



HPC Accelerator

Bacula Systems Documentation

Contents

1	Overview	2
2	HPC Accelerator: Scope	3
3	HPC Accelerator: Features	3
4	HPC Accelerator: Architecture	3
5	HPC Accelerator: Installation	4
5.1	Steps	5
5.2	Result	6
6	HPC Accelerator: Configuration	6
7	HPC Accelerator: Operations	8
7.1	Backup	8
7.2	Restore	15
8	HPC Accelerator: Best Practices	18
9	HPC Accelerator: Limitations	18

Contents

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

The following article aims at presenting the reader with information about HPC Accelerator (previously named Automatic Big Filesystem Split Plugin, also referred to as the HPC Accelerator). The document describes what the plugin is, defines the scope of its operations, and presents its main features.

1 Overview

This plugin serves as a Bacula solution for systems that function as large NAS devices, particularly in scenarios where NDMP and Snapshot accelerator technologies are not available.

Also, in the subchapters below, it focuses on more detailed topics.

2 HPC Accelerator: Scope

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

HPC Accelerator is available since **Bacula Enterprise 18.0**.

3 HPC Accelerator: Features

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

The main feature of the **HPC Accelerator** is to offer Full, Incremental and Differential backups with automatic parallelism.

4 HPC Accelerator: Architecture

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

HPC Accelerator is a Bacula File Daemon plugin.

This plugin facilitates the automated parallel backups of large directory trees by initiating several Bacula backup jobs concurrently, thereby distributing the backup workload among all jobs to achieve optimal performance.

Below is a simplified vision of the architecture of this plugin within a generic **Bacula Enterprise** deployment:




Fig. 1: HPC Accelerator Architecture

5 HPC Accelerator: Installation

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

This article describes how to install **HPC Accelerator**.

The installation process involves executing the `install-big-fs.sh` script several times.

- The first execution should be with the `configure` option:

```
root@user:~# /opt/bacula/scripts/install-big-fs.sh configure
```

- The second execution should be with the `install` option:

```
root@user:~# /opt/bacula/scripts/install-big-fs.sh install
```

- At this point, a configuration sample located at `/opt/bacula/big-fs/bacula-dir.conf.sample` is created. This file contains example resource configurations for a Client, Fileset, Job, and restricted Console. These resources should be added to the `/opt/bacula/etc/bacula-dir.conf` individually, or this file may be included in the `bacula-dir.conf` using the `@ include` option, as follows: `@/opt/bacula/big-fs/bacula-dir.conf.sample`

- After reloading the Bacula director's configuration, the `install` script may be optionally executed a third time with the `test` option:

```
root@user:~# /opt/bacula/scripts/install-big-fs.sh test
```

to check if the installation is working correctly.

Note: If the Bacula Director already has a Client resource, the Client in `bacula-dir.conf.sample` should not be copied into the Director's configuration as the Client resource should not be duplicated. However it is important to set the `Maximum Concurrent Jobs = <number>` to a value larger than 1. Ideally this number should be equal to or greater than the number of workers configured on the plugin command line in the Fileset.

5.1 Steps

1. Run:

```
root@big-fs-test:/opt/bacula# scripts/install-big-fs.sh configure

Which unix group of users should have access to the big-fs tools [adm]
changed group of '/opt/bacula/big-fs' from root to adm
mode of '/opt/bacula/big-fs' retained as 0750 (rwxr-x---)

Enter the Bacula Director Name [big-fs-test-dir]:

Enter the Bacula Director Address [big-fs-test]:

Enter the Bacula Director Port [9101]:

Enter the Bacula FileDaemon name [big-fs-test-fd]:

INFO: Creating configuration template for the Director
/opt/bacula/big-fs/bacula-dir.conf.sample will help you to setup
a Job with the Bacula Enterprise Automatic Big Filesystem Split Plugin.

The template can be included in your Director configuration and
you need to review all items marked as "might need to be adjusted"

INFO: Creating configuration template for the Storage daemon
/opt/bacula/big-fs/bacula-sd.conf.sample will help you to setup
storage for the Bacula Enterprise Automatic Big Filesystem Split Plugin.

The template indicate how to setup storage devices for each worker.
This setup allows virtual full backup to be executed.

To do this setup in Bweb
- Go in Bweb, click New, Storage, Virtual Disk Changer, 50 devices, ...
- Add devices to the existing default disk changer DiskAutochanger it_
```

