



Bacula Enterprise Management

Bacula Systems Documentation

Contents

1	Daily Management	2
2	Logs Management	5
3	Jobs Monitoring	7
4	Storage Space Management	25
5	PostgreSQL Catalog Administration	38
6	Tape/Volume Management	40
7	Tape Autochanger Usage	56

Contents

The following article aims at presenting various information on how to manage Bacula Enterprise.

As a central part of your system, Bacula services will report critical issues with network, disk space, or Catalog database in its Job logs.

It is important to do some system monitoring to allow for correct backup operations:

- monitor storage space
- monitor Job results
- monitor your infrastructure

You need to monitor space because Bacula cannot predict whether there will be enough room on tape or disks. You must ensure the available space is adequate for storing Jobs, but also on the Director as full filesystems will prohibit the backup from running correctly. For example, the spooling data is by default in `/opt/bacula/working`, the same for the spooled attributes. If this file system is full, Jobs will not succeed. You might want to use a tool like [Nagios](#) in order to create alarms if the pool space or partition space for `/opt` is low. Also remember that filesystems filled close to capacity are prone to have performance issues.

Read more:

1 Daily Management

The following chapter provides details on jobs.

1.1 Job Levels, Types and Statuses

The following chapter presents descriptions on job levels, job types and job statuses.

Job Levels

The table below describes several levels for a job.

Level	Description
Backup levels	
F	Full backup: All files
I	Incremental: Files modified since last backup
D	Differential: Files modified since last full backup
S	Since: Not used
F	Virtual full backup
Verification levels	
C	Verify from Catalog
V	Verify: Init database
O	Verify volume to Catalog entries
d	Verify disk attributes to Catalog
A	Verify data on volume
Others	
B	Base level job
—	None: for Restore and Admin

See also:

Go to:

- [Job Types](#)
- [Job Statuses](#)

Go back to the [Job Levels, Types and Statuses](#) chapter.

Go back to the [Daily Management](#) chapter.

Go back to the main [Bacula Enterprise](#) chapter.

Job Types

The table below describes several types for a job.

Type	Description
B	Backup Job
V	Verify Job
R	Restore Job
D	Admin job
C	Copy of a Job
c	Copy Job
M	A previous backup job that was migrated
g	Migration Job
A	Archive Job
S	Scan Job
U	Console program
I	Internal system “job”

See also:

Go back to:

- [Job Levels](#)

Go to:

- [Job Statuses](#)

Go back to the [Job Levels, Types and Statuses](#) chapter.

Go back to the [Daily Management](#) chapter.

Go back to the main [Bacula Enterprise](#) chapter.

Job Statuses

The table below describes several statuses for a job.

Status	Description
A	Job canceled by user
B	Job blocked
C	Job created but not yet running
D	Verify differences
E	Job terminated in error
F	Job waiting on File daemon
I	Incomplete Job
L	Committing data (last despool)
M	Job waiting for Mount
R	Job running
S	Job waiting on the Storage daemon
T	Job terminated normally
W	Job terminated normally with warnings
a	SD despooling attributes
c	Waiting for Client resource
d	Waiting for maximum jobs

I – Incomplete Jobs

An incomplete Job is a Backup Job which has been ended before completion in a situation where the Storage Daemon was able and allowed (per configuration) to commit data to Volumes and to the Catalog that allows the Job to be resumed later on. This usually happens when there is a communication problem between SD and FD, or when the File Daemon fails for some reason.

Note that backups done from dynamically created data sets, for example database dumps or ephemeral file system snapshots, may not be restartable despite the Storage Daemon considering the data to be sufficient for this purpose.

See the console commands `restart` and `stop`.

See also:

Go back to:

- [Job Levels](#)
- [Job Types](#)

Go back to the [Job Levels, Types and Statuses](#) chapter.

Go back to the [Daily Management](#) chapter.

Go back to the main [Bacula Enterprise](#) chapter.

Go back to the [Daily Management](#) chapter.

Go back to the main [Bacula Enterprise](#) chapter.

See also:

Go to:

- [Logs Management](#)
- [Jobs Monitoring](#)
- [Storage Space Management](#)
- [PostgreSQL Catalog Administration](#)
- [Tape/Volume Management](#)
- [Tape Autochanger Usage](#)

Go back to the [main Bacula Enterprise Management page](#).

2 Logs Management

The following chapter provides information on how to manage logs.

2.1 Message Format Explanation

The messages generated by any Bacula component are all structured the same way:

- A message starts with an optional time stamp. The time stamp can be shown or hidden depending on command or options used. TODO details and links to list/llist, -td
- The name of the component generating the message follows. To be able to distinguish different message sources, convention is to use suffixes of *-fd*, *-sd*, *-dir* for File Daemon, Storage Daemon, and Director names, respectively.
- The next field of the message is the Job Id that caused the message. If messages are created that are not related to a Job, they get the Job Id 0 (zero) assigned. Such messages can refer to routine activity, such as volume pruning triggered by `status` commands, but they can also indicate noteworthy events, such as failing authentication due to an on-going brute force attack. Proper message resources will ensure that important messages are forwarded to responsible administrators and routine messages do not clutter logging facilities.
- The messages text follows. A messages Text can contain arbitrary characters (in particular, line breaks). Messages are intended for human readers, and a certain expertise with both Bacula and the systems involved is assumed.

In recent versions of the Bacula software, a structured message characterisation has been introduced, which consists of a severity level keyword and a coded identifier at the beginning of a messages text. An example:

```
2022-11-14 18:18:10 bsys-demo-sd JobId 1984: Warning: [SW0201] Out of freespace.
↳caused End of Volume "Vol-0045" at 245 on device "DiskAutochanger_Dev0" (/opt/
↳bacula/archive). Write of 64512 bytes got 3851.
```

The text will remain complete and human readable, but the severity indication `Warning` and the message cause identifier `SW0201` make automated processing much easier.

The identifier will always consist of two letters and four digits:

- The first letter will be D, S or F for **Director**, **Storage Daemon** or **File Daemon**, respectively.
- The second letter will indicate the severity level:

Character	Severity Level
A	Abort
F	Fatal
E	Error
W	Warning
S	Security
I	Informational
D	Debug
O	Ok, i.e. normal completion

Note that this list is a bit less granular than the message types documented in Messages Resource.

- The numerical identifier of the message will never change, even if the text will be adapted or translated.

There is no particular structure in assignment of the numbers for each message, so filtering or processing by ranges is not suitable. Also, message id numbers will not be allocated sequentially.

Currently, not all messages have been converted to the new format, and this will only be completed after some more releases of the software.

Go back to the [Logs Management](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

See also:

Go back to:

- *Daily Management*

Go to:

- *Jobs Monitoring*
- *Storage Space Management*
- *PostgreSQL Catalog Administration*
- *Tape/Volume Management*
- *Tape Autochanger Usage*

Go back to the *main Bacula Enterprise Management page*.

3 Jobs Monitoring

Job monitoring is helpful when the Backup Administrator needs to decide what to do if a Job fails or if a warning occurs, as well as in sorting and prioritizing backup issues to resolve them in accordance with your environment. Joblogs and Bacula log file can be parsed by a monitoring software in order to re-run a job or take any action you may consider necessary (email/scripts/snapshots/reschedule the Job) if an error or a warning occurs.

3.1 Example Output

When you parse joblogs you can see:

```
Non-fatal FD errors:    0
SD Errors:              0
```

Non-fatal FD errors: 0 AND SD Errors: 0 Here everything went smoothly, no issues to report, probably no actions should be taken.

In the next case, you can see 1 Non-fatal FD errors and a “Backup OK – with warnings” termination:

```
Non-fatal FD errors:    1
SD Errors:              0
FD termination status: OK
SD termination status: OK
Termination:           Backup OK -- with warnings
```

Human investigation is advised for this job.

It should help to have a look at the joblog (`l1ist joblog jobid=<jobid>`). For the above Job, we might see the cause of the warnings, for example:

```
time: 2016-12-16 12:24:24
logtext: debianserver-fd JobId 18:      Could not stat "/home/bacula": ERR=No such file_
↳or directory
```

Here we can see the `/home/bacula` directory specified in the job FileSet is not on the Client’s filesystem to be backed up, as requested by the backup Job’s FileSet.

Note: It is important to regularly check **all** of your joblogs to find potential problems and in case of error, update your error parsing scripts/software to take them into account and take appropriate action(s). A Job is correctly done when Backup Termination is “*Backup OK*” and Non-fatal FD/SD errors are equal to 0.

3.2 Third Party Scripts

Since Bacula started out as an open-source project, and this project and its surrounding community are still thriving, there are many third party scripts that can also help in the day to day monitoring of your Bacula Enterprise environment.

One such example of a third party script is called *baculabackupreport.py*.

This is a Python script that will send an HTML email report of the Jobs that have been run in the last 24 hours (default). When configured to run in a cron job, this script has many features to help you to have a clear, concise overview of your Bacula environment on a daily basis.

The latest version of this script may be found on Github here: <https://github.com/waa>

Subject: **Bacula - 14 jobs in the past 24 hours: 0 bad, 0 with errors, for all clients, all jobs, all job types, and all job statuses (2 jobs queued/running)** 15:31
 To: Bill Arlofski

Job ID	Job Name	Client	Status	Errors	Type	Level	Files	Bytes	Start Time	End Time	Run Time
47928	Speedy	—	Running	0	Backup	Inc	—	—	2022-06-01 15:24:39	Still Running	—
47927	Copy-Speedy (Speedy)	—	Running	0	Copy Ctrl:	—	—	—	2022-06-01 15:24:39	Still Running	—
47926	Catalog	speedy-fd	OK	0	Copy of 47919	Full	1,171	992,928,597	2022-06-01 02:45:24	2022-06-01 02:45:40	0:00:16
47925	Copy-Catalog (Catalog)	—	OK	0	Copy Ctrl: 47919 -> 47926	—	—	—	2022-06-01 15:24:18	2022-06-01 15:24:38	0:00:20
47924	Speedy-Etc	speedy-fd	OK	0	Backup	Full	1,911	7,299,561	2022-06-01 15:23:49	2022-06-01 15:23:53	0:00:04
47923	Verify_Catalog (Catalog)	speedy-fd	OK	0	Verify of 47919	Data	1,171	2,263,313,645	2022-06-01 02:46:11	2022-06-01 02:46:17	0:00:06
47922	Verify_Catalog (Catalog)	speedy-fd	OK	0	Verify of 47919	VD2C	1,171	2,263,484,065	2022-06-01 02:45:53	2022-06-01 02:46:09	0:00:16
47921	Verify_Catalog (Catalog)	speedy-fd	OK	0	Verify of 47919	W2C	1,171	0	2022-06-01 02:45:48	2022-06-01 02:45:51	0:00:03
47920	RestoreCatalog	speedy-fd	OK	0	Restore	—	1,171	2,263,313,645	2022-06-01 02:45:43	2022-06-01 02:45:45	0:00:02
47919	Catalog	speedy-fd	OK	0	Backup Copied to 47926	Full	1,171	992,772,086	2022-06-01 02:45:24	2022-06-01 02:45:40	0:00:16
47918	SpeedyVMs	speedy-fd	OK	0	Backup	Inc	10	23,260,116,071	2022-05-31 23:00:02	2022-05-31 23:10:10	0:10:08
47917	Speedy-Etc	speedy-fd	OK	0	Backup	Inc	7	149,356	2022-05-31 23:00:01	2022-05-31 23:00:04	0:00:03
47916	Speedy	speedy-fd	OK	0	Backup	Inc	9,239	15,465,671,816	2022-05-31 23:00:01	2022-05-31 23:10:29	0:10:28
47915	pi	pi-fd	OK	0	Backup	Inc	56	14,976,676	2022-05-31 23:00:01	2022-05-31 23:00:23	0:00:22

Fig. 1: Example of a report showing some backup jobs, some copy control jobs, and two running jobs. Notice that all JobIds are URL links to either BWeb Management Suite (Bacula Enterprise) or Baculum (Bacula Community).

Read more:

Daemons Monitoring

In order to monitor daemons, use the Monitor configuration file. The Monitor configuration file is a stripped down version of the Director configuration file, mixed with a Console configuration file. It simply contains the information necessary to contact Directors, Clients, and Storage daemons you want to monitor.

For a general discussion of configuration file and resources including the data types recognized by **Bacula**, see the Configuration chapter.

The following Monitor Resource definition must be defined:

- Monitor – to define the Monitor’s name used to connect to all the daemons and the password used to connect to the Directors. Note, you must not define more than one Monitor resource in the Monitor configuration file.
- At least one *Client*, *Storage* or *Director* resource, to define the daemons to monitor.

Monitor Resource

The Monitor resource defines the attributes of the Monitor running on the network. The parameters you define here must be configured as a Director resource in Clients and Storages configuration files, and as a Console resource in Directors configuration files.

Monitor Start of the Monitor records.

Name = <name> Specify the Director name used to connect to Client and Storage, and the Console name used to connect to Director. This record is required.

DisplayAdvancedOptions = <boolean> Display advanced options in the tray monitor (for backup and restore operations)

CommandDirectory = <directory> Directory where the tray monitor will look at a regular interval to find commands to execute.

Refresh Interval = <time> Specifies the time to wait between status requests to each daemon. It can't be set to less than 1 second, or more than 10 minutes, and the default value is **5 seconds**.

Director Resource

The Director resource defines the attributes of the Directors that are monitored by this Monitor.

As you are not permitted to define a Password in this resource, to avoid obtaining full Director privileges, you must create a Console resource in Director's configuration file, using the Console Name and Password defined in the Monitor resource. To avoid security problems, you should configure this Console resource to allow access to no other daemons, and permit the use of only two commands: **status** and **.status** (see below for an example).

You may have multiple Director resource specifications in a single Monitor configuration file.

Director Start of the Director records.

Name = <name> The Director name used to identify the Director in the list of monitored daemons. It is not required to be the same as the one defined in the Director's configuration file. This record is required.

Port = <port-number> Specify the port to use to connect to the Director. This value will most likely already be set to the value you specified on the **--with-baseport** option of the **./configure** command. This port must be identical to the **DIRport** specified in the **Director** resource of the Director's configuration file. The default is **9101** so this record is not normally specified.

Address = <address> Where the address is a host name, a fully qualified domain name, or a network address used to connect to the Director. This record is required.

Client Resource

The Client resource defines the attributes of the Clients that are monitored by this Monitor.

You must create a Director resource in the Client's configuration file, using the Director Name defined in the Monitor resource. To avoid security problems, you should set the **Monitor** directive to **Yes** in this Director resource.

You may have multiple Director resource specifications in a single Monitor configuration file.

Client (or FileDaemon) Start of the Client records.

Name = <name> The Client name used to identify the Director in the list of monitored daemons. It is not required to be the same as the one defined in the Client's configuration file. This record is required.

Address = <address> Where the address is a host name, a fully qualified domain name, or a network address in dotted quad notation for a **Bacula** File daemon. This record is required.

