/ndmp/@version@/stretch-64/
↪stretch ndmp
```

After that, a run of `apt-get update` is needed. Then, the plugin can be installed using `apt-get install bacula-enterprise-ndmp`

#### RPMS based Linux distribution

On RHEL or alternative distributions, extend the repository file for your package manager to contain a section for the plugin - `/etc/yum.repos.d/bacula.repo`:

```
[Bacula]
name=Bacula Enterprise
baseurl=https://www.baculasystems.com/dl/@customer@/rpms/bin/@version@/rhel7-64/
enabled=1
protect=0
gpgcheck=0

[BaculaNDMPPlugin]
name=Bacula Enterprise NDMP Plugin
baseurl=https://www.baculasystems.com/dl/@customer@/rpms/ndmp/@version@/rhel7-64/
enabled=1
protect=0
gpgcheck=0
```

Then, perform a `yum update` or `dnf update` and after that the package can be installed with `yum install bacula-enterprise-ndmp` or `dnf install bacula-enterprise-ndmp`.

Manual installation of the packages, can be done after downloading the proper files from the Bacula Systems provided download area, and then using the low-level package manager (`rpm` or `dpkg`) to perform the plugin installation.

### 4.3 NetApp OnTAP Python API Installation

To install the NetApp OnTAP Python API, run the following commands as root:

```
# pip3 install netapp-ontap --target /opt/bacula/python
```

Unfortunately, the information needed to access the OnTAP API is usually different from the NDMP information, so it requires to specify a few new parameters on the plugin command line such as:

- `ontap_user`
- `ontap_password`
- `ontap_host`
- `ontap_vserver`
- `ontap_ignore_ssl`
- `ontap_profile`

To test the connectivity with the OnTAP API interface of your NetApp system, you may use the following command to list the volumes on the vservers:

```
# /opt/bacula/bin/snapmgr.py -e "vserver" -u "user" -f "host" -p "password" -x -V  
vol1  
svm1_root  
backup_data1_1
```

If you have an error such as `401 Client Error: Unauthorized for url`, you may have to use another interface to query the OnTAP data.

## 5 Configuration

The NDMP Plugin configuration is done on the **client side** (File Daemon) and takes place in the **`bacula-fd.conf`** file.

### 5.1 NDMP Plugin Configuration

The following chapter presents the information on NDMP Plugin configuration.

#### Automatic Objects Integration

Since Bacula version 16.0.7, a new solution has been introduced, so that each object can be backed up separately with different Jobs to maximize the throughput and the resiliency. It is highly recommended to use this new solution for that purpose - Automatic Object Integration (Scan Plugin). See an example for NDMP.

## Custom Volume Format

Starting with Bacula Enterprise 10.2.4, the volume format is automatically detected from the NDMP system. The `volume_format` option should not be required anymore.

The NDMP Plugin should be aware of the volume structure in order to detect if the administrator wants to restore to a new volume (`where=/dev/vol_tmp`) or inside a subdirectory of the targeted volume (`where=/tmp`).

The NDMP Plugin will automatically detect the following structure:

- `/vol/` (used on NetApp)
- `/root_vdm/` (used on some EMC hardware)

If your volumes use a different naming scheme, you **must** use the `volume_format` option.

```
/dev/volume_home      -> volume_format=/dev/  
/rootvolume/volume_tmp -> volume_format=/rootvolume/  
/VG/volume_var        -> volume_format=/VG/
```

The actual volume format and the allowed flexibility is vendor specific. In general, if volumes can be created with plain names and no path in the NAS configuration frontend, a fixed volume format will be in use.

```
Fileset {  
  Name = NDMPFS  
  ...  
  Include {  
    Plugin = "ndmp:host=nasbox user=root pass=root file=/dev/vol1 volume_format=/dev/"  
  }  
}
```

For example, on NetApp, the volume format is:

```
/vol/vol1  
/vol/vol2  
...
```

The common part to distinguish volumes and directories is `/vol/`. On Solaris with ZFS, volumes are stored in `/dev/`.