(continues on next page)

(continued from previous page)

```
↪ has 3 devices by default
```

```
root@big-fs-test:/opt/bacula# scripts/install-big-fs.sh install
Installing big-fs plugin
```

2. Update the Director's configuration with the help of the newly generated `/opt/bacula/big-fs/bacula-dir.conf.sample`.
3. Open `bconsole` and run `reload`:
4. Test the setup:

```
root@big-fs-test:/opt/bacula# scripts/install-big-fs.sh test
1000 OK: 10002 big-fs-test-dir Version: 18.0.0 (20 November 2023)
INFO: Connection to the Director OK
INFO: Connection from the Director to the Client OK
INFO: Plugin installed correctly
INFO: Job found on the Director
INFO: Fileset configured on the Director
```

5.2 Result

HPC Accelerator is installed.

6 HPC Accelerator: Configuration

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

The following article presents the configuration of the plugin.

The plugin consists of the following distinct parameters:

1. `workers=<nb_workers>` This parameter indicates the maximum number of parallel backup jobs that are allowed. The default setting is 3.

Note: The number of parallel workers is ultimately limited by the Maximum Concurrent Jobs setting in the Job, Client, Director, Storage, `bigfs-backint.conf` file specified by the `conffile=` plugin option, and other resources within Bacula.

2. `files=<nb_files>` This parameter defines the maximum number of files that can be processed before initiating another worker job. The default setting is 20.
3. `conffile=<path>` When this parameter is set, the plugin will search for its configuration at the specified location rather than the default location `/opt/bacula/big-fs/big-fs-backint.conf`.

4. `ctrlfile=<path>` This parameter indicates the base path for control files. FIFOs and `.ctl` files are created and used for establishing connection and communication with the bacula-fd plugin. The default value is `/tmp/fs-backint`.

The following block is an example of the HPC Accelerator configuration file and its relevant Fileset:

```
Fileset {
    Name = "big-fs-fileset"
    Description = "Automatic Big Filesystem Split FileSet"
    Include {
        File = /etc/
        File = /home/user/
        Options {
            Signature = MD5
        }
        Plugin = "big-fs: workers=5 files=20 conffile=/opt/bacula/etc/big-fs-
↪job.conf ctrlfile=/opt/bacula/tmp/fs-backint"
    }
}
```

Below, there is an example of an HPC Accelerator configuration file: `/opt/bacula/etc/big-fs-job.conf`:

```
job=big-fs-job
client=hostname-fd
bconsole="/opt/bacula/etc/bconsole -n -c /opt/bacula/etc/bconsole.conf"
maximumconcurrentjobs=10
ctrlfile=/opt/bacula/tmp/fs-backint
waitjobcompletion=yes
```

Note:

- The setting `maximumconcurrentjobs` is a Bacula-wide configuration that is necessary for the plugin to operate effectively.
 - The setting `maximumconcurrentjobs` should be set such that it meets the condition `maximumconcurrentjobs >= workers`.
 - When multiple backint jobs are executed on the same system, each job must have its own `fs-backint` configuration file with the relevant `job=<jobname>` in it.
 - When multiple backint jobs are executed on the same system, each job must either have its own restricted console with the proper JobACL and FilesetACL, or add each of these ACLs to the same restricted Console configuration for each additional Job and Fileset pair.
 - To ensure that each job uses its own configuration file, the plugin parameter `conffile=<path>` should be employed.
 - The parameter `ctrlfile=<path>` must be the same in both Fileset and `fs-backint.conf` configurations.
-

7 HPC Accelerator: Operations

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

The following chapter describes details regarding the **HPC Accelerator** operations.

7.1 Backup

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

Perform a backup by running the job via bconsole.