Port = <port-number> Where the port is a port number at which the **Bacula** File daemon can be contacted. The default is **9102**.

Password = <password> This is the password to be used when establishing a connection with the File services, so the Client configuration file on the machine to be backed up must have the same password defined for this Director. This record is required.

Storage Resource

The Storage resource defines the attributes of the Storages that are monitored by this Monitor.

You must create a Director resource in the Storage's configuration file, using the Director Name defined in the Monitor resource. To avoid security problems, you should set the **Monitor** directive to **Yes** in this Director resource.

You may have multiple Director resource specifications in a single Monitor configuration file.

Storage Start of the Storage records.

Name = <name> The Storage name used to identify the Director in the list of monitored daemons. It is not required to be the same as the one defined in the Storage's configuration file. This record is required.

Address = <address> Where the address is a host name, a fully qualified domain name, or a network address in dotted quad notation for a **Bacula** Storage daemon. This record is required.

Port = <port> Where port is the port to use to contact the storage daemon for information and to start jobs. This same port number must appear in the Storage resource of the Storage daemon's configuration file. The default is **9103**.

Password = <password> This is the password to be used when establishing a connection with the Storage services. This same password also must appear in the Director resource of the Storage daemon's configuration file. This record is required.

Tray Monitor Security

There is no security problem in relaxing the permissions on **tray-monitor.conf** as long as FD, SD and DIR are configured properly, so the passwords contained in this file only gives access to the status of the daemons. It could be a security problem if you consider the status information as potentially dangerous.

Concerning Director's configuration: In **tray-monitor.conf**, the password in the Monitor resource must point to a restricted console in **bacula-dir.conf** (details here). So, if you use this password with bconsole, you'll only have access to the status of the director (commands **status** and **.status**). It could be a security problem if there is a bug in the ACL code of the director.

Concerning File and Storage Daemons' configuration: In **tray-monitor.conf**, the Name in the Monitor resource must point to a Director resource in **bacula-fd/sd.conf**, with the Monitor directive set to **yes** (once again, see the documentation). It could be a security problem if there is a bug in the code which check if a command is valid for a Monitor (this is very unlikely as the code is pretty simple).

Click [here](#) to see the Tray Monitor sample configuration.

Tray Monitor Sample Configuration

An example Tray Monitor configuration file might be the following:

```
#
# Bacula Tray Monitor Configuration File
#
Monitor {
    Name = rufus-mon # password for Directors
    RefreshInterval = 10 seconds
}
Client {
    Name = rufus-fd
    Address = rufus
    Port = 9102
    # password for FileDaemon
    Password = "FYpq4yyI1y562EMS35bA0J0QC0M2L3t5cZ0bxT3XQxgxpTn"
}
Storage {
    Name = rufus-sd
    Address = rufus
    Port = 9103
    # password for StorageDaemon
    Password = "9usxgc307dMbe7jbD16v0PXlhD64UVasIDD0DH2WAujcDsc6"
}
Director {
    Name = rufus-dir
    port = 9101
    address = rufus
}
```

Sample File Daemon's Director Record

[Click here](#) to see the full example.

```
#
# Restricted Director, used by tray-monitor to get the
#
status of the file daemon
#
Director {
    Name = rufus-mon
    Password = "FYpq4yyI1y562EMS35bA0J0QC0M2L3t5cZ0bxT3XQxgxpTn"
    Monitor = yes
}
```

Sample Storage Daemon's Director Record

[Click here to see the full example.](#)

```
#
# Restricted Director, used by tray-monitor to get the
#
# status of the storage daemon
#
Director {
    Name = rufus-mon
    Password = "9usxgc307dMbe7jbD16v0PXlhD64UVasIDD0DH2WAujcDsc6"
    Monitor = yes
}
```

Sample Director's Console Record

[Click here to see the full example.](#)

```
#
# Restricted console used by tray-monitor to get the status of the director
#
Console {
    Name = Monitor
    Password = "GN0uRo7PTUm1MbqrJ2Gr1p0fk0HQJTxwnFyE4WSST3MWZseR"
    CommandACL = status, .status
}
```

Go back to the *Daemons Monitoring* chapter.

Go back to the *Jobs Monitoring* chapter.

Go back to the *main Bacula Enterprise Management page*.

See also:

Go to:

- *Using Bacula Catalog to Grab Information*
- *BWeb with Graphite*
- *Bacula Simple Network Management Protocol (bsnmp)*
- *Send Joblogs, Alerts, Mount Requests via Email (bsmtp)*

Go back to the *Jobs Monitoring*.

Go back to the *main Bacula Enterprise Management page*.

Using Bacula Catalog to Grab Information

Bacula Catalog contains lots of information about your IT infrastructure, how many files, their size, the number of video or music files etc. Using **Bacula** catalog during the day to get them permit to save resources on your servers.

In this chapter, you will find tips and information to measure bacula efficiency and report statistics.

Job Statistics

If you want to have statistics on your backups to provide some SLA indicators, you could use a few SQL queries on the table to report how many:

- jobs have run
- jobs have been successful
- files have been backed up
- and so on.

However, these statistics are accurate only if your job retention is greater than your statistics period, i.e, if jobs are purged from the catalog, you won't be able to use them.

Since the introduction of the `jobhisto` catalog DB table, you can use the `update stats [days=num]` console command to fill this table with new Job records. If you want to be sure to take in account only **good jobs**, e.g. if one of your important job has failed but you have fixed the problem and restarted it on time, you probably want to delete the first *bad* job record and keep only the successful one. For that simply let your staff do the job, and update table after two or three days depending on your organization using the `[days=num]` option.

These statistics records aren't used for restoring, but mainly for capacity planning, billings, etc.

The **BWeb** interface provides a statistics module that can use this feature. You can also use third party tools (like Talend) or extract information by yourself.

The **Statistics Retention** = Director directive defines the length of time that **Bacula** will keep statistics job records in the Catalog database after the Job End time. When this time period expires, and if user runs `prune stats` command, **Bacula** will prune (remove) Job records that are older than the specified period.

For example, you can use the following **RunScript** console commands in your nightly **BackupCatalog** job to maintain statistics.

```
Job {
  Name = BackupCatalog
  ...
  RunScript {
    Console = "update stats days=3"
    Console = "prune stats yes"
    RunsWhen = After
    RunsOnClient = no
  }
}
```

Go back to [Using Bacula Catalog to Grab Information](#).

Go back to the [Jobs Monitoring](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

See also:

Go back to:

- *Daemons Monitoring*

Go to:

- *BWeb with Graphite*
- *Bacula Simple Network Management Protocol (bsnmp)*
- *Send Joblogs, Alerts, Mount Requests via Email (bsmtp)*


Go back to the *Jobs Monitoring*.

Go back to the *main Bacula Enterprise Management page*.

BWeb with Graphite

The following article explains how to customize your Bacula Enterprise installation with BWeb to use the powerful features of the Graphite visualization engine. Since version 12.0.0 BWeb can integrate Graphite widgets.

The workflow is as follows:



Go through the steps:

Graphite Installation

We will not explain in detail how to set up a Graphite server. The project's website <https://graphiteapp.org/> has a Getting Started section where you can find further information.

See also:

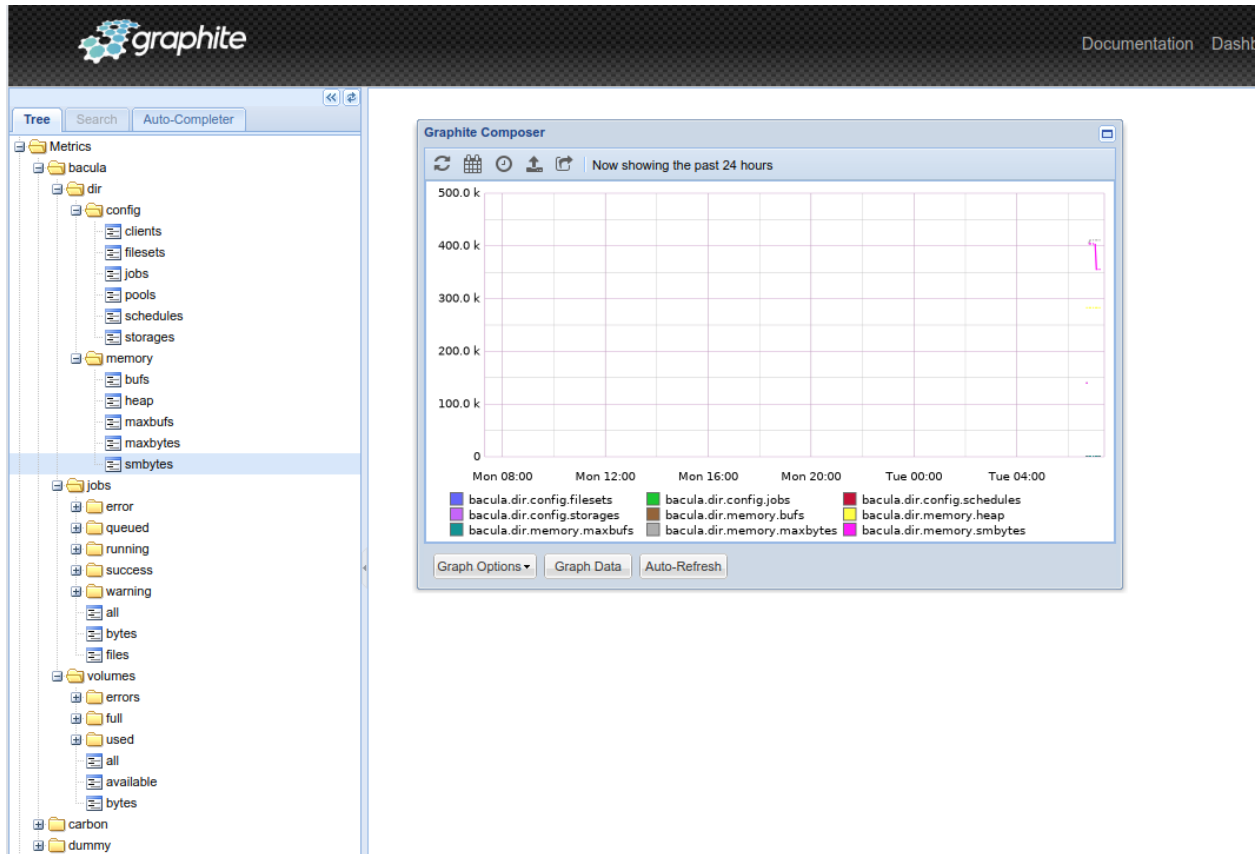
Go to:

- *Install Bacula Enterprise and BWeb Version 12.0.0 or Higher*
- *Configure a Statistics Resource in BWeb*
- *Browse Statistics Data in Graphite Web GUI and Export URL*
- *Create Widget in BWeb and Add Graphite Metrics to Dashboard*

Go back to the *BWeb with Graphite* chapter.

Go back to the *Jobs Monitoring*.

Go back to the *main Bacula Enterprise Management page*.



Install Bacula Enterprise and BWeb Version 12.0.0 or Higher

Read the documents:

- [BaculaEnterpriseInstallation](#)
- [BWeb Installation](#)

and follow the instructions therein to install the Bacula Enterprise components.

See also:

Go back to:

- [Graphite Installation](#)

Go to:

- [Configure a Statistics Resource in BWeb](#)
- [Browse Statistics Data in Graphite Web GUI and Export URL](#)
- [Create Widget in BWeb and Add Graphite Metrics to Dashboard](#)

Go back to the *BWeb with Graphite* chapter.

Go back to the *Jobs Monitoring*.

Go back to the *main Bacula Enterprise Management page*.

Configure a Statistics Resource in BWeb

In the BWeb configuration overview screen (Workset) you can select one of the three major Bacula Components: Director, Storage Daemon, and File Daemon. Once you have clicked on the button, the “Actions” box on the left contains a “List Statistics” button that will bring you to the Statistics resource overview page for the selected Bacula component.

Use the “Plus” symbol to add a new Statistics resource. It needs a Name (mandatory), optional description, you can set the interval with which Bacula will send information to the statsd process on your Graphite server. In the Type dropdown field select “Graphite” and enter metrics that you want to export e.g. bacula.jobs.*. A prefix can be useful to distinguish between your Bacula File Daemons for instance. Open the Graphite Options to enter the FQDN or IP address of your Graphite server and the port where Carbon receiver listens to plaintext.

Save the Statistics resource and restart the Bacula daemon that will then send statistics data to the Graphite host and port with the time interval specified.

It is possible to check if *statsd* is running fine with bconsole:

```
.status dir statistics
```

In BWeb this can be done with the web console which you will have to activate in the top menu, Configuration, Configure BWeb, Edit, and then check “Web Bconsole”. This will give you a new icon next to the search field in the top menu where you can access a bconsole to issue the above `.status` command.

See also:

Go back to:

- [Graphite Installation](#)
- [Install Bacula Enterprise and BWeb Version 12.0.0 or Higher](#)

Go to:

- [Browse Statistics Data in Graphite Web GUI and Export URL](#)
- [Create Widget in BWeb and Add Graphite Metrics to Dashboard](#)

Go back to the *BWeb with Graphite* chapter.

Go back to the *Jobs Monitoring*.

Go back to the *main Bacula Enterprise Management page*.

Browse Statistics Data in Graphite Web GUI and Export URL

In the Graphite web interface you will find the new Bacula statistics in the Metrics section node tree structure on the left. Drill down to the metrics that you want to display and modify the graph to your liking. Once done, click the “Short URL” button in the Graphite Composer window to retrieve the URL to embed the statistics graph into the BWeb web interface.

See also:

Go back to:

- [Graphite Installation](#)
- [Install Bacula Enterprise and BWeb Version 12.0.0 or Higher](#)
- [Configure a Statistics Resource in BWeb](#)

Go to:

- [Create Widget in BWeb and Add Graphite Metrics to Dashboard](#)


Go back to the *BWeb with Graphite* chapter.

Go back to the *Jobs Monitoring*.


Go back to the *main Bacula Enterprise Management page*.

Create Widget in BWeb and Add Graphite Metrics to Dashboard

In BWeb choose Configuration | Page Composer from the top menu. The button “Add new box” will allow you to create a custom graph based on the Graphite metrics: you choose a name for the new box/widget and an ID, then you enter the URL exported in step 4 and you can specify a custom refresh interval (each box in the BWeb dashboard can have its own refresh interval).



Confirm with “Add box” and drag and drop the new box into the page composer main window to modify your BWeb dashboard. You can change the position and size of the widgets.



See also:

Go back to:

- *Graphite Installation*
- *Install Bacula Enterprise and BWeb Version 12.0.0 or Higher*
- *Configure a Statistics Resource in BWeb*
- *Browse Statistics Data in Graphite Web GUI and Export URL*

Go back to the *BWeb with Graphite* chapter.

Go back to the *Jobs Monitoring*.

Go back to the *main Bacula Enterprise Management page*.