```
/dev/Volume01  
/dev/Volume02  
/dev/Volume03
```

If volumes are organized like in this example:

```
/FS1/  
/FS2/  
/FS3/
```

The `volume_format` should be configured as `/FS` – note that there is **no trailing slash** here!

## Fileset Examples

This section will present various File Set examples.

```
Fileset {
  Name = NDMP
  Include {
    Plugin = "ndmp:host=nasbox user=root pass=root volume=/vol/vol1"
  }
}
```

```
Fileset {
  Name = NDMP_home
  Include {
    Plugin="ndmp:host=nas user=root pass=root volume=/vol/vol2 file=/home"
  }
}
```

```
Fileset {
  Name = NDMP_home
  Include {
    Plugin="ndmp:host=nas user=root pass=root volume=/vol/vol2 type=SMTAPE"
  }
}
```

```
Fileset {
  Name = NDMP_snapshot
  Include {
    Plugin="ndmp:host=nas user=root pass=root volume=/vol/vol2/.snapshot/snap1"
  }
}
```

```
Fileset {
  Name = NDMP_Incremental4ever
  Include {
    Plugin="ndmp:profile=netapp volume=/vol/vol2 use_base_date"
  }
}
```

```
Fileset {
  Name = NDMP_SMTAPE
  Include {
    # The plugin line should be on a single line
    Plugin = "ndmp:host=nasbox user=root pass=root volume=/svm1/vol2
              ontap_user=admin ontap_password=pass
              ontap_host=ontapsrv ontap_vserver=svm1 ontap_ignore_ssl"
  }
}
```

## Using ndmp.conf to Store Username and Password

If exposing the access credentials of the NAS on the Plugin command line is undesirable, it is possible to use file `ndmp.conf`, stored on the File Daemon host, to save your credentials.

In the following example, the profile `root` will refer to the NAS host `nasbox`, the connection will use MD5 as authentication method (“/m”), the username will be `root` and the password will be `password`.

```
# cat /opt/bacula/etc/ndmp.conf
[--root]
-D nasbox/m,root,password
```

The format of the **-D** parameter is the following:

```
HOST[:PORT] [/FLAGS] [, USERNAME, PASSWORD]
```

Where:

Option	Description
HOST	Is the host name or IP address of the NAS
:PORT	Optional port number. If not given the port number is 10000.
USERN	A user name that will be recognized by the NAS. Whether this is a general user name or a special account within the NAS is implementation dependent.
PASSW	The password corresponding to the USERNAME. The password field should not contain special characters such as: “, ’, !, \$, / or #.
/	Optional flags to indicate desired NDMP version or authentication method. The default version is negotiated to be the highest possible.
FLAGS	The default authentication method is text (NDMP_AUTH_TEXT).
2	Use NDMP version 2.
3	Use NDMP version 3.
4	Use NDMP version 4.
n	Use no authentication (NDMP_AUTH_NONE).
t	Use text authentication (NDMP_AUTH_TEXT). The user name and password are sent over the network as clear text (unencrypted).
m	Use MD5 challenge/response authentication (NDMP_AUTH_MD5). The remote NAS is asked for a challenge. The password is used as the shared secret and is never sent over the network.

In the following example, the profile “11” will refer to the NAS host “10.1.1.11”, the connection will use MD5 as authentication method (“/m”), the username will be `root` and the password will be `mypassword`. A second profile “12” is also shown in this example `ndmp.conf` file.

```
[root@lxbbackup ~]# cat /opt/bacula/etc/ndmp.conf
[--11]
-D 10.1.1.11/m,root,mypassword

[--12]
-D 10.1.1.12/m,root,mypassword2
```

In the following example two **Fileset** resources are configured – one for each of the two NAS profiles configured in the `ndmp.conf` example above.