Below is an example of a Fileset and a Job:

Fileset:

```
Fileset {
  Name = "big-fs-plugin"
  Include {
    File=/etc
    File=/home
    Options {
      signature=MD5
    }
    Plugin = "big-fs: workers=5"
  }
}
```

Job:

```
Job {
  Name = "BigFsPluginJob"
  Type = Backup
  Client = host-fd
  FileSet = "big-fs-plugin"
  Storage = File
  Messages = Standard
  Pool = Default
  Maximum Concurrent Jobs = 5
  Write Bootstrap = "/opt/bacula/working_dir/bootstrap.bsr"
  AllowIncompleteJobs = no
}
```

(continues on next page)

(continued from previous page)

```
    Accurate = yes  
}
```

Joblog:

```
run job=job.user-fd.big-fs level=Full  
Using Catalog "MyCatalog"  
Run Backup job  
JobName:   job.user-fd.big-fs  
Level:     Full  
Client:    user-fd  
FileSet:   fs.user-fd.big-fs  
Pool:      Default (From Job resource)  
Storage:   File (From Job resource)  
When:      2025-08-18 11:16:25  
Priority: 10  
OK to run? (Yes/mod/no): yes  
Job queued. JobId=1  
Connecting to Director user:8101  
1000 OK: 10003 user-dir Version: 18.2.0 (22 May 2025)  
Enter a period to cancel a command.  
Automatically selected Client: user-fd  
Connecting to Client user-fd at user:8102  
You have messages.  
wait  
messages  
18-Aug 11:16 user-dir JobId 1: Start Backup JobId 1, Job=job.user-fd.big-fs.  
↪2025-08-18_11.16.25_04  
18-Aug 11:16 user-dir JobId 1: Connected to Storage "File" at user:8103 with_  
↪TLS  
18-Aug 11:16 user-dir JobId 1: Using Device "FileStorage" to write.  
18-Aug 11:16 user-dir JobId 1: Connected to Client "user-fd" at user:8102_  
↪with TLS  
18-Aug 11:16 user-fd JobId 1: Connected to Storage at user:8103 with TLS  
18-Aug 11:16 user-sd JobId 1: Wrote label to prelabeled Volume "TestVolume001  
↪" on File device "FileStorage" (<file stroage>)  
18-Aug 11:16 user-fd JobId 1: big-fs: Controljob started a fs-backint session  
18-Aug 11:16 user-fd JobId 1: Connecting to fs-backint process...  
18-Aug 11:16 user-fd JobId 1: fs-backint process connected.  
18-Aug 11:16 user-dir JobId 2: Start Backup JobId 2, Job=job.user-fd.big-fs.  
↪2025-08-18_11.16.29_06  
18-Aug 11:16 user-dir JobId 2: Using Control JobId 1, Job=job.user-fd.big-fs.  
↪2025-08-18_11.16.25_04  
18-Aug 11:16 user-dir JobId 2: Connected to Storage "File" at user:8103 with_  
↪TLS  
18-Aug 11:16 user-dir JobId 2: Using Device "FileStorage" to write.  
18-Aug 11:16 user-dir JobId 2: Connected to Client "user-fd" at user:8102_  
↪with TLS  
18-Aug 11:16 user-fd JobId 2: Connected to Storage at user:8103 with TLS  
18-Aug 11:16 user-fd JobId 2: Connecting to fs-backint process...  
18-Aug 11:16 user-fd JobId 2: fs-backint process connected.  
18-Aug 11:16 user-dir JobId 3: Start Backup JobId 3, Job=job.user-fd.big-fs.
```

(continues on next page)