See also:

Go back to:

- *Daemons Monitoring*
- *Using Bacula Catalog to Grab Information*

Go to:

- *Bacula Simple Network Management Protocol (bsnmp)*
- *Send Jobs, Alerts, Mount Requests via Email (bsmtp)*

Go back to the *Jobs Monitoring*.

Go back to the *main Bacula Enterprise Management page*.

Bacula Simple Network Management Protocol (bsnmp)

SNMP or Simple Network Management Protocol is a widely used protocol that helps to organize and collect information about managed devices on IP networks. SNMP exposes data from the managed devices in the form of variables that are organized in a management information base (MIB). This MIB is designed to describe status and configuration information. These variables can then be remotely accessed by centralized monitoring applications.

A given MIB offers a convention of variables that structure information in a way that is useful for the target system to be monitored. Different target systems that are used for the same purpose will share this structure, while reporting different values depending on their activities.

Bacula Enterprise activities can be remotely monitored through this protocol. Bacula Enterprise provides a simple script that can send SNMP notifications (SNMP INFORM) to a remote SNMP service. This service must be provided by any third-party application that offers this kind of monitoring service. The generated notifications can contain key data on how jobs are behaving, as job configuration data, job status and other values. This document describes how to setup this integration.

Note: This feature was added with Bacula version 16.0.

It is currently available as packaged add-on for a limited number of supported systems.

Instructions to install, test, and integrate the bsnmp job INFORM sender.

The functionality is provided by a very simple Python script which requires some MIBs to be available. Configuration is provided in a plain text file in ini format.

Script Installation

1. Put the file bsnmp into a location on the DIR host where the DIR can execute programs in. We propose to use /opt/bacula/scripts
2. Make the file executable by the Bacula DIR user account, usually bacula:

```
chown bacula. /opt/bacula/scripts/bsnmp
chmod u+x /opt/bacula/scripts/bsnmp
```

3. Ensure python3 is available: python3 --version should not fail.

4. Ensure the needed dependencies are installed. This can be somewhat difficult and may depend on distribution and python installation. In typical cases, with package based installations (i.e., no pip-based ones and no virtual environments and such things):

```
$ python3
Python 3.7.3 (default, Jan 22 2021, 20:04:44)
[GCC 8.3.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import argparse
>>> import configparser
>>> from pysnmp.hlapi import *
Traceback (most recent call last):
File "<stdin>", line 1, in <module>
ModuleNotFoundError: No module named 'pysnmp'
>>>
```

indicates the module providing snmp functionality is not installed. On a Debian (based) system, `apt-get install python3-pysnmp4 python3-pysmi` might install what is needed (proper dependencies, also for other distributions, will be sorted out later).

5. The python snmp module needs a local directory writeable by itself. This needs to be created for the Bacula DIR user, as the home directory for this user is not writable:

```
mkdir /opt/bacula/.pysnmp"
chown bacula: /opt/bacula/.pysnmp/
```

See also:

Go to:

- [MIB Installation](#)
- [Script Configuration](#)
- [Verification of Operations](#)
- [Example snmptrapd Configuration for Testing Purposes](#)
- [Integration with DIR Messages Resource](#)

Go back to the [Bacula Simple Network Management Protocol \(bsnmp\)](#) chapter.

Go back to the [Jobs Monitoring](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

MIB Installation

1. MIBs for basic snmp infrastructure may be needed. Installation from distribution repos or tools will usually be sufficient. On a Debian (based) system, adding or enabling the non-free distribution repository and running `apt-get install snmp-mibs-downloader` will usually be sufficient.
2. The two provided MIB files should be copied to the system wide MIB repository:

```
cp -iv *-MIB.txt /usr/share/snmp/mibs/
chown root. /usr/share/snmp/mibs/ITS-*
```

See also:

Go back to:

- [Script Installation](#)

Go to:

- [Script Configuration](#)
- [Verification of Operations](#)
- [Example snmptrapd Configuration for Testing Purposes](#)
- [Integration with DIR Messages Resource](#)

Go back to the [Bacula Simple Network Management Protocol \(bsnmp\)](#) chapter.

Go back to the [Jobs Monitoring](#) chapter.

Go back to the [main Bacula Enterprise Management](#) page.

Script Configuration

1. The script needs a configuration file where it can be read by the Bacula DIR user. We suggest copying the provided example configuration to a suitable location:

```
cp -iv bsnmp.conf /opt/bacula/etc/  
chown bacula. /opt/bacula/etc/bsnmp.conf
```

2. The configuration needs to be adapted. Using your favourite plain text editor, edit it:

```
[thisengine]  
username = USM user name on TRAP/INFORM receiver system  
authkey = password for the above user  
  
[target]  
ipv4 host = DNS or IP address of trap receiver. IPv4 only  
udp port = 162
```

The information to put will, naturally, be very site specific and may need to be provided by a different team/person. Note that, at this time:

- *only* SNMPv3 is supported (no, we do not plan to introduce non-encrypted, non-authenticated messaging!)
- *only* the USM scheme is supported (we *may* support more in the future, but USM still seems to be the typical case)
- *only* HMAC-SHA1 authentication and AES-CFB-128 encryption are supported. This will probably be extended in the future
- *only* UDP transport is supported. This should not be a problem. The port number is optional, default is 162, which is the assigned port number, and thus should not need modification in most cases.

See also:

Go back to:

- [Script Installation](#)
- [MIB Installation](#)

Go to:

- [Verification of Operations](#)
- [Example snmptrapd Configuration for Testing Purposes](#)

- [Integration with DIR Messages Resource](#)

Go back to the [Bacula Simple Network Management Protocol \(bsnmp\)](#) chapter.

Go back to the [Jobs Monitoring](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

Verification of Operations

Assuming the installation and configuration has been done per the above suggestions, running, as the bacula user,

```
/opt/bacula/scripts/bsnmp -c /opt/bacula/etc/bsnmp.cfg -T
```

should result in no console output, and an SNMP INFORM message received at the configured receiver.

Any console output should be reported so we can assist or fix what needs fixing.

See also:

Go back to:

- [Script Installation](#)
- [MIB Installation](#)
- [Script Configuration](#)

Go to:

- [Example snmptrapd Configuration for Testing Purposes](#)
- [Integration with DIR Messages Resource](#)

Go back to the [Bacula Simple Network Management Protocol \(bsnmp\)](#) chapter.

Go back to the [Jobs Monitoring](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

Example snmptrapd Configuration for Testing Purposes

1. The following suggestions are intended to be useful for reasonably skilled system administrators willing to read and understand a manual page as well as willing to spend a bit of time experimenting and learning. If more directed help is needed, ask the experts (TM).

We assume a management system ingesting SNMP INFORMs is already in place. It would not be Bacula specific anyway.

2. Install all possibly needed MIBs on the designated snmp receiver host as outlined above.
3. Create an snmptrapd user. Manual pages are useful, **snmp** understanding is as well. In short: Put appropriate createuser lines into `/var/lib/snmp/snmptrapd.conf` and (re) start snmptrapd. The lines will be converted to local, engine-specific auth lines. `createuser informsubmit SHA informpass AES` is what was used for the default config file we provide.
4. For easier readability, starting snmptrapd with all MIBs loaded may be useful. For that purpose, editing the systemd unit file is probably the simplest approach. Thus:

```
systemctl show snmptrapd.service | grep ExecStart
```

to find the actual program call, and

```
systemctl edit snmptrapd.service
```

to add the needed `-m ALL` option:

```
[Service]
ExecStart=
ExecStart=/usr/sbin/snmptrapd -Lsd -f -m ALL -p /run/snmptrapd.pid
```

5. If things work out, after the `bsnmp -T` call, information such as

```
Jul 18 13:57:37 bsys-docdev snmptrapd[23559]: 2022-07-18 13:57:37 gnom.os.its-
↳ lehmann.de [UDP: [192.168.0.55]:52424->[192.168.0.44]:162]:
DISMAN-EXPRESSION-MIB::sysUpTimeInstance = Timeticks: (0) 0:00:00.00          SNMPv2-
↳ MIB::snmpTrapOID.0 = OID: ITS-Bacula-MIB::baculaNotifyJob          ITS-Bacula-
↳ MIB::baculaJobName = STRING: Test-Job-Ignore          ITS-Bacula-
↳ MIB::baculaJobResult = INTEGER: stopped(4)
```

should be logged.

See also:

Go back to:

- [Script Installation](#)
- [MIB Installation](#)
- [Script Configuration](#)
- [Verification of Operations](#)

Go to:

- [Integration with DIR Messages Resource](#)

Go back to the [Bacula Simple Network Management Protocol \(bsnmp\)](#) chapter.

Go back to the [Jobs Monitoring](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

Integration with DIR Messages Resource

1. Identify the configuration that defines the **Messages** for the Jobs that should trigger SNMP INFORMs to be sent.
2. Modify the **Messages** Resource so that the `bsnmp` script call is added as a new **Mail** message destination. In a fairly default installation, this could be

```
Messages {
  Name = "Standard"
  MailCommand = "/opt/bacula/bin/bsmtmp -h <hidden> -f \"(Bacula) <%r>\" -s \
↳ "Bacula: %t %e of %n %l\" %r"
  OperatorCommand = "/opt/bacula/bin/bsmtmp -h <hidden> -f \"(Bacula) <%r>\" -s \
↳ "Bacula: Intervention needed for %j\" %r"
  Mail = "admin" = All,!Debug,!Saved,!Skipped
  MailCommand = "/opt/bacula/scripts/bsnmp -c /opt/bacula/etc/bsnmp.conf"
  Mail = snmp-dummy = All,!Debug,!Saved,!Skipped
  Append = "/opt/bacula/log/bacula.log" = All,!Debug,!Saved,!Skipped
```

(continues on next page)

```
Console = All,!Debug,!Saved,!Skipped
Operator = "admin" = Mount
Catalog = All,Events,!Debug,!Saved
}
```

The key points to observe are

- The original Mail sending can be left as is, but all the entries for the new delivery method must follow the existing ones. So, MailCommand and Mail lines are added after the existing mail related entries.
 - The actual address used with the bsnmp script is ignored, but needs to be provided per syntax definition. Put whatever you like.
 - The actual message types used are not critical, but make sure that final Job reports get delivered. File lists, debug output, events and any other message types are, currently, ignored.
3. Verify configuration syntax as usual (bacula-dir -t)
 4. Reload configuration. In bconsole, reload command.
 5. Run a Job. After finishing, an SNMP INFORM should be delivered. When testing with snmptrapd (see above), information like

```
Jul 17 23:06:58 bsys-docdev snmptrapd[23559]: 2022-07-17 23:06:58 gnom.os.its-
↳ lehmman.de [UDP: [192.168.0.55]:45198->[192.168.0.44]:162]:
DISMAN-EXPRESSION-MIB::sysUpTimeInstance = Timeticks: (0) 0:00:00.00      SNMPv2-
↳ MIB::snmpTrapOID.0 = OID: ITS-Bacula-MIB::baculaNotifyJob      ITS-Bacula-
↳ MIB::baculaJobName = STRING: moria-home      ITS-Bacula-MIB::baculaJobResult =_
↳ INTEGER: ok(0)
```

should be logged.

See also:

Go back to:

- [Script Installation](#)
- [MIB Installation](#)
- [Script Configuration](#)
- [Verification of Operations](#)
- [Example snmptrapd Configuration for Testing Purposes](#)

Go back to the *Bacula Simple Network Management Protocol (bsnmp)* chapter.

Go back to the *Jobs Monitoring* chapter.

Go back to the *main Bacula Enterprise Management page*.

See also:

Go back to:

- [Daemons Monitoring](#)
- [Using Bacula Catalog to Grab Information](#)
- [BWeb with Graphite](#)

Go to:

- *Send Joblogs, Alerts, Mount Requests via Email (bsmtp)*

Go back to the *Jobs Monitoring*.

Go back to the *main Bacula Enterprise Management page*.

Send Joblogs, Alerts, Mount Requests via Email (bsmtp)

bsmtp is a simple mail transport program that permits more flexibility than the standard mail programs typically found on Unix systems. It can even be used on Windows machines.

It is called:

```
Usage: bsmtp [-f from] [-h mailhost] [-s subject] [-c copy] [recipient ...]
  -c          set the Cc: field
  -dnn       set debug level to nn
  -f          set the From: field
  -h          use mailhost:port as the bsmtp server
  -l         limit the lines accepted to nn
  -s         set the Subject: field
  -?         print this message.
```

If the `-f` option is not specified, bsmtp will use your userid. If the option `-h` is not specified bsmtp will use the value in the environment variable `bsmtpSERVER` or if there is none `localhost`. By default port 25 is used.

If a line count limit is set with the `-l` option, bsmtp will not send an email with a body text exceeding that number of lines. This is especially useful for large restore job reports where the list of files restored might produce very long mails your mail-server would refuse or crash. However, be aware that you will probably suppress the job report and any error messages unless you check the log file written by the Director (see the messages resource in this manual for details).

recipients is a space separated list of email recipients.

The body of the email message is read from standard input.

An example of the use of bsmtp would be to put the following statement in the Messages resource of your `bacula-dir.conf` file. Note, these commands should appear on a single line each.

```
mailcommand = "/home/bacula/bin/bsmtp -h mail.domain.com -f \"\\(Bacula\\) %r\"
              -s \"Bacula: %t %e of %c %l\" %r"
operatorcommand = "/home/bacula/bin/bsmtp -h mail.domain.com -f \"\\(Bacula\\) %r\"
                  -s \"Bacula: Intervention needed for %j\" %r"
```

Where you replace `/home/bacula/bin` with the path to your Bacula binary directory, and you replace `mail.domain.com` with the fully qualified name of your bsmtp (email) server, which normally listens on port 25. For more details on the substitution characters (e.g. `%r`) used in the above line, please see the documentation of the Mail Command in the Messages Resource chapter of the Bacula Enterprise Main manual.

It is highly recommended that you test one or two cases by hand to make sure that the mailhost that you specified is correct and that it will accept your email requests. Since bsmtp always uses a TCP connection rather than writing in the spool file, you may find that your from address is being rejected because it does not contain a valid domain, or because your message is caught in your spam filtering rules. Generally, you should specify a fully qualified domain name in the from field, and depending on whether your bsmtp gateway is Exim or Sendmail, you may need to modify the syntax of the from part of the message. Please test.

When running bsmtp by hand, you will need to terminate the message by entering a `ctl-d` in column 1 of the last line.

If you are getting incorrect dates (e.g. 1970) and you are running with a non-English language setting, you might try adding a `LANG="en_US"` immediately before the bsmtp call.

In general, `bsmtp` attempts to cleanup email addresses that you specify in the `from`, `copy`, `mailhost`, and `recipient` fields, by adding the necessary `<` and `>` characters around the address part. However, if you include a `display-name` (see RFC 5332), some SMTP servers such as Exchange may not accept the message if the `display-name` is also included in `<` and `>`. As mentioned above, you must test, and if you run into this situation, you may manually add the `<` and `>` to the `Bacula mailcommand` or `operatorcommand` and when `bsmtp` is formatting an address if it already contains a `<` or `>` character, it will leave the address unchanged.