```
Fileset {
  Name = NDMP_11
  Include {
    Plugin = "ndmp: host=10.1.1.11 profile=11 volume=/vol/LAN_backup type=smtape"
  }
}

Fileset {
  Name = NDMP_12
  Include {
    Plugin = "ndmp: host=10.1.1.12 profile=12 volume=/vol/DMZ_backup type=smtape"
  }
}
```

---

**Note:** The `ndmp.conf` file is different from the `ontap.conf` file described here [ontapconf](#).

---

## SMTAPE Incremental/Differential Configuration

When using SMTAPE with Incremental or Differential backup level, the NDMP Plugin has to manage snapshots via the NetApp OnTAP API. The REST API used by the NDMP Plugin is available on NetApp 9.6 and later.

## Saving Access Control Lists (ACLs)

When using the NDMP Plugin, all metadata, including ACLs, should be backed up and restored correctly, as this functionality is provided by vendor built-in tools. Thus, problems in this area are in scope of vendor support and not due to particular backup or restore software.

## ACLs Using CIFS

If a CIFS volume is mounted on Linux system, it may not be possible to backup ACLs. In order to reliably back up ACLs on CIFS volumes these must be mounted on a Windows server with a **Bacula** File Daemon installed. If a large amount of data is involved, avoiding to send the data over the network twice may be achieved by installing the Windows server in a KVM host on the Storage Daemon server. By doing so, it will be possible to backup Windows ACLs and reduce the network load.

This virtual server should be started only during the NAS backup. Contact Bacula Systems for more information on this configuration.

Newer Linux distributions with recent Samba tools may actually allow to back up and restore security-related file metadata, though this is not fully supported by Bacula at this time. Please contact **Bacula Systems** support for more information.

In particular for file systems which are accessed both by Windows systems via CIFS and by NFS, and which are storing file ACLs for both access methods, a non-NDMP backup preserving all ACL metadata is often impossible, or would cause considerable resource consumption.

## ACLs Using NFS

To be able to backup file Access Control Lists (ACLs) through NFS mounts, ACL support needs to be enabled with the mount command as follows:

```
# mount -t nfs -o soft,intr,proto=tcp,acl nas:/vol/vol1 /nas/vol/vol1
```

In addition, it is needed to add the **ACL Support** option to the Bacula File Set used:

```
Fileset {  
  Name = FS_NAS_VOL1  
  Options {  
    ACL Support = yes  
    ...  
  }  
  File = /nas/vol/vol1  
}
```

## Plugin Options

Table 2: NDMP Plugin Options

Option	Re- quired	De- fault	Info	Example
host	Yes		NAS Hostname	host=192.168.0.1
user	Yes		Username	user=root
pass	Yes		Password	pass=password
port	No	10000	Connection Port	port=10010
auth	No	md5	md5, text, none	auth=text
profile	No		Profile name	profile=root
volume	Yes		Volume to backup	volume=/vol/vol0
vol- ume_format	No		Specific Volume format	volume_format=/vol/
file	No	/	Directory to backup	file=/home
timeout	No	5mins	Timeout for various NDMP related commands	timeout=5h
type	No	dump	dump, tar, SMTAPE, config 1	type=dump
data_port_range	No		Data port range	data_port_range=2000-3000
abort_on_error	No		Abort job on error	abort_on_error
debug	No		Enable debug	debug
use_cab <sup>2</sup>	No		Enable the NetApp Cluster Aware Backup Protocol	use_cab
use_hist	No		Store HIST information in the catalog - need to adjust timeout	use_hist
limit_hist	No		Limit the path depth of the NDMP history stored in the catalog	limit_hist=3
use_base_date <sup>3</sup>	No		Enable Incremental forever using BASE_DATE variable	use_base_date
hist_retention	No		Prune HIST files stored in working/ndmp after the specified time	hist_retention=30days
max_level <sup>3</sup>	No	9	Maximum number of DUMP level allowed by the NDMP vendor	max_level=32
on- tap_profile <sup>4</sup>	No		ontap.conf reference	ontap_profile=name1
ontap_host	No		OnTAP hostname	ontap_host=netapp.lan
ontap_user	No		OnTAP username	ontap_user=admin
on- tap_password	No		OnTAP password	ontap_password=xxx
on- tap_vserver	No		OnTAP vserver name	ontap_vserver=svm1
on- tap_use_https	No		Use OnTAP HTTPS access	ontap_use_https

1 The `config` option is deprecated since Bacula Enterprise 10.2.4

2 `use_cab` option is available as of Bacula Enterprise 16.0.2

3 `use_base_date` and `max_level` options are available as of Bacula Enterprise 16.0

4 OnTAP options are available as of Bacula Enterprise 12.2.4

Table 3: NDMP Plugin Restore Options

Option	Required	Default	Info	Example
dest_volume	No		Destination volume to restore to	dest_volume=/vol/vol1

## 5.2 NetApp Configuration

### OnTAP Configuration

The OnTAP parameters can be stored in `/opt/bacula/etc/ontap.conf` and be referenced with `ontap_profile` parameter on the plugin command line.