(continued from previous page)

```
→2025-08-18_11.16.32_08
18-Aug 11:16 user-dir JobId 3: Using Control JobId 1, Job=job.user-fd.big-fs.
→2025-08-18_11.16.25_04
18-Aug 11:16 user-dir JobId 3: Connected to Storage "File" at user:8103 with
→TLS
18-Aug 11:16 user-dir JobId 3: Using Device "FileStorage" to write.
18-Aug 11:16 user-dir JobId 3: Connected to Client "user-fd" at user:8102
→with TLS
18-Aug 11:16 user-fd JobId 3: Connected to Storage at user:8103 with TLS
18-Aug 11:16 user-fd JobId 3: Connecting to fs-backint process...
18-Aug 11:16 user-fd JobId 3: fs-backint process connected.
18-Aug 11:16 user-dir JobId 4: Start Backup JobId 4, Job=job.user-fd.big-fs.
→2025-08-18_11.16.35_10
18-Aug 11:16 user-dir JobId 4: Using Control JobId 1, Job=job.user-fd.big-fs.
→2025-08-18_11.16.25_04
18-Aug 11:16 user-dir JobId 4: Connected to Storage "File" at user:8103 with
→TLS
18-Aug 11:16 user-dir JobId 4: Using Device "FileStorage" to write.
18-Aug 11:16 user-dir JobId 4: Connected to Client "user-fd" at user:8102
→with TLS
18-Aug 11:16 user-fd JobId 4: Connected to Storage at user:8103 with TLS
18-Aug 11:16 user-fd JobId 4: Connecting to fs-backint process...
18-Aug 11:16 user-fd JobId 4: fs-backint process connected.
18-Aug 11:16 user-dir JobId 5: Start Backup JobId 5, Job=job.user-fd.big-fs.
→2025-08-18_11.16.39_12
18-Aug 11:16 user-dir JobId 5: Using Control JobId 1, Job=job.user-fd.big-fs.
→2025-08-18_11.16.25_04
18-Aug 11:16 user-dir JobId 5: Connected to Storage "File" at user:8103 with
→TLS
18-Aug 11:16 user-dir JobId 5: Using Device "FileStorage" to write.
18-Aug 11:16 user-dir JobId 5: Connected to Client "user-fd" at user:8102
→with TLS
18-Aug 11:16 user-fd JobId 5: Connected to Storage at user:8103 with TLS
18-Aug 11:16 user-fd JobId 5: Connecting to fs-backint process...
18-Aug 11:16 user-fd JobId 5: fs-backint process connected.
18-Aug 11:16 user-sd JobId 3: Elapsed time=00:00:23, Transfer rate=1.026 M
→Bytes/second
18-Aug 11:16 user-sd JobId 2: Elapsed time=00:00:26, Transfer rate=873.9 K
→Bytes/second
18-Aug 11:16 user-sd JobId 4: Elapsed time=00:00:20, Transfer rate=1.147 M
→Bytes/second
18-Aug 11:16 user-sd JobId 5: Elapsed time=00:00:16, Transfer rate=2.074 M
→Bytes/second
18-Aug 11:16 user-sd JobId 3: Sending spooled attrs to the Director.
→Despooling 242,249 bytes ...
18-Aug 11:16 user-sd JobId 2: Sending spooled attrs to the Director.
→Despooling 280,617 bytes ...
18-Aug 11:16 user-sd JobId 4: Sending spooled attrs to the Director.
→Despooling 215,459 bytes ...
18-Aug 11:16 user-sd JobId 5: Sending spooled attrs to the Director.
→Despooling 332,153 bytes ...
18-Aug 11:16 user-dir JobId 3: Bacula user-dir 18.2.0 (22May25):
```

(continues on next page)

(continued from previous page)

Build OS: x86_64-pc-linux-gnu
JobId: 3
Job: job.user-fd.big-fs.2025-08-18_11.16.32_08
Backup Level: Incremental, since=1970-01-01 01:00:01
Client: "user-fd" 18.2.0 (22May25) x86_64-pc-linux-gnu,
FileSet: "fs.user-fd.big-fs" 2025-08-18 11:16:23
Pool: "Default" (From Job resource)
Catalog: "MyCatalog" (From Client resource)
Storage: "File" (From Job resource)
Scheduled time: 18-Aug-2025 11:16:32
Start time: 18-Aug-2025 11:16:35
End time: 18-Aug-2025 11:16:59
Elapsed time: 24 secs
Priority: 10
FD Files Written: 884
SD Files Written: 884
FD Bytes Written: 23,467,218 (23.46 MB)
SD Bytes Written: 23,619,179 (23.61 MB)
Rate: 977.8 KB/s
Software Compression: None
Comm Line Compression: 50.7% 2.0:1
Snapshot/VSS: no
Encryption: no
Accurate: no
Volume name(s): TestVolume001
Volume Session Id: 3
Volume Session Time: 1755508583
Last Volume Bytes: 102,703,523 (102.7 MB)
Non-fatal FD errors: 0
SD Errors: 0
FD termination status: OK
SD termination status: OK
Termination: Backup OK