See also:

Go back to:

- [Daemons Monitoring](#)
- [Using Bacula Catalog to Grab Information](#)
- [BWeb with Graphite](#)
- [Bacula Simple Network Management Protocol \(*bsnmp*\)](#)

Go back to the [Jobs Monitoring](#).

Go back to the [main Bacula Enterprise Management page](#).

See also:

Go back to:

- [Daily Management](#)
- [Logs Management](#)

Go to:

- [Storage Space Management](#)
- [PostgreSQL Catalog Administration](#)
- [Tape/Volume Management](#)
- [Tape Autochanger Usage](#)

Go back to the [main Bacula Enterprise Management page](#).

4 Storage Space Management

The following chapter provides information on how to manage storage space.

4.1 Volume Recycling

By default, once **Bacula** starts writing a Volume, it may append to the volume, but it will not overwrite the existing data thus destroying it. However when **Bacula recycles** a Volume, the Volume becomes available for being reused, and **Bacula** can at some later time overwrite the previous contents of that Volume. At that point all previous data on that Volume will be lost. If the Volume is a tape, the tape will be rewritten from the beginning. If the Volume is a disk file, the file will be truncated before being rewritten.

You may not want Bacula to automatically recycle (reuse) Volumes. Doing so may require a large number of Volumes though. However, it is also possible to manually recycle Volumes so that they may be reused. For more on manual recycling, see the section entitled [Manually Recycling Volumes](#).

Most people prefer to have a Pool of Volumes that are used for daily backups and recycled once a week, another Pool of Volumes that are used for Full backups once a week and recycled monthly, and finally a Pool of Volumes that are

used once a month and recycled after a year or two. With a scheme like this, the number of Volumes in your pool or pools remains constant.

By properly defining your Volume Pools with appropriate Retention periods, **Bacula** can manage the recycling (such as defined above) automatically.

Automatic Volume Recycling

Automatic recycling of Volumes is controlled by four records in the resource definition in the Director's configuration file. These four records are:

- AutoPrune = yes
- VolumeRetention = <time>
- Recycle = yes
- RecyclePool = <APool>
- ScratchPool = <APool>

The above first three directives are all you need assuming that you fill each of your Volumes then wait the Volume Retention period before reusing them, providing there is some non-pruned Jobs or Files on the Volume. **Recycle Pool** and **Scratch Pool** directives are not required. If not defined, **Bacula** will recycle Volumes and keep them in the current Pool.

If you want **Bacula** to stop using a Volume and recycle it before it is full, you will need to use one or more additional directives such as:

- **Use Volume Once=yes**
- **Volume Use Duration=ttt**
- **Maximum Volume Jobs=nnn**
- **Maximum Volume Bytes=mmm**

See below and the Basic Volume Management chapter for complete examples.

Automatic recycling of Volumes is performed by **Bacula** only when it wants a new Volume and no appendable Volumes are available in the Pool. It will then search the Pool for any Volumes with the flag set and the Volume Status is. At that point, it will choose the oldest purged volume and recycle it.

If there are no volumes with status, then the recycling occurs in two steps:

The first is that the Catalog for a Volume must be pruned of all Jobs (i.e. Purged) and Files contained on that Volume.

The second step is the actual recycling of the Volume. Only Volumes marked or will be considered for pruning. The Volume will be purged, all Jobs and Files associated to this Volume are pruned from Catalog, if the **VolumeRetention** period has expired. When a Volume is marked as , it means that no Catalog records for Jobs and Files reference that Volume, and the Volume can be recycled and reused. Please note a Volume can be reused even though the **VolumeRetention** period has not expired if the Jobs and Files associated to the Volume had been already pruned. Until recycling actually occurs, the Volume data remains intact. If no Volumes can be found for recycling for any of the reasons stated above, **Bacula** will request operator intervention (i.e. it will ask you to label a new volume).

A key point mentioned above, that can be a source of frustration, is that **Bacula** will only recycle purged Volumes if there is no other appendable Volume available, otherwise, it will always write to an appendable Volume before recycling even if there are Volumes marked as . This preserves your data as long as possible. So, if you wish to “force” **Bacula** to use a purged Volume, you must first ensure that no other Volume in the Pool is marked. If necessary, you can manually set a volume to. The reason for this is that **Bacula** wants to preserve the data on your old Volumes (even though purged from the catalog) as long as absolutely possible before overwriting it. There are also a number of directives such as **Volume Use Duration** that will automatically mark a volume as and thus no longer appendable.

Recycling Algorithm

When a Job needs a Volume and no appendable, recycled or purged is available, **Bacula** starts the recycling algorithm pruning Jobs and Files for the oldest Volume in the pool used by the job. After the Volume is pruned, and if the **Recycle** flag is on (**Recycle=yes**) for that Volume, **Bacula** will relabel it and write new data on it.

As mentioned above, there are two key points for getting a Volume to be recycled. First, the Volume must no longer be marked Append (there are a number of directives to automatically make this change), and second since the last write on the Volume, one or more of the Retention periods must have expired so that there are no more catalog backup job records that reference that Volume. Once both those conditions are satisfied, the volume can be marked Purged and hence recycled.

The algorithm described below assumes that AutoPrune is enabled, that Recycling is turned on, and that you have defined appropriate Retention periods, or used the defaults for all these items.

- If the request is for an Autochanger device, look only for Volumes in the Autochanger (i.e. with InChanger set and that have the correct Storage device).
- Search the Pool for a Volume with VolStatus=Append (if there is more than one, the Volume with the oldest date last written is chosen. If two have the same date then the one with the lowest MediaId is chosen).
- Search the Pool for a Volume with VolStatus=Recycle and the InChanger flag is set true (if there is more than one, the Volume with the oldest date last written is chosen. If two have the same date then the one with the lowest MediaId is chosen).
- Search the Pool for a Volume with VolStatus=Purged and try recycling any purged Volumes.
- Prune volumes applying Volume retention period (Volumes with VolStatus or are pruned). Note, when all the File and Job records are pruned from a Volume, the Volume will be marked and this may be prior to the expiration of the Volume retention period.
- Prune the oldest Volume if **RecycleOldestVolume=**. **yes** (the Volume with the oldest LastWritten date and VolStatus equal to or is chosen). This record ensures that all retention periods are properly respected.
- Purge the oldest Volume if **PurgeOldestVolume=yes** (the Volume with the oldest LastWritten date and VolStatus equal to or is chosen). We strongly recommend against the use of **PurgeOldestVolume** as it can quite easily lead to loss of current backup data.
- Search for a Volume in the Scratch Pool (the default Scratch pool or another Scratch Pool defined for the pool) and if found move it to the current Pool for the Job and use it. Note, when the Scratch Volume is moved into the current Pool, the basic Pool defaults are applied as if it is a newly labeled Volume (equivalent to an **update volume from pool** command).
- If we were looking for Volumes in the Autochanger, go back to step 2 above, but this time, look for any Volume whether or not it is in the Autochanger.
- Attempt to create a new Volume if automatic labeling is enabled and the maximum number of Volumes is not reached.
- Give up and ask operator.

The above occurs when **Bacula** has finished writing a Volume or when no Volume is present in the drive.

On the other hand, if you have inserted a different Volume after the last job, and **Bacula** recognizes the Volume as valid, the Storage daemon will request authorization from the Director to use this Volume. In this case, if you have set **Recycle CurrentVolume = yes=** and the Volume is marked as Used or Full, Bacula will prune the volume and if all jobs were removed during the pruning (respecting the retention periods), the Volume will be recycled and re-used.

If you want to strictly observe the **Volume Retention Period** under all circumstances – i.e. you have a Volume with critical data, you must set **Recycle=no**. However, in doing so, you must manually do the recycling of the Volume for it to be used again.

The recycling algorithm in this case is:

- If the VolStatus is or is set, the volume will be used.
- If **Recycle Current Volume** is set and the volume is marked, or **Bacula** will prune the volume (applying the retention period). If all Jobs are pruned from the volume, it will be recycled.

This permits users to manually change the Volume every day and load volumes in an order different from what is in the catalog, and if the volume does not contain a current copy of your backup data, it will be used.

A few points to keep in mind:

1. If a pool doesn't have maximum volumes defined then **Bacula** will prefer to demand new volumes over forcibly purging older volumes.
2. If volumes become free through pruning, then they get marked as and are immediately available for recycling – these will be used in preference to creating new volumes.
3. If the Job, File, and Volume retention periods are different, then it's common to see a volume with no files or jobs listed in the database, but which is still not marked as **Purged**, if **Autoprune** is set to **no** in the pool resource.

See also:

Go to [Recycle Status](#).

Go back to the [Volume Recycling](#) chapter.

Go back to the [Storage Space Management](#) chapter.

Go back to the main [Bacula Enterprise](#) chapter.

Recycle Status

Each Volume inherits the Recycle status (yes or no) from the Pool resource record when the Media record is created (normally when the Volume is labeled). This Recycle status is stored in the Media record of the Catalog. Using the Console program, you may subsequently change the Recycle status for each Volume. For example in the following output from **list volumes**:

VolumeNa	Media	VolSta	VolByte	LastWritte	VolRet	Rec
File0001	File	Full	4190055	2002-05-25	14400	1
File0002	File	Full	1896460	2002-05-26	14400	1
File0003	File	Full	1896460	2002-05-26	14400	1
File0004	File	Full	1896460	2002-05-26	14400	1
File0005	File	Full	1896460	2002-05-26	14400	1
File0006	File	Full	1896460	2002-05-26	14400	1
File0007	File	Purged	1896466	2002-05-26	14400	1

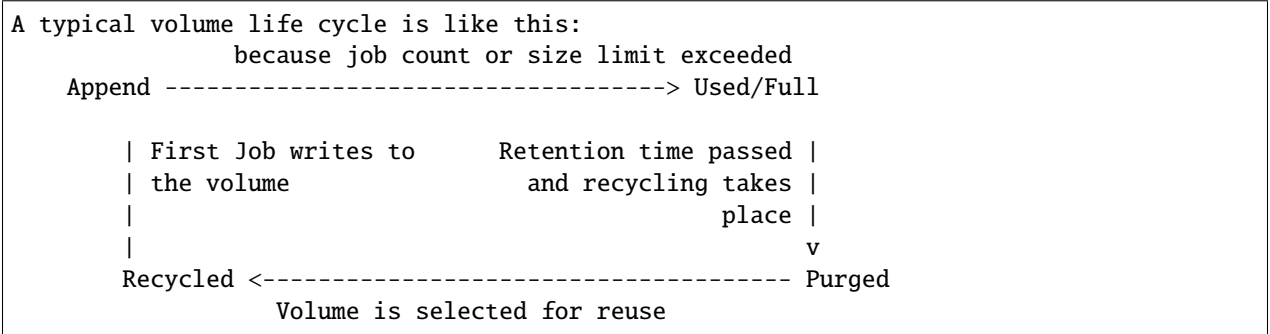
all the volumes are marked as recyclable, and the last Volume, **File0007** has been purged, so it may be immediately recycled. The other volumes are all marked recyclable and when their Volume Retention period (14400 seconds or four hours) expires, they will be eligible for pruning, and possibly recycling. Even though Volume **File0007** has been purged, all the data on the Volume is still recoverable. A purged Volume simply means that there are no entries in the Catalog. Even if the Volume Status is changed to , the data on the Volume will be recoverable. The data is lost only when the Volume is re-labeled and re-written.

To modify Volume **File0001** so that it cannot be recycled, you use the **update volume pool=File** command in the console program, or simply **update** and **Bacula** will prompt you for the information.

VolumeNa	Media	VolSta	VolByte	LastWritte	VolRet	Rec
File0001	File	Full	4190055	2002-05-25	14400	0
File0002	File	Full	1897236	2002-05-26	14400	1
File0003	File	Full	1896460	2002-05-26	14400	1
File0004	File	Full	1896460	2002-05-26	14400	1
File0005	File	Full	1896460	2002-05-26	14400	1
File0006	File	Full	1896460	2002-05-26	14400	1
File0007	File	Purged	1896466	2002-05-26	14400	1

In this case, **File0001** will never be automatically recycled. The same effect can be achieved by setting the Volume Status to Read-Only.

As you have noted, the Volume Status (VolStatus) column in the catalog database contains the current status of the Volume, which is normally maintained automatically by **Bacula**. To give you an idea of some of the values it can take during the life cycle of a Volume, here is a picture of the process.



See also:

Go to [Recycling Algorithm](#).

Go back to the [Volume Recycling](#) chapter.

Go back to the [Storage Space Management](#) chapter.

Go back to the main [Bacula Enterprise](#) chapter.

Manually Recycling Volumes

Although automatic recycling of Volumes is implemented in version 1.20 and later (see the [Automatic Recycling of Volumes](#)), you may want to manually force reuse (recycling) of a Volume.

Assuming that you want to keep the Volume name, but you simply want to write new data on the tape, the steps to take are:

- Use the **update volume** command in the Console to ensure that the field is set to **1**
- Use the **purge jobs volume** command in the Console to mark the Volume as **Purged**. Check by using **list volumes**.

Once the Volume is marked Purged, it will be recycled the next time a Volume is needed.

If you wish to reuse the volume by giving it a new name, follow the following steps:

- Use the **purge jobs volume** command in the Console to mark the Volume as **Purged**. Check by using **list volumes**.
- Use the Console **relabel** command to relabel the Volume.

Please note that the **relabel** command applies only to tape Volumes.

For **Bacula** versions prior to 1.30 or to manually relabel the Volume, use the instructions below:

- Use the **delete volume** command in the Console to delete the Volume from the Catalog.
- If a different tape is mounted, use the **unmount** command, remove the tape, and insert the tape to be renamed.
- Write an EOF mark in the tape using the following commands:

```
mt -f /dev/nst0 rewind
mt -f /dev/nst0 weof
```

where you replace **/dev/nst0** with the appropriate device name on your system.

- Use the **label** command to write a new label to the tape and to enter it in the catalog.

Be aware that the **delete** command can be dangerous. Once it is done, to recover the File records, you must either restore your database as it was before the **delete** command, or use the **bscan** utility program to scan the tape and recreate the database entries.

Specificities of Tape Recycling

Most people will want **Bacula** to fill a tape and when it is full, a new tape will be mounted, and so on. However, as an extreme example, it is possible for **Bacula** to write on a single tape, and every night to rewrite it. To get this to work, you must do two things: first, set the **VolumeRetention** to less than your save period (one day), and the second item is to make **Bacula** mark the tape as full after using it once. This is done using **mt**. If this latter record is not used and the tape is not full after the first time it is written, **Bacula** will simply append to the tape and eventually request another volume. Using the tape only once, forces the tape to be marked after each use, and the next time **Bacula** runs, it will recycle the tape.