```
# cat /opt/bacula/etc/ontap.conf
[name1]
ontap_user=admin
ontap_password=MyPassword
ontap_vserver=svm1
ontap_ignore_ssl=True
ontap_host=10.0.2.3
```

```
Plugin = "ndmp:ontap_profile=name1 volume=/svm1/v1 host=netapp user=root pass=xxx.
↪ type=smtape"
```

### NetApp Cluster Aware Backup (CAB) Configuration

With NetApp NDMP systems, it is possible to configure the Bacula Enterprise NDMP Plugin to contact the node hosting a volume for the NDMP session via the NDMP CAB Extension.

```
Fileset {
    Name = "ndmp-CAB"
    Include {
        Plugin = "ndmp: user=x password=y host=svm1 volume=/volume1 use_cab"
    }
}
```

### NetApp Cluster Backup Using OnTAP API

With NetApp NDMP systems, it is also possible to configure the Bacula Enterprise NDMP Plugin to contact the node hosting a volume for the NDMP session via the OnTap API. The CAB extension is the recommended configuration to use.

---

**Note:** Available since Bacula Enterprise 16.0.1

---

For this configuration, it is required to configure the following:

- Install the python OnTap API via `pip3 install netapp-ontap --target=/opt/bacula/python`
- Create and configure the `/opt/bacula/etc/ontap.conf` file with OnTAP API related information
- Create and configure the `/opt/bacula/etc/ndmp.conf` file with NDMP related information

- One NDMP LIF per cluster node (role data)
- Use the DUMP method

In the following example, the OnTAP configuration should allow Bacula to query the OnTAP api.

```
# cat /opt/bacula/etc/ontap.conf
[svm1]
ontap_host=netapp-ontap-api.lan
ontap_user=bacula
ontap_password=AVeryGoodPassword?
ontap_vserver=svm1
ontap_ignore_ssl=True
```

In the following ndmp configuration file, the NDMP protocol must be configured on one of the node of the cluster.

```
# cat /opt/bacula/etc/ndmp.conf
[--svm1]
-D netapp-cluster1.lan/m,bacula,Xlkdfkdd
```

The snapmgr.py tool can be used to test the NetApp cluster configuration with the following command:

```
# /opt/bacula/bin/snapmgr.py -o svm1 -M
volume=/svm1/volume1 aggr=ONTAP97_02_FC_1
volume=/svm1/svm1_root node=ONTAP97-01 aggr=ONTAP97_01_FC_1 ip=10.0.110.93
```

The final Fileset plugin command line should look like the following:

```
Fileset {
    Name = "ndmp-ONTAP"
    Include {
        Plugin = "ndmp: profile=svm1 ontap_profile=svm1 host=svm1 volume=/svm1/volume1"
    }
}
```

Some FlexGroup volumes may cause issues with the detection, it is possible to assign statically the node to connect to in the ontap.conf file.

```
# cat /opt/bacula/etc/ontap.conf
[svm1]
ontap_user = bacula
. . .