18-Aug 11:16 user-dir JobId 4: Bacula user-dir 18.2.0 (22May25):

Build OS: x86_64-pc-linux-gnu
JobId: 4
Job: job.user-fd.big-fs.2025-08-18_11.16.35_10
Backup Level: Incremental, since=1970-01-01 01:00:01
Client: "user-fd" 18.2.0 (22May25) x86_64-pc-linux-gnu,,
FileSet: "fs.user-fd.big-fs" 2025-08-18 11:16:23
Pool: "Default" (From Job resource)
Catalog: "MyCatalog" (From Client resource)
Storage: "File" (From Job resource)
Scheduled time: 18-Aug-2025 11:16:35
Start time: 18-Aug-2025 11:16:38
End time: 18-Aug-2025 11:16:59
Elapsed time: 21 secs
Priority: 10
FD Files Written: 790
SD Files Written: 790
FD Bytes Written: 22,822,584 (22.82 MB)

(continues on next page)

(continued from previous page)

SD Bytes Written: 22,958,663 (22.95 MB)
Rate: 1086.8 KB/s
Software Compression: None
Comm Line Compression: 48.8% 2.0:1
Snapshot/VSS: no
Encryption: no
Accurate: no
Volume name(s): TestVolume001
Volume Session Id: 4
Volume Session Time: 1755508583
Last Volume Bytes: 102,703,523 (102.7 MB)
Non-fatal FD errors: 0
SD Errors: 0
FD termination status: OK
SD termination status: OK
Termination: Backup OK

18-Aug 11:16 user-dir JobId 2: Bacula user-dir 18.2.0 (22May25):

Build OS: x86_64-pc-linux-gnu
JobId: 2
Job: job.user-fd.big-fs.2025-08-18_11.16.29_06
Backup Level: Incremental, since=1970-01-01 01:00:01
Client: "user-fd" 18.2.0 (22May25) x86_64-pc-linux-gnu,,
FileSet: "fs.user-fd.big-fs" 2025-08-18 11:16:23
Pool: "Default" (From Job resource)
Catalog: "MyCatalog" (From Client resource)
Storage: "File" (From Job resource)
Scheduled time: 18-Aug-2025 11:16:29
Start time: 18-Aug-2025 11:16:32
End time: 18-Aug-2025 11:16:59
Elapsed time: 27 secs
Priority: 10
FD Files Written: 1,018
SD Files Written: 1,018
FD Bytes Written: 22,544,261 (22.54 MB)
SD Bytes Written: 22,722,278 (22.72 MB)
Rate: 835.0 KB/s
Software Compression: None
Comm Line Compression: 52.8% 2.1:1
Snapshot/VSS: no
Encryption: no
Accurate: no
Volume name(s): TestVolume001
Volume Session Id: 2
Volume Session Time: 1755508583
Last Volume Bytes: 102,703,523 (102.7 MB)
Non-fatal FD errors: 0
SD Errors: 0
FD termination status: OK
SD termination status: OK
Termination: Backup OK

(continues on next page)