An example Pool resource that does this is:

```
Pool {
  Name = DDS-4
  Use Volume Once = yes
  Pool Type = Backup
  AutoPrune = yes
  VolumeRetention = 12h # expire after 12 hours
  Recycle = yes
}
```

Regular Schedule Tape Usage Example

This example is meant to show you how one could define a fixed set of volumes that **Bacula** will rotate through on a regular schedule. There are an infinite number of such schemes, all of which have various advantages and disadvantages. Note: these volumes may either be tape volumes or disk volumes.

We start with the following assumptions:

- A single tape has more than enough capacity to do a full save.
- There are ten tapes that are used on a daily basis for incremental backups. They are pre-labeled Daily1 ... Daily10.
- There are four tapes that are used on a weekly basis for full backups. They are labeled Week1 ... Week4.
- There are 12 tapes that are used on a monthly basis for full backups. They are numbered Month1 ... Month12
- A full backup is done every Saturday evening (tape inserted Friday evening before leaving work).

- No backups are done over the weekend (this is easy to change).
- The first Friday of each month, a Monthly tape is used for the Full backup.
- Incremental backups are done Monday - Friday (actually Tue-Fri mornings).

We start the system by doing a Full save to one of the weekly volumes or one of the monthly volumes. The next morning, we remove the tape and insert a Daily tape. Friday evening, we remove the Daily tape and insert the next tape in the Weekly series. Monday, we remove the Weekly tape and re-insert the Daily tape. On the first Friday of the next month, we insert the next Monthly tape in the series rather than a Weekly tape, then continue. When a Daily tape finally fills up, **Bacula** will request the next one in the series, and the next day when you notice the email message, you will mount it and **Bacula** will finish the unfinished incremental backup.

What does this give? Well, at any point, you will have the last complete Full save plus several Incremental saves. For any given file you want to recover (or your whole system), you will have a copy of that file every day for at least the last 14 days. For older versions, you will have at least three and probably four Friday full saves of that file, and going back further, you will have a copy of that file made on the beginning of the month for at least a year.

So you have copies of any file (or your whole system) for at least a year, but as you go back in time, the time between copies increases from daily to weekly to monthly.

What would the **Bacula** configuration look like to implement such a scheme?

```
Schedule {
    Name = "NightlySave"
    Run = Level=Full Pool=Monthly 1st sat at 03:05
    Run = Level=Full Pool=Weekly 2nd-5th sat at 03:05
    Run = Level=Incremental Pool=Daily tue-fri at 03:05
}
Job {
    Name = "NightlySave"
    Type = Backup
    Level = Full
    Client = LocalMachine
    FileSet = "File Set"
    Messages = Standard
    Storage = DDS-4
    Pool = Daily
    Schedule = "NightlySave"
}
# Definition of file storage device
Storage {
    Name = DDS-4
    Address = localhost
    SDPort = 9103
    Password = XXXXXXXXXXXXXXXX
    Device = FileStorage
    Media Type = 8mm
}
FileSet {
    Name = "File Set"
    Include {
        Options { signature=MD5 }
        File = ffffffffffffffffffff
    }
    Exclude { File=*.o }
}
```

(continues on next page)

```

}
Pool {
    Name = Daily
    Pool Type = Backup
    AutoPrune = yes
    VolumeRetention = 10d    # recycle in 10 days
    Maximum Volumes = 10
    Recycle = yes
}
Pool {
    Name = Weekly
    Use Volume Once = yes
    Pool Type = Backup
    AutoPrune = yes
    VolumeRetention = 30d # recycle in 30 days (default)
    Recycle = yes
}
Pool {
    Name = Monthly
    Use Volume Once = yes
    Pool Type = Backup
    AutoPrune = yes
    VolumeRetention = 365d # recycle in 1 year
    Recycle = yes
}

```

See also:

Go to:

- [Pruning](#)
- [Pruning and Recycling Example](#)

Go back to the [Storage Space Management](#) chapter.Go back to the [main Bacula Enterprise Management](#) page.

4.2 Pruning

As **Bacula** writes files to a Volume, it keeps a list of files, jobs, and volumes in a database called the Catalog. Among other things, the database helps **Bacula** to decide which files to back up in an incremental or differential backup, and helps you locate files on past backups when you want to restore something. However, the Catalog will grow larger and larger as time goes on, and eventually it can become unacceptably large.

Bacula's process for removing entries from the Catalog is called Pruning.

The default is Automatic Pruning, which means that once a Job record reaches a certain age (e.g. 30 days old) and a pruning occurs, it will be removed from the Catalog. Note that Job records that are required for current restore won't be removed automatically, and File records are needed for VirtualFull and Accurate backups. Once a job has been pruned, you can still restore it from the backup Volume, provided that the Volume has not been recycled, but one additional step is required: scanning the volume with **bscan**. The alternative to Automatic Pruning is Manual Pruning, in which you explicitly tell **Bacula** to erase the catalog entries for a volume. You'd usually do this when you want to reuse a **Bacula** volume, because there's no point in keeping a list of files that **USED TO BE** on a volume. Or, if the catalog is starting to get too big, you could prune the oldest jobs to save space. Manual pruning is done with the `prune` command.

For many deployments, Automatic Pruning is usually fine and does not interfere with normal operations. However, with larger deployments with several millions of files, pruning can take a long time, and can possibly hold up other Jobs.

Additionally, there may be race conditions when Volume pruning is done while backup Jobs are running. This is especially true in large environments, and even more so if Volume truncation is being performed.

The recommended solution for medium and large environment (more than 50 Clients and millions of files) is to *disable automatic pruning* everywhere, and perform pruning at a quiet time outside of the backup window. The customer has to create a specific Director's Storage resource pointing to a free SD device used in the Autochanger (or all the devices). The requirement is to have this specific device pointing back to Autochanger it belongs to by using an Autochanger directive.

There are three main resources, on the Director side, where automatic pruning is configured in Bacula:

- the *Job{}* resource
- the *Client{}* resource
- the *Pool{}* resource

For a *Job{}* resource, there are three directives used to configure automatic prune:

- **Prune Jobs** = <yes|no> (default is **no**)
- **Prune Files** = <yes|no> (default is **no**)
- **Prune Volumes** = <yes|no> (default is **no**)

For *Client{}* and *Job{}* resources, there is only one directive used to configure automatic prune:

- **AutoPrune** = <yes|no> (default is **yes**)

If automatic pruning is disabled in these resources (**AutoPrune = no**) and the directives **Prune Jobs**, **Prune Files** or **Prune Volumes** are set to **yes** in a *Job{}* resource, it will override the settings in the *Client{}* resource (for Jobs and Files) and in the *Pool{}* resource (for Volumes).

Read more:

Disabling Automatic Pruning

In order to disable automatic pruning of Jobs and files, you should:

- In *Job{}* resources, avoid the usage of **Prune Jobs = yes** and **Prune Files = yes**.
- In all *Client{}* resources, configure **AutoPrune = no**.

In order to disable automatic pruning of volumes, you should:

- In the *Job{}* resources, avoid the usage of **Prune Volumes = yes**.
- In all *Pool{}* resources, configure **AutoPrune = no**.

Once your Jobs, Clients and Pools have all been edited, you must tell the Director to reload its configuration.

Run the following command to validate your modifications

```
/opt/bacula/bin/bacula-dir -t -u bacula -g bacula
```

Then reload the configuration from **bconsole** with:

```
reload
```

After disabling automatic pruning for Jobs, files and Volumes in all the defined resources, Bacula will no longer perform any automatic pruning. As you might suspect, with no automatic pruning, the Catalog database would grow forever, so we need to perform pruning manually or in a more automated way by using an *Admin* type Job. *Please contact support@baculasystems.com for manual ways to prune or automatic pruning jobs and scripts.*

Should you wish to have volumes pruned for specific pools, you can do so by using separated `prune` commands for the different pools, such as the following:

```
echo "prune expired volume pool=Pool_1 yes" | /opt/bacula/bin/bconsole > /tmp/prunePool_
↪1.log
echo "prune expired volume pool=Pool_2 yes" | /opt/bacula/bin/bconsole > /tmp/prunePool_
↪2.log
echo "prune expired volume pool=Pool_3 yes" | /opt/bacula/bin/bconsole > /tmp/prunePool_
↪3.log
```

See also:

Go to *Pruning Directives*.

Go back to the *Pruning* chapter.

Go back to the *Storage Space Management* chapter.

Go back to the main *Bacula Enterprise* chapter.

Pruning Directives

There are three pruning durations. All apply to catalog database Jobs and Files records and not to the actual data in a Volume. The pruning (or retention) durations are for: Volumes (Jobs and Files records in the Volume), Jobs (Job records), and Files (File records). The durations inter-depend a bit because if **Bacula** prunes a Volume, it automatically removes all the Job records, and, consequently, all the File records. Also when a Job record is pruned, all the File records for that Job are also pruned (deleted) from the catalog.

Having the File records in the database means that you can examine all the files backed up for a particular Job. They take the most space in the catalog (probably 90-95% of the total). When the File records are pruned, the Job records can remain, and you can still examine what Jobs ran, but not the details of the Files backed up. In addition, without the File records, you cannot use the Console **restore** command to restore specific files.

When a Job record is pruned, the Volume (Media record) for that Job can still remain in the database, and if you do a **list volumes**, you will see the volume information, but the Job records (and its File records) will no longer be available.

In each case, pruning removes information about where older files are, but it also prevents the catalog from growing to be too large. You choose the retention periods in function of how many files you are backing up and the time periods you want to keep those records online, and the size of the database. You can always re-insert the records (with 98% of the original data) by using **bscan** to scan in a whole Volume or any part of the volume that you want.

By setting **AutoPrune** to **yes** you will permit **Bacula** to automatically prune Volumes in the Pool when a Job needs a Volume. Volume pruning means removing Jobs and Files records from the catalog. It does not shrink the size of the Volume or affect the Volume data until the Volume gets overwritten. When a Job requests a volume and there are no Volumes with Volume Status `,` or available, **Bacula** will begin volume pruning. This means that all Jobs that are older than the **VolumeRetention** period will be pruned from every Volume that has Volume Status or and has Recycle set to **yes**. Pruning consists of deleting the corresponding Job, File, and JobMedia records from the catalog database. No change to the physical data on the Volume occurs during the pruning process. When all files are pruned from a Volume (i.e. no records in the catalog), the Volume will be marked as implying that no Jobs remain on the volume. The Pool records that control the pruning are described below.

AutoPrune = `<yes|no>` If AutoPrune is set to **yes** **(default)**, **Bacula** will automatically apply the Volume retention period when running a Job and it needs a new Volume but no appendable volumes are available. At that point, **Bacula**

will prune all Volumes that can be pruned (i.e. AutoPrune set) in an attempt to find a usable volume. If during the autopruning, all files are pruned from the Volume, it will be marked with VolStatus . The default is **yes**. Note, that although the File and Job records may be pruned from the catalog, a Volume will be marked Purged (and hence ready for recycling) if the Volume status is Append, Full, Used, or Error. If the Volume has another status, such as Archive, Read-Only, Disabled, Busy, or Cleaning, the Volume status will not be changed to Purged.

Volume Retention = <time-period-specification> The Volume Retention record defines the length of time that **Bacula** will guarantee that the Volume is not reused (i.e. recycled). As long as a Volume has the status **Append**, it will not be recycled even though the retention period may have expired, until the status is changed to some other status such as Full, Used, Purged, . . .

The retention period when applied starts at the time of the last write to the Volume. For a Volume to be recycled, the volume must not have a status of **Append** and the retention period must have expired. Even in that case, the Volume will not be recycled if any other appendable Volume is available for the Job to complete. As noted above, the Volume status will be marked **Purged** by Bacula prior to it being recycled.

Note, when all the Job records that are on the Volume have been removed, the Volume will be marked (i.e. it has no more valid Jobs stored on it), and the Volume may be recycled even if the **Volume Retention** period has not expired.

When this time period expires, and if **AutoPrune** is set to **yes**, and a new Volume is needed, but no appendable Volume is available, **Bacula** will prune (remove) Job records that are older than the specified Volume Retention period even if the Job or File retention has not been reached. Normally this should not happen since the Volume Retention period should always be set greater than the Job Retention period, which should be greater than the File Retention period.

The Volume Retention period takes precedence over any Job Retention period you have specified in the Client resource. It should also be noted, that the Volume Retention period is obtained by reading the Catalog Database Media record rather than the Pool resource record. This means that if you change the VolumeRetention in the Pool resource record (in bacula-dir.conf), you must ensure that the corresponding change is made in the catalog by using the **update pool** command. Doing so will insure that any new Volumes will be created with the changed Volume Retention period. Any existing Volumes will have their own copy of the Volume Retention period that can only be changed on a Volume by Volume basis using the **update volume** command.

When all Job catalog entries are removed from the volume, its VolStatus is set to . The files remain physically on the Volume until the volume is overwritten.

Retention periods are specified in seconds, minutes, hours, days, weeks, months, quarters, or years on the record. See the Configuration chapter of this manual for additional details of time specification.

The default Volume Retention period is 1 year.

Recycle = <yes|no> This statement tells **Bacula** whether or not the particular Volume can be recycled (i.e. rewritten). If Recycle is set to **no**, then even if **Bacula** prunes all the Jobs on the volume and it is marked , it will not consider the volume for recycling. If Recycle is set to **yes** (the default) and all Jobs have been pruned, the volume status will be set to and the volume may then be reused when another volume is needed. If the volume is reused, it is relabeled with the same Volume Name, however all previous data will be lost.

It is also possible to “force” pruning of all Volumes in the Pool associated with a Job by adding **Prune Files=yes** to the Job resource.

See also:

Go to *Disabling Automatic Pruning*.

Go back to the *Pruning* chapter.

Go back to the *Storage Space Management* chapter.

Go back to the main *Bacula Enterprise* chapter.

Notes

- When a volume is pruned, all the Jobs and files associated with this volume are pruned as well. When a Job is pruned, all the files associated with the Job are also pruned. If you have your Job and file retentions greater than your Volume retention, your Jobs and files will be pruned as soon as the Volumes they are stored on is pruned.
- The pruning process is closely related to retention periods configuration. Bacula will only prune Jobs, files or Volumes whose their retention time has expired. There is one exception: a Volume can be pruned and reused if there are no more Jobs and files associated to it in the Catalog, even if the volume retention time has not expired. This is usually desirable, provided job retention times match actual requirements. Please check `BEPPSchedulesAndRetentions` section of this guide for retention period recommendations.
- If a single `Pool{}` or `Client{}` resource has an `Autoprune = yes` set, automatic pruning is activated.