[svm1.flexgroup]
volume1 = 1.2.3.4
volume2 = 1.2.3.4
```

## 5.3 NAS Volume Configuration and Testing

Starting with Bacula Enterprise 10.2.4, the configuration will be automatically stored with the dump file. It is no longer required to specify this plugin command in the Fileset.

By specifying **type=config** as a NDMP plugin argument, Bacula will backup a configuration overview of all the NAS volumes to the file `/@ndmp/<host>.config`. To display the configuration, this file can be restored, and Bacula will display the NAS volume configuration in the Job log as shown in the example:

```
Plugin = "ndmp:host=nasbox user=root pass=pass type=config"
```

```
...
JobId 3: QR  File system /vol/vol2
JobId 3: QR    physdev
JobId 3: QR    unsupported 0x0
JobId 3: QR    type        WAFL
JobId 3: QR    status      online
JobId 3: QR    space       805306 total, 118784 used, 8052948 avail
JobId 3: QR    inodes      288238 total, 100 used
JobId 3: QR    empty default env
JobId 3: QR
JobId 3: QR  File system /vol/vol1
JobId 3: QR    physdev
JobId 3: QR    unsupported 0x0
JobId 3: QR    type        WAFL
JobId 3: QR    status      online
JobId 3: QR    space       8589938 total, 6594560 used, 858334 avail
JobId 3: QR    inodes      307450 total, 1703 used
JobId 3: QR    empty default env
JobId 3: QR
...
```

The NetApp master configuration file is located by default in `/vol/vol0/etc`. If the complete NAS setup should be restorable from scratch, the directory `/vol/vol0/etc` must be included in the backup Fileset.

The **type=config** option can also be used to check the Network configuration and the connection information between the Bacula Enterprise NDMP Plugin and the NAS. If this backup job is not successful, check the configuration, in particular network and firewall related, and the access password.

## 6 Operations

The following chapter describes details regarding NDMP Plugin operations.

### 6.1 Backup

Backups created by the NDMP Plugin can be freely copied or migrated using the respective Job types: Copy Job and Migration Job of **Bacula Enterprise** as of version 10.

### 6.2 Restore

#### File Level Recovery

Using snapshot options available directly on the NAS allows to restore any file to any point in time represented by a snapshot. Since snapshots have little to do with data protection, using NDMP with native format or NetApp SMTAPE option can be used to protect your NAS system from a disaster situation. Combining both techniques provides data protection with some flexibility.

On NDMP systems that support DAR, the plugin option `use_hist` can be used to enable the “Single Item Restore” feature. With the `use_hist` option, all the filenames included in a `dump` are stored in the Bacula catalog, and a `restore` session will have the ability to choose individual files or directories to restore. Only the selected files are restored on the NDMP destination system. The correct `dump` file is automatically selected during the file selection process.

#### Restoring Using NFS/CIFS

Restoring through NFS/CIFS can be done as usual with `bconsole`, `BAT` or `BWeb`.

#### Restoring Using NDMP

To perform a full restore of an NDMP protocol backup, select one of the `*.dump`, `*.tar`, or `*.smtape` files in the `/@ndmp/` virtual directory. It is possible to change the `where` parameter to restore to an alternate location. However, the `RegexWhere` option is not available with this method.

The `where` parameter may refer to a directory inside the original volume, or a directory inside an alternate volume. The **volume\_format** option permits to configure how the NDMP Plugin will distinguish volumes and directories (see `customvolumeformat` for more information).

In a case that there is a need to explicitly define an alternate volume to restore to, the `dest_volume` plugin option can be used. This option can be used by modifying the plugin arguments at restore time.

Below, additional preparations for SMTAPE data stream restores are described.

In the following example, if the `dest_volume` is, during the restore, set to `/vol/vol1`, the NDMP stream will be restored to the directory `/restored` in volume `vol1`.

```
* restore where=/restored
```

In the following example, when `dest_volume` is not set, during the restore, the NDMP stream will be restored to the original volume, into the directory `/tmp`.