(continued from previous page)

```
18-Aug 11:16 user-dir JobId 4: Begin pruning Jobs older than 6 months .
18-Aug 11:16 user-dir JobId 4: No Jobs found to prune.
18-Aug 11:16 user-dir JobId 4: Begin pruning Files.
18-Aug 11:16 user-dir JobId 4: No Files found to prune.
18-Aug 11:16 user-dir JobId 4: End auto prune.
```

```
18-Aug 11:16 user-dir JobId 5: Bacula user-dir 18.2.0 (22May25):
Build OS:          x86_64-pc-linux-gnu
JobId:             5
Job:               job.user-fd.big-fs.2025-08-18_11.16.39_12
Backup Level:      Incremental, since=1970-01-01 01:00:01
Client:            "user-fd" 18.2.0 (22May25) x86_64-pc-linux-gnu,,
FileSet:           "fs.user-fd.big-fs" 2025-08-18 11:16:23
Pool:              "Default" (From Job resource)
Catalog:           "MyCatalog" (From Client resource)
Storage:           "File" (From Job resource)
Scheduled time:    18-Aug-2025 11:16:39
Start time:        18-Aug-2025 11:16:42
End time:          18-Aug-2025 11:16:59
Elapsed time:      17 secs
Priority:           10
FD Files Written:  1,207
SD Files Written:  1,207
FD Bytes Written:  32,979,914 (32.97 MB)
SD Bytes Written:  33,191,053 (33.19 MB)
Rate:              1940.0 KB/s
Software Compression: None
Comm Line Compression: 50.5% 2.0:1
Snapshot/VSS:      no
Encryption:         no
Accurate:           no
Volume name(s):     TestVolume001
Volume Session Id:  5
Volume Session Time: 1755508583
Last Volume Bytes:  102,703,523 (102.7 MB)
Non-fatal FD errors: 0
SD Errors:          0
FD termination status: OK
SD termination status: OK
Termination:        Backup OK
```

```
18-Aug 11:16 user-dir JobId 3: Begin pruning Jobs older than 6 months .
18-Aug 11:16 user-dir JobId 3: No Jobs found to prune.
18-Aug 11:16 user-dir JobId 3: Begin pruning Files.
18-Aug 11:16 user-dir JobId 3: No Files found to prune.
18-Aug 11:16 user-dir JobId 3: End auto prune.
```

```
18-Aug 11:16 user-dir JobId 5: Begin pruning Jobs older than 6 months .
18-Aug 11:16 user-dir JobId 5: No Jobs found to prune.
18-Aug 11:16 user-dir JobId 5: Begin pruning Files.
18-Aug 11:16 user-dir JobId 5: No Files found to prune.
18-Aug 11:16 user-dir JobId 5: End auto prune.
```

(continues on next page)

(continued from previous page)

```
18-Aug 11:16 user-dir JobId 2: Begin pruning Jobs older than 6 months .
18-Aug 11:16 user-dir JobId 2: No Jobs found to prune.
18-Aug 11:16 user-dir JobId 2: Begin pruning Files.
18-Aug 11:16 user-dir JobId 2: No Files found to prune.
18-Aug 11:16 user-dir JobId 2: End auto prune.

18-Aug 11:16 user-fd JobId 1: JobId=2 Job Status is OK
18-Aug 11:16 user-fd JobId 1: JobId=3 Job Status is OK
18-Aug 11:17 user-fd JobId 1: JobId=4 Job Status is OK
18-Aug 11:17 user-fd JobId 1: JobId=5 Job Status is OK
18-Aug 11:17 user-sd JobId 1: Elapsed time=00:00:32, Transfer rate=7 Bytes/
↪second
18-Aug 11:17 user-sd JobId 1: Sending spooled attrs to the Director.↵
↪Despooling 117 bytes ...
18-Aug 11:17 user-dir JobId 1: Bacula user-dir 18.2.0 (22May25):
  Build OS:          x86_64-pc-linux-gnu
  JobId:             1
  Job:               job.user-fd.big-fs.2025-08-18_11.16.25_04
  Backup Level:      Full
  Client:            "user-fd" 18.2.0 (22May25) x86_64-pc-linux-gnu,,
  FileSet:           "fs.user-fd.big-fs" 2025-08-18 11:16:23
  Pool:              "Default" (From Job resource)
  Catalog:           "MyCatalog" (From Client resource)
  Storage:           "File" (From Job resource)
  Scheduled time:    18-Aug-2025 11:16:25
  Start time:        18-Aug-2025 11:16:27
  End time:          18-Aug-2025 11:17:00
  Elapsed time:      33 secs
  Priority:           10
  FD Files Written:   1
  SD Files Written:   1
  FD Bytes Written:   0 (0 B)
  SD Bytes Written:   251 (251 B)
  Rate:              0.0 KB/s
  Software Compression: None
  Comm Line Compression: 13.7% 1.2:1
  Snapshot/VSS:      no
  Encryption:         no
  Accurate:           yes
  Volume name(s):     TestVolume001
  Volume Session Id:  1
  Volume Session Time: 1755508583
  Last Volume Bytes:  102,704,274 (102.7 MB)
  Non-fatal FD errors: 0
  SD Errors:          0
  FD termination status: OK
  SD termination status: OK
  Termination:        Backup OK

18-Aug 11:17 user-dir JobId 1: Begin pruning Jobs older than 6 months .
18-Aug 11:17 user-dir JobId 1: No Jobs found to prune.
```

(continues on next page)

(continued from previous page)

```
18-Aug 11:17 user-dir JobId 1: Begin pruning Files.
18-Aug 11:17 user-dir JobId 1: No Files found to prune.
18-Aug 11:17 user-dir JobId 1: End auto prune.
```

Note: When an HPC Accelerator job is started, the initial HPC Accelerator job that starts is the `control` job which will backup a single empty record and then trigger additional worker jobs. It is anticipated that this control job backs up 1 file and 0 bytes.