See also:

Go back to:

- [Volume Recycling](#)

Go to:

- [Pruning and Recycling Example](#)

Go back to the [Storage Space Management](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

4.3 Pruning and Recycling Example

Perhaps the best way to understand the various resource records that come into play during automatic pruning and recycling is to run a Job that goes through the whole cycle. If you add the following resources to your Director's configuration file:

```
Schedule {
  Name = "30 minute cycle"
  Run = Level=Full Pool=File Messages=Standard Storage=File
  hourly at 0:05
  Run = Level=Full Pool=File Messages=Standard Storage=File
  hourly at 0:35
}
Job {
  Name = "Filetest"
  Type = Backup
  Level = Full
  Client=XXXXXXXXXX
  FileSet="Test Files"
  Messages = Standard
  Storage = File
  Pool = File
  Schedule = "30 minute cycle"
}
# Definition of file storage device
Storage {
  Name = File
  Address = XXXXXXXXXXXX
  SDPort = 9103
```

(continues on next page)

```

Password = XXXXXXXXXXXXXXXX
Device = FileStorage
Media Type = File
}
FileSet {
  Name = "File Set"
  Include {
    Options { signature=MD5 }
    File = ffffffff
  }
  Exclude { File=*.o }
}
Pool {
  Name = File
  Use Volume Once = yes
  Pool Type = Backup
  LabelFormat = "File"
  AutoPrune = yes
  VolumeRetention = 4h
  Maximum Volumes = 12
  Recycle = yes
}

```

Where you will need to replace the **fffffff**'s by the appropriate files to be saved for your configuration. For the FileSet Include, choose a directory that has one or two megabytes maximum since there will probably be approximately eight copies of the directory that **Bacula** will cycle through.

In addition, you will need to add the following to your Storage daemon's configuration file:

```

Device {
  Name = FileStorage
  Media Type = File
  Archive Device = /tmp
  LabelMedia = yes;
  Random Access = Yes;
  AutomaticMount = yes;
  RemovableMedia = no;
  AlwaysOpen = no;
}

```

With the above resources, **Bacula** will start a Job every half hour that saves a copy of the directory you chose to /tmp/File0001 ... /tmp/File0012. After 4 hours, **Bacula** will start recycling the backup Volumes (/tmp/File0001 ...). You should see this happening in the output produced. **Bacula** will automatically create the Volumes (Files) the first time it uses them.

To turn it off, either delete all the resources you've added, or simply comment out the record in the resource.

See also:

Go back to:

- [Volume Recycling](#)
- [Pruning](#)

Go back to the [Storage Space Management](#) chapter.

Go back to the *main Bacula Enterprise Management page*.

See also:

Go back to:

- *Daily Management*
- *Logs Management*
- *Jobs Monitoring*

Go to:

- *PostgreSQL Catalog Administration*
- *Tape/Volume Management*
- *Tape Autochanger Usage*

Go back to the *main Bacula Enterprise Management page*.

5 PostgreSQL Catalog Administration

5.1 Catalog Backup

The default backup script is doing an SQL dump. You can change this to other more advanced backup methods supported by PostgreSQL, see <https://www.postgresql.org/docs/current/continuous-archiving.html> for further suggestions. Make sure to test the full Catalog restore from time to time.

The `pg_dump` tool can run concurrently with backup jobs but we advise starting the Catalog Backup job when no other jobs are running. There are two reasons for this. Running jobs have a specific status set in the Catalog, and when Catalog is restored Director checks it and marks those jobs as failed. Secondly, jobs running while the Catalog is backed up will continue writing to the volumes, and the Catalog dump will not hold all the information backup volumes keep. *Bscan* is a tool one can use to recreate Catalog records from backup volumes.

5.2 Cleaning up Catalog Database

There are two tables in the Catalog database that occupy most of the space. The File table stores records of all the files that were backed up, and the Path table stores records of all the directory paths. The Path table is shared between all clients and jobs, while the File table will store all files backed up by each backup job.

When files are stored or read the Path table is loaded into memory, and if it grows too much it may affect the performance of the database server.

Usually, Bacula will prune the records from the Catalog when backup jobs are out of retention. Pruning jobs and related files can leave orphaned records. To avoid having too many orphans it is advised to use the `dbcheck` program periodically to clean the Catalog.

5.3 Vacuum and Analyze

When a record is deleted or updated PostgreSQL doesn't remove the old row from the table but marks it as a dead tuple. Daily operations and orphan removal specifically leave dead tuples in a table. For the database to reuse the table rows occupied by dead tuples it is required to vacuum a table. Regular vacuum ensures the database files do not grow out of proportion keeping dead tuples on disk.

Analyze collects the statistics about column values in tables. These statistics are used to find the best query plan. Rows manipulation changes statistics so regular analyze ensures optimal query performance.

Usually, it would be sufficient to let autovacuum and autoanalyze to deal with vacuuming and statistics. But for large tables, the default settings may be inadequate. Automatic vacuum is controlled by two configuration options with the below defaults.

```
autovacuum_vacuum_threshold = 50
autovacuum_vacuum_scale_factor = 0.2
```

The autovacuum is triggered when:

```
pg_stat_user_tables.n_dead_tup >
(pg_class.reltuples x autovacuum_vacuum_scale_factor)
+ autovacuum_vacuum_threshold
```

Analyze is run after the autovacuum automatically and is similarly controlled with:

```
autovacuum_analyze_threshold = 50
autovacuum_analyze_scale_factor = 0.1
```

Where `autovacuum_analyze_scale_factor` specifies a fraction of the table that is changed before autoanalyze is triggered.

As the number of files kept in the Catalog grows the gap between subsequent autovacuum/autoanalyze executions grows too which results in longer runtimes and possibly inability to run autovacuum/autoanalyze at all. Then we advise running it manually on daily basis and especially after pruning operations. You can schedule this task with cron or use the following Admin job during non-production hours (jobs can run concurrently).

```
Job {
  Name = "VACUUM"
  Type = Admin
  JobDefs = DefaultJob
  Runscript {
    Command = "vacuumdb -a -z -q"
    RunsOnClient = no
    RunsWhen = Before
  }
}
```

To avoid using `trust` in the `pg_hba.conf`, you may want to use the `.pgpass` file (with `0600` permission bits) in the bacula user's home directory.

```
% ls -l /opt/bacula/.pgpass
-rw----- 1 bacula 201 32 Feb  3 10:19 /opt/bacula/.pgpass

% cat /opt/bacula/.pgpass
*:*:bacula:bacula:yourpassword
```

Note: For large Catalogs, we also advise not to rely on automatic pruning but to schedule maintenance jobs outside of the backup window to deal with records pruning. Please see *Pruning* for more information. If you run pruning and vacuum in a single Admin job, please take note Bacula cannot guarantee the order of execution for RunScript sections. If there is a need to execute tasks in a set order, please use a shell script and run the script from the Admin job.

A different approach would be to tweak settings for the File table to ensure vacuum and analyze is run automatically. Eliminating `autovacuum_vacuum_scale_factor` from the equation above allows triggering autovacuum as soon as `autovacuum_vacuum_threshold` is crossed. It is important to set the `autovacuum_vacuum_threshold` to a value that would enable vacuum runs automatically after pruning.

For example, if 1 million file records are pruned or updated daily the following settings seem appropriate.

```
ALTER TABLE file SET (autovacuum_vacuum_scale_factor = 0,
autovacuum_vacuum_threshold = 1000000,
autovacuum_analyze_scale_factor = 0,
autovacuum_analyze_threshold = 250000,
autovacuum_vacuum_cost_delay = 0);
```

See also:

- [dbcheck](#)

See also:

Go back to:

- [Daily Management](#)
- [Logs Management](#)
- [Jobs Monitoring](#)
- [Storage Space Management](#)

Go to:

- [Tape/Volume Management](#)
- [Tape Autochanger Usage](#)

Go back to the [main Bacula Enterprise Management page](#).

6 Tape/Volume Management

This document describes utility programs written to aid **Bacula** users and developers in dealing with Volumes external to **Bacula**.

6.1 Requirements

- Each of the following programs requires a **valid Storage Daemon configuration file** (actually, the only part of the configuration file that these programs need is the **Device** resource definitions). This permits the programs to find the configuration parameters for your archive device (generally a tape drive). By default, they read **bacula-sd.conf** in the current directory, but you may specify a different configuration file using the **-c** option.
- Each of these programs require a **device-name** where the Volume can be found. In the case of a tape, this is the physical device name such as **/dev/nst0** or **/dev/rmt/0ubn** depending on your system. For the program to work, it must find the identical name in the Device resource of the configuration file. See below for specifying Volume names.

Note: If you have **Bacula** running and you want to use one of these programs, you will either need to stop the Storage Daemon, or **unmount** any tape drive you want to use, otherwise the drive will be **busy** because **Bacula** is using it.

- If you are attempting to read or write an archive file rather than a tape, the **device-name** should be the full path to the archive location including the filename. The filename (last part of the specification) will be stripped and used as the Volume name, and the path (first part before the filename) must have the same entry in the configuration file. So, the path is equivalent to the archive device name, and the filename is equivalent to the volume name.
- In general, you must specify the Volume name to each of the programs below (with the exception of **btape**). The best method to do so is to specify a **bootstrap** file on the command line with the **-b** option. As part of the bootstrap file, you will then specify the Volume name or Volume names if more than one volume is needed. For example, suppose you want to read tapes **tape1** and **tape2**. First construct a **bootstrap** file named say, **list.bsr** which contains:

```
Volume=test1|test2
```

where each Volume is separated by a vertical bar. Then simply use:

```
./bls -b list.bsr /dev/nst0
```

In the case of **Bacula** Volumes that are on files, you may simply append volumes as follows:

```
./bls /tmp/test1\|test2
```

where the backslash () was necessary as a shell escape to permit entering the vertical bar (|).

And finally, if you feel that specifying a Volume name is a bit complicated with a bootstrap file, you can use the **-V** option (on all programs except **bcopy**) to specify one or more Volume names separated by the vertical bar (|). For example,

```
./bls -V Vol001 /dev/nst0
```

You may also specify an asterisk (*) to indicate that the program should accept any volume. For example:

```
./bls -V\* /dev/nst0
```

Read more:

bls

bls can be used to do an ls type listing of a **Bacula** tape or file. It is called:

```
Usage: bls [options] <device-name>
-b <file> specify a bootstrap file
-c <file> specify a config file
-d <level> specify debug level
-e <file> exclude list
-i <file> include list
-j list jobs
-k list blocks
(no j or k option) list saved files
-L dump label
-p proceed inspite of errors
-v be verbose
-V specify Volume names (separated by |)
-E Check records to detect errors
-? print this message
```

For example, to list the contents of a tape:

```
./bls -V Volume-name /dev/nst0
```

Or to list the contents of a file:

```
./bls /tmp/Volume-name
```

or

```
./bls -V Volume-name /tmp
```

Note that, in the case of a file, the Volume name becomes the filename, so in the above example, you will replace the **Volume-name** with the name of the volume (file) you wrote.

Normally if no options are specified, bls will produce the equivalent output to the `ls -l` command for each file on the tape. Using other options listed above, it is possible to display only the Job records, only the tape blocks, etc. For example:

```
./bls /tmp/File002
bls: butil.c:148 Using device: /tmp
drwxrwxr-x 3 k k 4096 02-10-19 21:08 /home/kern/bacula/k/src/dird/
drwxrwxr-x 2 k k 4096 02-10-10 18:59 /home/kern/bacula/k/src/dird/CVS/
-rw-rw-r-- 1 k k 54 02-07-06 18:02 /home/kern/bacula/k/src/dird/CVS/Root
-rw-rw-r-- 1 k k 16 02-07-06 18:02 /home/kern/bacula/k/src/dird/CVS/Repository
-rw-rw-r-- 1 k k 1783 02-10-10 18:59 /home/kern/bacula/k/src/dird/CVS/Entries
-rw-rw-r-- 1 k k 97506 02-10-18 21:07 /home/kern/bacula/k/src/dird/Makefile
-rw-r--r-- 1 k k 3513 02-10-18 21:02 /home/kern/bacula/k/src/dird/Makefile.in
-rw-rw-r-- 1 k k 4669 02-07-06 18:02 /home/kern/bacula/k/src/dird/README-config
-rw-r--r-- 1 k k 4391 02-09-14 16:51 /home/kern/bacula/k/src/dird/authenticate.c
-rw-r--r-- 1 k k 3609 02-07-07 16:41 /home/kern/bacula/k/src/dird/autoprunce.c
-rw-rw-r-- 1 k k 4418 02-10-18 21:03 /home/kern/bacula/k/src/dird/bacula-dir.conf
...
-rw-rw-r-- 1 k k 83 02-08-31 19:19 /home/kern/bacula/k/src/dird/.cvsignore
```

(continues on next page)

```
bls: Got EOF on device /tmp
84 files found.
```

Listing Jobs

If you are listing a Volume to determine what Jobs to restore, normally the **-j** option provides you with most of what you will need as long as you don't have multiple clients. For example,

```
./bls -j -V Test1 -c stored.conf DDS-4
bls: butil.c:258 Using device: "DDS-4" for reading.
11-Jul 11:54 bls: Ready to read from volume "Test1" on device "DDS-4" (/dev/nst0).
Volume Record: File:blk=0:1 SessId=4 SessTime=1121074625 JobId=0 DataLen=165
Begin Job Session Record: File:blk=0:2 SessId=4 SessTime=1121074625 JobId=1 Level=F
↳Type=B
Begin Job Session Record: File:blk=0:3 SessId=5 SessTime=1121074625 JobId=5 Level=F
↳Type=B
Begin Job Session Record: File:blk=0:6 SessId=3 SessTime=1121074625 JobId=2 Level=F
↳Type=B
Begin Job Session Record: File:blk=0:13 SessId=2 SessTime=1121074625 JobId=4 Level=F
↳Type=B
End Job Session Record: File:blk=0:99 SessId=3 SessTime=1121074625 JobId=2 Level=F Type=B
Files=168 Bytes=1,732,978 Errors=0 Status=T
End Job Session Record: File:blk=0:101 SessId=2 SessTime=1121074625 JobId=4 Level=F
↳Type=B
Files=168 Bytes=1,732,978 Errors=0 Status=T
End Job Session Record: File:blk=0:108 SessId=5 SessTime=1121074625 JobId=5 Level=F
↳Type=B
Files=168 Bytes=1,732,978 Errors=0 Status=T
End Job Session Record: File:blk=0:109 SessId=4 SessTime=1121074625 JobId=1 Level=F
↳Type=B
Files=168 Bytes=1,732,978 Errors=0 Status=T
11-Jul 11:54 bls: End of Volume at file 1 on device "DDS-4" (/dev/nst0), Volume "Test1"
11-Jul 11:54 bls: End of all volumes.
```

shows a full save followed by two incremental saves.