```
* restore where=/tmp
```

When restoring an NDMP data stream, it is possible that the filer device can not apply all metadata, in particular ACLs. Such situations need to be resolved after the restore itself, usually by checking and possibly adjusting permissions, ownership and ACLs of restored files through a system mounting the restored file system. In those situations, it might be important to observe if permissions or ownership are inherited through directory hierarchies.

## Restoring Selected Files from NDMP Streams

If the `use_hist` option has been applied with NDMP based backups, information about individual files will be added to Bacula's Catalog. To restore selected files, not complete volumes, you would navigate into the file system tree below the `@ndmp/` directory, pick individual files as usual, and finalize the restore preparation in the usual way, too. It is required to restore to an NDMP Plugin in this case, and the `where=` option will still be pointing to a directory inside the target volume. The volume to restore to would also be indicated by the plugin option `dest_volume`.

---

**Note:** Such single file restores will still read the complete NDMP data stream and send it to the filer device, which is responsible to apply the file selection during restoration.

---

Note that **SMTAPE** format data streams do not provide single file restore capabilities.

## Restoring Files without NDMP Host

In some cases, it might be necessary to extract files generated by the NDMP host by first restoring the NAS backup to a NDMP Plugin host machine, then using programs such as `restore` or `tar` on the restored data stream. For that purpose, select the respective NDMP data stream (not the individual files, in case `use_hist` is applied) and the `where=` option with the following format, i.e. with a greater-than sign directly followed by the absolute target path on the FD host:

```
* restore where=">/tmp"
```

NDMP files would then be extracted to the **local** `/tmp/` directory.

---

**Note:** The quotes in the above example must be used only when using command line argument, they are not needed when editing options via the restore menu.

---

## Restoring with NetApp SMTAPE

The SMTAPE option using the path to the root of the volume (e.g., `/vol/foo`) will back up and restore complete volumes. This is closer to backing up and restoring a hard disk partition using `dd` or `partimage`. Bacula has no knowledge of the internal structure of the SMTAPE data because, when the volume is brought back online after the restore, the filer will simply recognize it as a WAFL filesystem and proceed on its own.

To restore a NetApp volume with the SMTAPE option, the target volume must have the type Data Protection (DP) and needs to be restricted, which can be done with the `volume` command on the NetApp host:

```
netapp> volume create -volume restore -aggregate agg1_data -size 5GB
  -state online -policy default -type DP -autosize-mode grow_shrink
  -snapshot-policy none -foreground true