7.2 Restore

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

Restore operations for the HPC Accelerator job sets are handled by Bacula in the same manner as any other any other restore. Users are not required to modify their standard restore workflow.

In the case of point-in-time restores using, for instance `bconsole`'s option 6, "Select backup for a client before a specified time", Bacula will automatically make sure that all worker backup jobs that had been run concurrently as part of a HPC Accelerator backup set are pulled into the restore directory tree. This guarantees that the restore tree from which files are selected is complete.

```
* restore fileset=big-fs-plugin select all done
Automatically selected Catalog: MyCatalog
Using Catalog "MyCatalog"
Automatically selected Client: host-fd
+-----+-----+-----+-----+-----+-----+-----+
| jobid | level | jobfiles | jobbytes | starttime           | volumename       |
+-----+-----+-----+-----+-----+-----+-----+
| 1     | F     | 19,249   | 205,589 | 2023-12-15 15:55:46 | TestVolume001   |
| 2     | I     | 645      | 7,423   | 2023-12-15 15:55:50 | TestVolume001   |
| 3     | I     | 203      | 2,272   | 2023-12-15 15:55:54 | TestVolume001   |
| 4     | I     | 33       | 302     | 2023-12-15 15:56:01 | TestVolume001   |
+-----+-----+-----+-----+-----+-----+-----+
You have selected the following JobIds: 1,2,3,4

Building directory tree for JobId(s) 1,2,3,4 ...
↳ ++++++
20,126 files inserted into the tree and marked for extraction.
Bootstrap records written to .../working/host-dir.restore.1.bsr

The Job will require the following (*=>InChanger):
Volume(s)           Storage(s)           SD Device(s)
=====
TestVolume001       File                 FileStorage
```

(continues on next page)

Volumes marked with "*" are in the Autochanger.

20,130 files selected to be restored.

Using Catalog "MyCatalog"

Run Restore job

JobName: RestoreFiles
Bootstrap: /opt/bacula//working/host-dir.restore.1.bsr
Where: /tmp/bacula-restores
Replace: Always
FileSet: Full Set
Backup Client: host-fd
Restore Client: host-fd
Storage: File
When: 2023-12-15 15:56:06
Catalog: MyCatalog
Priority: 10
Plugin Options: *None*

OK to run? (Yes/mod/no): yes

Job queued. JobId=5

You have messages.

* messages

15-dec 15:56 host-dir JobId 5: Start Restore Job RestoreFiles.2023-12-15_15.
→56.06_12

15-dec 15:56 host-dir JobId 5: Restoring files from JobId(s) 1,2,3,4

15-dec 15:56 host-dir JobId 5: Found a volume cycle in the bootstrap, fixing
→automatically the reading process

15-dec 15:56 host-dir JobId 5: Connected to Storage "File" at host:8103 with
→TLS

15-dec 15:56 host-dir JobId 5: Using Device "FileStorage" to read.

15-dec 15:56 host-dir JobId 5: Connected to Client "host-fd" at host:8102
→with TLS

15-dec 15:56 host-fd JobId 5: Connected to Storage at host:8103 with TLS

15-dec 15:56 host-sd JobId 5: Ready to read from volume "TestVolume001" on
→File device "FileStorage" (/opt/bacula/archive).

15-dec 15:56 host-sd JobId 5: Forward spacing Volume "TestVolume001" to
→addr=242

15-dec 15:56 host-sd JobId 5: Elapsed time=00:00:01, Transfer rate=3.336 M
→Bytes/second

15-dec 15:56 host-dir JobId 5: Connected to Storage "File" at host:8103 with
→TLS

15-dec 15:56 host-dir JobId 5: Using Device "FileStorage" to read.