Adding the **-v** option will display virtually all information that is available for each record:

Listing Blocks

Normally, except for debugging purposes, you will not need to list **Bacula** blocks (the “primitive” unit of **Bacula** data on the Volume). However, you can do so with:

```
./bls -k /tmp/File002
bls: butil.c:148 Using device: /tmp
Block: 1 size=64512
Block: 2 size=64512
...
Block: 65 size=64512
Block: 66 size=19195
```

(continues on next page)

(continued from previous page)

```
bls: Got EOF on device /tmp
End of File on device
```

By adding the `-v` option, you can get more information, which can be useful in knowing what sessions were written to the volume:

```
./bls -k -v /tmp/File002
Volume Label:
Id : Bacula 0.9 mortal
VerNo : 10
VolName : File002
PrevVolName :
VolFile : 0
LabelType : VOL_LABEL
LabelSize : 147
PoolName : Default
MediaType : File
PoolType : Backup
HostName :
Date label written: 2002-10-19 at 21:16
Block: 1 blen=64512 First rec FI=VOL_LABEL SessId=1 SessTim=1035062102 Strm=0 rlen=147
Block: 2 blen=64512 First rec FI=6 SessId=1 SessTim=1035062102 Strm=DATA rlen=4087
Block: 3 blen=64512 First rec FI=12 SessId=1 SessTim=1035062102 Strm=DATA rlen=5902
Block: 4 blen=64512 First rec FI=19 SessId=1 SessTim=1035062102 Strm=DATA rlen=28382
...
Block: 65 blen=64512 First rec FI=83 SessId=1 SessTim=1035062102 Strm=DATA rlen=1873
Block: 66 blen=19195 First rec FI=83 SessId=1 SessTim=1035062102 Strm=DATA rlen=2973
bls: Got EOF on device /tmp
End of File on device
```

Armed with the SessionId and the SessionTime, you can extract just about anything.

If you want to know even more, add a second `-v` to the command line to get a dump of every record in every block.

```
./bls -k -v -v /tmp/File002
bls: block.c:79 Dump block 80f8ad0: size=64512 BlkNum=1
Hdrcksum=b1bdfd6d cksum=b1bdfd6d
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=VOL_LABEL Strm=0 len=147 p=80f8b40
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=SOS_LABEL Strm=-7 len=122 p=80f8be7
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=1 Strm=UATTR len=86 p=80f8c75
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=2 Strm=UATTR len=90 p=80f8cdf
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=3 Strm=UATTR len=92 p=80f8d4d
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=3 Strm=DATA len=54 p=80f8dbd
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=3 Strm=MD5 len=16 p=80f8e07bls: block.c:92
↪Rec: VId=1 VT=1035062102 FI=4 Strm=UATTR len=98 p=80f8e2b
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=4 Strm=DATA len=16 p=80f8ea1
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=4 Strm=MD5 len=16 p=80f8ec5
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=5 Strm=UATTR len=96 p=80f8ee9
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=5 Strm=DATA len=1783 p=80f8f5d
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=5 Strm=MD5 len=16 p=80f9668
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=6 Strm=UATTR len=95 p=80f968c
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=6 Strm=DATA len=32768 p=80f96ff
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=6 Strm=DATA len=32768 p=8101713
```

(continues on next page)

(continued from previous page)

```
bls: block.c:79 Dump block 80f8ad0: size=64512 BlkNum=2
Hdrcksum=9acc1e7f cksum=9acc1e7f
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=6 Strm=contDATA len=4087 p=80f8b40
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=6 Strm=DATA len=31970 p=80f9b4b
bls: block.c:92 Rec: VId=1 VT=1035062102 FI=6 Strm=MD5 len=16 p=8101841
...
```

Normally, except for deep debugging purposes, you will not need to verify **Bacula** checksum blocks. However, you can do so with:

```
./bls -j -E -V File002 FileStorage
bls: butil.c:297-0 Using device: "FileChgr1-Dev2" for reading.
14-Feb 09:43 bls JobId 0: Ready to read from volume "TestVolume001" on file device
↪ "FileChgr1-Dev2" (/tmp/regress/tmp).
Volume Record: File:blk=0:215 SessId=1 SessTime=1423731064 JobId=0 DataLen=180
14-Feb 09:43 bls JobId 0: Error: block_util.c:451 Volume data error at 0:215!
Block checksum mismatch in block=1 len=64512: calc=34589ccc blk=8d8c5921
```

See also:

Go to:

- [btape](#)
- [bscan](#)
- [bcopy](#)
- [bextract](#)

Go back to the [Tape/Volume Management](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

btape

This program permits a number of elementary tape operations via a tty command interface. It works only with tapes and not with other kinds of **Bacula** storage media (DVD, File, etc). The **test** command, described below, can be very useful for testing older tape drive compatibility problems. Aside from initial testing of tape drive compatibility with **Bacula**, **btape** will be mostly used by developers writing new tape drivers.

btape can be dangerous to use with existing **Bacula** tapes because it will relabel a tape or write on the tape if so requested regardless that the tape may contain valuable data, so please be careful and use it only on blank tapes.

To work properly, **btape** needs to read the Storage daemon's configuration file. As a default, it will look for **bacula-sd.conf** in the current directory. If your configuration file is elsewhere, please use the **-c** option to specify where.

The physical device name must be specified on the command line, and this same device name must be present in the Storage daemon's configuration file read by **btape**

```
Usage: btape <options> <device_name>
-b <file> specify bootstrap file
-c <file> set configuration file to file
-d <nn> set debug level to nn
-p proceed inspite of I/O errors
-s turn off signals
```

(continues on next page)

```
-v be verbose
-? print this message.
```

Using btape to Verify Tape Drive

An important reason for this program is to ensure that a Storage daemon configuration file is defined so that **Bacula** will correctly read and write tapes.

It is highly recommended that you run the `test` command before running your first **Bacula** job to ensure that the parameters you have defined for your storage device (tape drive) will permit **Bacula** to function properly. You only need to mount a blank tape, enter the command, and the output should be reasonably self explanatory.

btape Commands

The full list of commands are:

Command	Description
=====	=====
autochanger test	autochanger
bsf	backspace file
bsr	backspace record
cap	list device capabilities
clear	clear tape errors
eod	go to end of Bacula data for append
eom	go to the physical end of medium
fill	fill tape, write onto second volume
unfill	read filled tape
fsf	forward space a file
fsr	forward space a record
help	print this command
label	write a Bacula label to the tape
load	load a tape
quit	quit btape
rawfill	use write() to fill tape
readlabel	read and print the Bacula tape label
rectest	test record handling functions
rewind	rewind the tape
scan read()	tape block by block to EOT and report
scanblocks	Bacula read block by block to EOT and report
speed	report drive speed
status	print tape status
test	General test Bacula tape functions
weof	write an EOF on the tape
wr	write a single Bacula block
rr	read a single record
qfill	quick fill command

The most useful commands are:

test test writing records and EOF marks and reading them back.

fill completely fill a volume with records, then write a few records on a second volume, and finally, both volumes will be read back. This command writes blocks containing random data, so your drive will not be able to compress the data, and thus it is a good test of the real physical capacity of your tapes.

readlabel read and dump the label on a Bacula tape.

cap list the device capabilities as defined in the configuration file and as perceived by the Storage daemon.

The **readlabel** command can be used to display the details of a Bacula tape label. This can be useful if the physical tape label was lost or damaged.

In the event that you want to relabel a **Bacula**, you can simply use the **label** command which will write over any existing label. However, please note for labeling tapes, we recommend that you use the **label** command in the **Console** program since it will never overwrite a valid **Bacula** tape.

Testing Tape Drive

See `TestTapeDriveDeviceBtapeUtility`.

To determine the best configuration of your tape drive, you can run the new **speed** command available in the **btape** program.

This command can have the following arguments:

file_size=n Specify the Maximum File Size for this test (between 1 and 5GB). This counter is in GB.

nb_file=n Specify the number of file to be written. The amount of data should be greater than your memory (*file_size * nb_file*).

skip_zero This flag permits to skip tests with constant data.

skip_random This flag permits to skip tests with random data.

skip_raw This flag permits to skip tests with raw access.

skip_block This flag permits to skip tests with Bacula block access.

```
*speed file_size=3 skip_raw
btape.c:1078 Test with zero data and bacula block structure.
btape.c:956 Begin writing 3 files of 3.221 GB with blocks of 129024 bytes.
+++++
btape.c:604 Wrote 1 EOF to "Drive-0" (/dev/nst0)
btape.c:406 Volume bytes=3.221 GB. Write rate = 44.128 MB/s
...
btape.c:383 Total Volume bytes=9.664 GB. Total Write rate = 43.531 MB/s
btape.c:1090 Test with random data, should give the minimum throughput.
btape.c:956 Begin writing 3 files of 3.221 GB with blocks of 129024 bytes.
+++++
btape.c:604 Wrote 1 EOF to "Drive-0" (/dev/nst0)
btape.c:406 Volume bytes=3.221 GB. Write rate = 7.271 MB/s
+++++
...
btape.c:383 Total Volume bytes=9.664 GB. Total Write rate = 7.365 MB/s
```

When using compression, the random test will give you the minimum throughput of your drive. The test using constant string will give you the maximum speed of your hardware chain. (cpu, memory, scsi card, cable, drive, tape).

You can change the block size in the Storage Daemon configuration file.

See also:

Go back to:

- [bls](#)

Go to:

- [bscan](#)
- [bcopy](#)
- [bextract](#)

Go back to the [Tape/Volume Management](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

bscan

The `bscan` program can be used to re-create a database (catalog) records from the backup information written to one or more Volumes. This is normally needed only if one or more Volumes have been pruned or purged from your catalog so that the records on the Volume are no longer in the catalog, or for Volumes that you have archived. Note, if you scan in Volumes that were previously purged, you will be able to do restores from those Volumes. However, unless you modify the Job and File retention times for the Jobs that were added by scanning, the next time you run any backup Job with the same name, the records will be pruned again. Since it takes a long time to scan Volumes this can be very frustrating. Recovered volume status is set to Archive and Volume Retention time is set to one month by default for volumes that were restored this way.

With some care, `bscan` can also be used to synchronize your existing catalog with a Volume. Although we have never seen a case of `bscan` damaging a catalog, since `bscan` modifies your catalog, we recommend that you do a simple ASCII backup of your database before running `bscan` just to be sure. See **Compacting Your Database** for the details of making a copy of your database.

`bscan` can also be useful in a disaster recovery situation, after the loss of a hard disk, if you do not have a valid **bootstrap** file for reloading your system, or if a Volume has been recycled but not overwritten, you can use `bscan` to re-create your database, which can then be used to **restore** your system or a file to its previous state.

`bscan` is also able to re-create/list records such as Restore Object and Plugin Objects.

It is called:

```
Usage: bscan [options] <bacula-archive>
-b bootstrap specify a bootstrap file
-c <file> specify configuration file
-d <nn> set debug level to nn
-m update media info in database
-n <name> specify the database name (default bacula)
-u <user> specify database user name (default bacula)
-P <password> specify database password (default none)
-h <host> specify database host (default NULL)
-p proceed inspite of I/O errors
-r list records
-s synchronize or store in database
-v verbose
-V <Volumes> specify Volume names (separated by |)
-w <dir> specify working directory (default from conf file)
-? print this message
```


If you are using *MySQL* or *PostgreSQL*, there is no need to supply a working directory since in that case, `bscan` knows where the databases are. However, if you have provided security on your database, you may need to supply either the database name (`-b` option), the user name (`-u` option), and/or the password (`-p`) options.

Note: before `bscan` can work, it needs at least a bare bones valid database. If your database exists but some records are missing because they were pruned, then you are all set. If your database was lost or destroyed, then you must first ensure that you have the SQL program running (or), then you must create the **Bacula** database (normally named `bacula`), and you must create the **Bacula** tables using the scripts in the `/opt/bacula/scripts/` directory. This is explained in the **Installation** chapter of the manual. Finally, before scanning into an empty database, you must start and stop the Director with the appropriate `bacula-dir.conf` file so that it can create the Client and Storage records which are not stored on the Volumes. Without these records, scanning is unable to connect the Job records to the proper client.

Forgetting for the moment the extra complications of a full rebuild of your catalog, let's suppose that you did a backup to Volumes "Vol001" and "Vol002", then sometime later all records of one or both those Volumes were pruned or purged from the database. By using `bscan` you can recreate the catalog entries for those Volumes and then use the `restore` command in the Console to restore whatever you want. A command something like:

```
bscan -c bacula-sd.conf -v -V Vol001\|Vol002 /dev/nst0
```

will give you an idea of what is going to happen without changing your catalog. Of course, you may need to change the path to the Storage daemon's conf file, the Volume name, and your tape (or disk) device name. This command must read the entire tape, so if it has a lot of data, it may take a long time, and thus you might want to immediately use the command listed below. Note, if you are writing to a disk file, replace the device name with the path to the directory that contains the Volumes. This must correspond to the Archive Device in the conf file.

Then to actually write or store the records in the catalog, add the `-s` option as follows:

```
bscan -s -m -c bacula-sd.conf -v -V Vol001\|Vol002 /dev/nst0
```

When writing to the database, if `bscan` finds existing records, it will generally either update them if something is wrong or leave them alone. Thus if the Volumes you are scanning are all or partially in the catalog already, no harm will be done to that existing data. Any missing data will simply be added.

If you have multiple tapes, you should scan them with:

```
bscan -s -m -c bacula-sd.conf -v -V Vol001\|Vol002\|Vol003 /dev/nst0
```

Since there is a limit on the command line length (511 bytes) accepted by `bscan`, if you have too many Volumes, you will need to manually create a bootstrap file. See the **Restoring Files with the Bootstrap File** of the *Bacula Enterprise Main manual*, in particular the section entitled **Bootstrap for bscan**. Basically, the `.bsr` file for the above example might look like:

```
Volume=Vol001  
Volume=Vol002  
Volume=Vol003
```

Note: `bscan` does not support supplying Volume names on the command line and at the same time in a bootstrap file. Please use only one or the other.

You should, always try to specify the tapes in the order they are written. If you do not, any Jobs that span a volume may not be fully or properly restored. However, `bscan` can handle scanning tapes that are not sequential. Any incomplete records at the end of the tape will simply be ignored in that case. If you are simply repairing an existing catalog, this may be OK, but if you are creating a new catalog from scratch, it will leave your database in an incorrect state. If you do not specify all necessary Volumes on a single `bscan` command, `bscan` will not be able to correctly restore the records

that span two volumes. In other words, it is much better to specify two or three volumes on a single `bscan` command (or in a `.bsr` file) rather than run `bscan` two or three times, each with a single volume.

Note, the restoration process using `bscan` is not identical to the original creation of the catalog data. This is because certain data such as Client records and other non-essential data such as volume reads, volume mounts, etc is not stored on the Volume, and thus is not restored by `bscan`. The results of `bscanning` are, however, perfectly valid, and will permit restoration of any or all the files in the catalog using the normal **Bacula** console commands. If you are starting with an empty catalog and expecting `bscan` to reconstruct it, you may be a bit disappointed, but at a minimum, you must ensure that your `bacula-dir.conf` file is the same as what it previously was – that is, it must contain all the appropriate Client resources so that they will be recreated in your new database **before** running `bscan`. Normally when the Director starts, it will recreate any missing Client records in the catalog. Another problem you will have is that even if the Volumes (Media records) are recreated in the database, they will not have their autochanger status and slots properly set. As a result, you will need to repair that by using the `update slots` command. There may be other considerations as well. Rather than `bscanning`, you should always attempt to recover you previous catalog backup.