netapp> volume restrict -volume restore
Volume "restore" is now restricted.
```

Then, it is needed to specify `where=/vol/restore` in the restore command.

**Warning:** With this method, any data on the volume would be overwritten by the previously saved data.

After restore is complete, the volume is in a read-only state. To make the volume writeable, you need to either submit the `snapmirror break` command or use the Data ONTAP APIs.

## 7 Best Practices

- While it is technically possible to specify multiple volumes in a single **Fileset** resource configuration (by using multiple plugin command lines), this is not necessarily the best way to perform multiple volume backups. It is strongly recommended that one **Fileset** definition is created per each filer volume being backed up.
- By default, if one of your volume fails to back up in a “multi-volume” backup job, the main Bacula job will terminate “Backup OK – with warnings.” The JobStatus for jobs that terminate “Backup OK” and “Backup OK – with warnings” are not differentiated in the Catalog. They are both ‘T’, so this means that you will have to carefully monitor your backup job logs in case some volume backups fail and pay attention to the JobErrors field in the job summaries.
- To address this issue, there is a plugin option called “abort\_on\_error”, which causes Bacula to immediately fail the job as soon as an error is detected while backing up a volume. However, if you use this option, and the backup of volume number 5 in a list of 10 volumes fails, then the whole job will be failed, and volumes 6-10 will not be backed up during that job’s run.
- A 1:1 configuration (one volume backed up per job) means that the “abort\_on\_error” option will make more sense to enable in each job, so you will immediately know when a volume fails to backup since the Bacula job will terminate with a “Backup failed” message and ‘f’ in the Catalog for the job.
- With a 1:1 volume/job configuration, re-running a specific volume backup job is simple to do after the cause of the failure is investigated and fixed.
- In the example about the 10 volumes, without a 1:1 configuration, there is no way to re-run a backup of just the one volume that failed to back up.

All the above best practice are applied when using NDMPAutomaticObjectsIntegration.

## 8 Limitations

- The NDMP Plugin uses the job name to track the NDMP backup LEVEL information. It is not possible to mix multiple Filesets with a single job and the same NAS host and the same volumes.
- The NDMP password should not contain special characters such as: “, ’, !, \$, / or #
- With Single Item Restore, the number of files to be restored in a single session should not be too big. Some filers may not react very well with very large file lists.
- The Single Item Restore feature is not able to look into the Bacula backup stream. If a single file is selected inside a 1 TB NDMP dump file, 1TB will be read from Bacula’s source volumes. The DIRECT variable cannot be set to “Yes”.
- The NDMP dump does not provide information about deleted files. Bacula’s Accurate Mode is not supported.
- With the Single Item Restore feature enabled, the NDMP HIST file needs to be analyzed in memory at the end of the NDMP backup session on the File Daemon system. The memory needed for this operation depends on the

number of files and the tree structure. On some advanced situation, it is possible to configure a custom NDMP “HIST file” scanner by creating a symlink pointing to /opt/bacula/bin/custom\_ndmp\_idx\_dump.

- With the Single Item Restore feature enabled, all the files and directories included in a NDMP dump are stored in the Bacula Catalog. To build the restore tree with the “restore” bconsole command, the Director needs to allocate an average of 170B per file in memory.
- The history file (HIST) is stored after each backup session in the working directory. This file is needed to perform a Single Item Restore session. If the file is not present in the directory, the restore may not be successful. The history file is included automatically in the backup. It is possible to restore the history file associated with a given NDMP dump prior to the file restore. When selected, the history file is automatically restored to the working directory.
- The Single Item Restore feature has been tested on NetApp and EMC ISILON systems with 50,000 files.
- The Single Item Restore feature has not been tested intensively with accented characters.
- Permission on directories that are automatically created during a restore with the Single Item feature may not be correct.
- The HIST option is supported only with the directory or node file history format. To correctly set this option and verify that it is enabled on EMC Isilon devices, the following commands in the Isilon’s CLI interface may be issued:

```
NAS1-4# isi ndmp settings variables modify /ifs/path/to/specific/share HIST D
```

```
NAS1-4# isi ndmp settings variables list
```

Path	Name	Value
/ifs/path/to/specific/share	HIST	D

- It is not possible to restore individual files from a dump during the local restore (with where=">/path/" option).
- The replace option is not respected when using the Single Item Restore feature.
- For NDMP Single Item Restores in bconsole, if you wish to use Enter a list of files to restore option (option 7), or the file= option on the bconsole “restore” command line, a file with the list of filenames to restore must also include the names of the NDMP dump that holds the file(s) that you wish to restore.
- This plugin does not support backups made using the NDMP Tar format provided by EMC Unity systems.
- The native **dump** format often has a serious limitation of eight for the number of Incremental Jobs that you can run between two Full or Differential backups. Making more than eight Incrementals will not allow you to restore your system accurately. All files will be correctly restored, but subsequent restores of Incremental backups will not recreate the file system as it appeared during the final incremental backup.

With some NAS vendors such as NetApp, it is possible to set more incremental levels using the `max_level` plugin command parameter (up to 32 with NetApp), or use the `BASE_DATE` incremental forever backup strategy with the `use_base_date` plugin command parameter.

- The NDMP Plugin is not compatible with the `sparse` Fileset option.
- Backups created using the NDMP Plugin are not compatible with Virtual Full jobs. Do not attempt to combine these two backup strategies as you will not be able to properly restore NDMP Plugin jobs from Virtual Full backups.
- The `restart` command has limitations with plugins, as it initiates the Job from scratch rather than continuing it. Bacula determines whether a Job is restarted or continued, but using the `restart` command will result in a new Job.