15-dec 15:56 host-fd JobId 5: Connected to Storage at host:8103 with TLS

15-dec 15:56 host-sd JobId 5: Ready to read from volume "TestVolume001" on
→File device "FileStorage" (/opt/bacula/archive).

15-dec 15:56 host-sd JobId 5: Forward spacing Volume "TestVolume001" to
→addr=242

15-dec 15:56 host-sd JobId 5: End of Volume "TestVolume001" at addr=4213600
→on device "FileStorage" (/opt/bacula/archive).

15-dec 15:56 host-sd JobId 5: Elapsed time=00:00:01, Transfer rate=110.8 K

(continues on next page)

(continued from previous page)

```
↪Bytes/second
15-dec 15:56 host-dir JobId 5: Connected to Storage "File" at host:8103 with↪
↪TLS
15-dec 15:56 host-dir JobId 5: Using Device "FileStorage" to read.
15-dec 15:56 host-fd JobId 5: Connected to Storage at host:8103 with TLS
15-dec 15:56 host-sd JobId 5: Ready to read from volume "TestVolume001" on↪
↪File device "FileStorage" (/opt/bacula/archive).
15-dec 15:56 host-sd JobId 5: Forward spacing Volume "TestVolume001" to↪
↪addr=193778
15-dec 15:56 host-sd JobId 5: Elapsed time=00:00:01, Transfer rate=34.77 K↪
↪Bytes/second
15-dec 15:56 host-dir JobId 5: Connected to Storage "File" at host:8103 with↪
↪TLS
15-dec 15:56 host-dir JobId 5: Using Device "FileStorage" to read.
15-dec 15:56 host-fd JobId 5: Connected to Storage at host:8103 with TLS
15-dec 15:56 host-sd JobId 5: Ready to read from volume "TestVolume001" on↪
↪File device "FileStorage" (/opt/bacula/archive).
15-dec 15:56 host-sd JobId 5: Forward spacing Volume "TestVolume001" to↪
↪addr=2967749
15-dec 15:56 host-sd JobId 5: Elapsed time=00:00:01, Transfer rate=5.551 K↪
↪Bytes/second
15-dec 15:56 host-dir JobId 5: Bacula host-dir 18.0.0 (20Nov23):
Build OS:          x86_64-pc-linux-gnu archlinux
JobId:             5
Job:               RestoreFiles.2023-12-15_15.56.06_12
Restore Client:    "host-fd" 18.0.0 (20Nov23) x86_64-pc-linux-gnu,
↪archlinux,
Where:             /tmp/bacula-restores
Replace:           Always
Start time:        15-dec-2023 15:56:09
End time:          15-dec-2023 15:56:12
Elapsed time:      3 secs
Files Expected:    20,130
Files Restored:    20,130
Bytes Restored:    215,586 (215.5 KB)
Rate:              71.9 KB/s
FD Errors:         0
FD termination status: OK
SD termination status: OK
Termination:       Restore OK

[...]
```

8 HPC Accelerator: Best Practices

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

- It is advised to set the number of workers to be double the amount of disks designated for backup in order to prevent overloading the disk with read operations.
- It is advised to use `Accurate = yes` in the Job resource definition or to run a Job with `Accurate = yes` at a regular interval. Errors occurring with workers may lead to incomplete files backups. The Accurate option guarantees that all files are properly backed up at a given point in time.

9 HPC Accelerator: Limitations

Enterprise

Bacula Enterprise Only

This solution is **only** available for Bacula Enterprise. For upgrade inquiries, please reach out to sales@baculasystems.com.

- Virtual Full jobs need to have one-to-one mapping with devices to avoid creating loops on said devices. See `bacula-sd.conf.sample` generated during installation.
- The `estimate` command is not compatible with the HPC Accelerator.
- This plugin does not support incomplete jobs, as highlighted in Fileset example at `BigFsOperationsBackup`.
- In the restore command line of `bconsole`, the number of files reported as inserted into the tree may not match the actual number of entries.
- The `restart` command has limitations with plugins, as it initiates the Job from scratch rather than continuing it. Bacula determines whether a Job is restarted or continued, but using the `restart` command will result in a new Job.