Using `bscan` to Compare Volume to Existing Catalog

If you wish to compare the contents of a Volume to an existing catalog without changing the catalog, you can safely do so if and only if you do **not** specify either the `-m` or the `-s` options in order to avoid modifying the catalog. However, not every record that is stored in the catalog is stored on the volumes, so we don't particularly recommend using `bscan` other than for testing. Probably the best way to compare what is on a Volume to what is in the Catalog is to use one of the Verify Job options.

Using `bscan` to Recreate Catalog from Volume

This is the mode for which `bscan` is most useful. You can either `bscan` into a freshly created catalog, or directly into your existing catalog (after having made an ASCII copy as described above). Normally, you should start with a freshly created catalog that contains no data.

Starting with a single Volume named `TestVolume1`, you run a command such as:

```
./bscan -V TestVolume1 -v -s -m -c bacula-sd.conf /dev/nst0
```

If there is more than one volume, simply append it to the first one separating it with a vertical bar. You may need to precede the vertical bar with a forward slash escape the shell – e.g. `TestVolume1\TestVolume2`. The `-v` option was added for verbose output (this can be omitted if desired). The `-s` option that tells `bscan` to store information in the database. The physical device name `/dev/nst0` is specified after all the options.

For example, after having done a full backup of a directory, then two incrementals, I reinitialized the SQLite database as described above, and using the `bootstrap.bsr` file noted above, I entered the following command:

```
./bscan -b bootstrap.bsr -v -s -c bacula-sd.conf /dev/nst0
```

which produced the following output:

```
bscan: bscan.c:182 Using Database: bacula, User: bacula
bscan: bscan.c:673 Created Pool record for Pool: Default
bscan: bscan.c:271 Pool type "Backup" is OK.
bscan: bscan.c:632 Created Media record for Volume: TestVolume1
bscan: bscan.c:298 Media type "DDS-4" is OK.
bscan: bscan.c:307 VOL_LABEL: OK for Volume: TestVolume1
bscan: bscan.c:693 Created Client record for Client: Rufus
bscan: bscan.c:769 Created new JobId=1 record for original JobId=2
```

(continues on next page)

```

bscan: bscan.c:717 Created FileSet record "Kerns Files"
bscan: bscan.c:819 Updated Job termination record for new JobId=1
bscan: bscan.c:905 Created JobMedia record JobId 1, MediaId 1
bscan: Got EOF on device /dev/nst0
bscan: bscan.c:693 Created Client record for Client: Rufus
bscan: bscan.c:769 Created new JobId=2 record for original JobId=3
bscan: bscan.c:708 Fileset "Kerns Files" already exists.
bscan: bscan.c:819 Updated Job termination record for new JobId=2
bscan: bscan.c:905 Created JobMedia record JobId 2, MediaId 1
bscan: Got EOF on device /dev/nst0
bscan: bscan.c:693 Created Client record for Client: Rufus
bscan: bscan.c:769 Created new JobId=3 record for original JobId=4
bscan: bscan.c:708 Fileset "Kerns Files" already exists.
bscan: bscan.c:819 Updated Job termination record for new JobId=3
bscan: bscan.c:905 Created JobMedia record JobId 3, MediaId 1
bscan: Got EOF on device /dev/nst0
bscan: bscan.c:652 Updated Media record at end of Volume: TestVolume1
bscan: bscan.c:428 End of Volume. VolFiles=3 VolBlocks=57 VolBytes=10,027,437

```

The key points to note are that `bscan` prints a line when each major record is created. Due to the volume of output, it does not print a line for each file record unless you supply the `-v` option twice or more on the command line.

In the case of a Job record, the new JobId will not normally be the same as the original Jobid. For example, for the first JobId above, the new JobId is 1, but the original JobId is 2. This is nothing to be concerned about as it is the normal nature of databases. `bscan` will keep everything straight.

Although `bscan` claims that it created a Client record for Client: Rufus three times, it was actually only created the first time. This is normal.

You will also notice that it read an end of file after each Job (Got EOF on device ...). Finally the last line gives the total statistics for the `bscan`.

If you had added a second `-v` option to the command line, **Bacula** would have been even more verbose, dumping virtually all the details of each Job record it encountered.

Now if you start **Bacula** and enter a `list jobs` command to the console program, you will get:

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| JobId | Name      | StartTime          | Type | Lvl | JobFiles | JobBytes | JobStat |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1     | kernsave | 2002-10-07 14:59 | B    | F   | 84       | 4180207 | T       |
| 2     | kernsave | 2002-10-07 15:00 | B    | I   | 15       | 2170314 | T       |
| 3     | kernsave | 2002-10-07 15:01 | B    | I   | 33       | 3662184 | T       |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

which corresponds virtually identically with what the database contained before it was re-initialized and restored with `bscan`. All the Jobs and Files found on the tape are restored including most of the Media record. The Volume (Media) records restored will be marked as **Full** so that they cannot be rewritten without operator intervention.

It should be noted that `bscan` cannot restore a database to the exact condition it was in previously because a lot of the less important information contained in the database is not saved to the tape. Nevertheless, the reconstruction is sufficiently complete, that you can run `restore` against it and get valid results.

An interesting aspect of restoring a catalog backup using `bscan` is that the backup was made while **Bacula** was running and writing to a tape. At the point the backup of the catalog is made, the tape **Bacula** is writing to will have say 10 files on it, but after the catalog backup is made, there will be 11 files on the tape **Bacula** is writing. This there is a difference

between what is contained in the backed up catalog and what is actually on the tape. If after restoring a catalog, you attempt to write on the same tape that was used to backup the catalog, **Bacula** will detect the difference in the number of files registered in the catalog compared to what is on the tape, and will mark the tape in error.

There are two solutions to this problem. The first is possibly the simplest and is to mark the volume as **Used** before doing any backups. The second is to manually correct the number of files listed in the Media record of the catalog. This procedure is documented elsewhere in the manual and involves using the `update volume` command in `bconsole`.

Using `bscan` to Correct Volume File Count

If the Storage daemon crashes during a backup Job, the catalog will not be properly updated for the Volume being used at the time of the crash. This means that the Storage daemon will have written say 20 files on the tape, but the catalog record for the Volume indicates only 19 files.

Bacula refuses to write on a tape that contains a different number of files from what is in the catalog. To correct this situation, you may run a `bscan` with the `-m` option (but **without** the `-s` option) to update only the final Media record for the Volumes read.

After `bscan`

If you use `bscan` to enter the contents of the Volume into an existing catalog, you should be aware that the records you entered may be immediately pruned during the next job, particularly if the Volume is very old or had been previously purged. To avoid this, after running `bscan`, you can manually set the volume status (`VolStatus`) to **Read-Only** by using the `update` command in the catalog. This will allow you to restore from the volume without having it immediately purged. When you have restored and backed up the data, you can reset the `VolStatus` to **Used** and the Volume will be purged from the catalog.

See also:

Go back to:

- [*bls*](#)
- [*btape*](#)

Go to:

- [*bcopy*](#)
- [*bextract*](#)

Go back to the [*Tape/Volume Management*](#) chapter.

Go back to the [*main Bacula Enterprise Management page*](#).

`bcopy`

The `bcopy` program can be used to copy one **Bacula** archive file to another. For example, you may copy a tape to a file, a file to a tape, a file to a file, or a tape to a tape. For tape to tape, you will need two tape drives. (a later version is planned that will buffer it to disk). In the process of making the copy, no record of the information written to the new Volume is stored in the catalog. This means that the new Volume, though it contains valid backup data, cannot be accessed directly from existing catalog entries. If you wish to be able to use the Volume with the Console restore command, for example, you must first `bscan` the new Volume into the catalog.

bcopy Command Options

```
Usage: bcopy [-d debug_level] <input-archive> <output-archive>
  -b bootstrap specify a bootstrap file
  -c <file> specify configuration file
  -dnn set debug level to nn
  -i specify input Volume names (separated by |)
  -o specify output Volume names (separated by |)
  -p proceed inspite of I/O errors
  -v verbose
  -w dir specify working directory (default /tmp)
  -? print this message
```

By using a **bootstrap** file, you can copy parts of a **Bacula** archive file to another archive.

One of the objectives of this program is to be able to recover as much data as possible from a damaged tape. However, the current version does not yet have this feature.

As this is a new program, any feedback on its use would be appreciated. In addition, I only have a single tape drive, so I have never been able to test this program with two tape drives.

See also:

Go back to:

- [bls](#)
- [btape](#)
- [bscan](#)

Go to:

- [bextract](#)

Go back to the [Tape/Volume Management](#) chapter.

Go back to the [main Bacula Enterprise Management page](#).

bextract

Normally, you will restore files by running a **Restore Job** from the **Console** program. However, **bextract** can be used to extract a single file or a list of files from a **Bacula** tape or file. In fact, **bextract** can be a useful tool to restore files to an empty system assuming you are able to boot, you have statically linked **bextract** and you have an appropriate **bootstrap** file.

Some of the current limitations of **bextract** are:

1. It cannot restore access control lists (ACL) that have been backed up along with the file data.
2. It cannot restore encrypted files.
3. The command line length is relatively limited, which means that you cannot enter a huge number of volumes. If you need to enter more volumes than the command line supports, please use a bootstrap file (see below).

It is called:

```
Usage: bextract [-d debug_level] <device-name> <directory-to-store-files>
  -b <file> specify a bootstrap file
  -dnn set debug level to nn
```

(continues on next page)

```

-dt print timestamp in debug output
-T send debug traces to trace file
-e <file> exclude list
-i <file> include list
-p proceed inspite of I/O errors
-t read data from volume, do not write anything
-v verbose
-V specify Volume names (separated by |)
-? print this message

```

where **device-name** is the Archive Device (raw device name or full filename) of the device to be read, and **directory-to-store-files** is a path prefix to prepend to all the files restored.

Note: On Windows systems, if you specify a prefix of say `d:/tmp`, any file that would have been restored to `c:/My Documents` will be restored to `d:/tmp/My Documents`. That is, the original drive specification will be stripped. If no prefix is specified, the file will be restored to the original drive.

Extracting with Include or Exclude Lists

Using the **-e** option, you can specify a file containing a list of files to be excluded. Wildcards can be used in the exclusion list. This option will normally be used in conjunction with the **-i** option (see below). Both the **-e** and the **-i** options may be specified at the same time as the **-b** option. The bootstrap filters will be applied first, then the include list, then the exclude list.

Likewise, and probably more importantly, with the **-i** option, you can specify a file that contains a list (one file per line) of files and directories to include to be restored. The list must contain the full filename with the path. If you specify a path name only, all files and subdirectories of that path will be restored. If you specify a line containing only the filename (e.g. **my-file.txt**) it probably will not be extracted because you have not specified the full path.

For example, if the file **include-list** contains:

```

/home/kern/bacula
/usr/local/bin

```

Then the command:

```

./bextract -i include-list -V Volume /dev/nst0 /tmp

```

will restore from the **Bacula** archive **/dev/nst0** all files and directories in the backup from `/home/kern/bacula` and from `/usr/local/bin`. The restored files will be placed in a file of the original name under the directory `/tmp` (i.e. `/tmp/home/kern/bacula/...` and `/tmp/usr/local/bin/...`).

Extracting With Bootstrap File

The **-b** option is used to specify a **bootstrap** file containing the information needed to restore precisely the files you want. Specifying a **bootstrap** file is optional but recommended because it gives you the most control over which files will be restored. For more details on the **bootstrap** file, please see **Restoring Files with the Bootstrap File** chapter of this document. Note, you may also use a bootstrap file produced by the `restore` command. For example:

```
./bextract -b bootstrap-file /dev/nst0 /tmp
```

The bootstrap file allows detailed specification of what files you want restored (extracted). You may specify a bootstrap file and include and/or exclude files at the same time. The bootstrap conditions will first be applied, and then each file record seen will be compared to the include and exclude lists.

Extracting From Multiple Volumes

If you wish to extract files that span several Volumes, you can specify the Volume names in the bootstrap file or you may specify the Volume names on the command line by separating them with a vertical bar. See the section above under the `bls` program entitled **Listing Multiple Volumes** for more information. The same techniques apply equally well to the `bextract` program or read the **Bootstrap** chapter of this document.

Test Extraction

If you wish to ensure that all blocks of a volume are valid, you can specify the **-t** flag. A bootstrap might be used with this option. It is also possible to setup and use a Verify Job.

See also:

Go back to:

- *bls*
- *btape*
- *bscan*
- *bcopy*

Go back to the *Tape/Volume Management* chapter.

Go back to the *main Bacula Enterprise Management page*.

See also:

Go back to:

- *Daily Management*
- *Logs Management*
- *Jobs Monitoring*
- *Storage Space Management*
- *PostgreSQL Catalog Administration*

Go to:

- *Tape Autochanger Usage*

Go back to the *main Bacula Enterprise Management page*.

7 Tape Autochanger Usage

7.1 Tape Drives: update slot inventory and labeling tapes

Labelling tapes

If you want to use Tape Autochangers, please consider labelling them by barcode. Normally you can add new Volumes into a “Scratch Pool” where other Pools pick them. You can assign them directly to a given Pool if you prefer.

If you have a tape drive that is idle during the label command, it should be chosen automatically. Assuming that your 1st Drive (drive=0) is busy with backups, you can, in this case, use the 2nd drive of your library . If not chosen automatically, you can select the 2nd drive by using “drive=1” on the command line.

Note: The tapes should be labeled as scratch pool tapes. To read more about configuring pools, [click here](#).

Updating Slots Inventory

When changing tape cartridges in your Library by adding or removing Tapes or moving them around you should issue an `update slots` command within **bconsole**. Please be aware that some Libraries might unload a used cartridge to an arbitrary free slot, if the selected slot for unloading is currently occupied.

```
update slots barcode drive=1 storage=Autochanger-LTO
```

After that Bacula knows where the cartridges are, and the barcode sticker of each individual tape has been read. Now you must do the labelling with an additional command. Please use a non busy drive to label your barcodes.

Sample output

Assuming new tapes are in slots 30,29,28,16,14, you should choose a Pool, Storage and Slot(s) (drive), otherwise you will be prompted interactively.

```
label barcode storage=Autochanger-LTO pool=Scratch slots=30,29,28,16,14 drive=1
```

The **bconsole** `status slots` command gives you a list of your current slots assignments.

```
status slots
```

Slot	Volume Name	Status	Media Type	Pool
1*	VTL001	Append	LTO	VTL
2				

See also:

Go back to:

- [Daily Management](#)
- [Logs Management](#)
- [Jobs Monitoring](#)
- [Storage Space Management](#)

- *PostgreSQL Catalog Administration*
- *Tape/Volume Management*

Go back to the *main Bacula Enterprise Management page*.