



MSSQL VDI Plugin

Bacula Systems Documentation

Contents

- 1 Overview** **4**
- 2 Features Summary** **4**
- 3 Scope** **4**
- 4 Installation** **4**
- 5 Configuration** **5**
- 6 Plugin Options** **8**
- 7 Backup** **11**
- 8 Full Backup** **11**
- 9 Differential Backup** **11**
- 10 Transaction Log Backup** **11**
- 11 Point in Time Restore (PITR)** **11**
- 12 MSSQL Database Configuration** **11**
- 13 Restore** **12**
- 14 Restore Options** **12**
- 15 Point In Time Restore** **14**
 - 15.1 LSN Information 15
- 16 Restore Scenarios Overview** **16**
 - 16.1 Restore Using Advanced Relocation 16
 - 16.2 Restore With Same Name 17
 - 16.3 Restore Database With a New Name 17
 - 16.4 Restore to Local Disk 19
- 17 Restore the “master” Database** **22**
- 18 Database in *restoring* State** **22**
- 19 Always On Availability Groups** **22**
- 20 Backup** **22**
- 21 Restore** **22**
- 22 Limitations** **23**

Contents

- *Overview*
- *Features Summary*
- *Scope*
- *Installation*
- *Configuration*
- *Plugin Options*
- *Backup*
- *Full Backup*
- *Differential Backup*
- *Transaction Log Backup*
- *Point in Time Restore (PITR)*
- *MSSQL Database Configuration*
- *Restore*
- *Restore Options*
- *Point In Time Restore*
- *Restore Scenarios Overview*
- *Restore the “master” Database*
- *Database in restoring State*
- *Always On Availability Groups*
- *Backup*
- *Restore*
- *Limitations*

1 Overview

This user's guide presents how to use the MSSQL Plugin feature with **Bacula Enterprise**.

2 Features Summary

Bacula Systems provides a plugin for Microsoft SQL Server for **Bacula Enterprise** named `mssql-fd.dll`. The MSSQL Plugin provides the following main features:

- Full and Differential support
- Incremental (Log) level support
- Database level backup
- Ability to include/exclude databases from the backup job
- Support for “Copy Only” backups
- Restore MSSQL backup files to disk
- Send the backup stream directly to the Storage Daemon without requiring much local disk space.
- *Point in time recovery restore*

The MSSQL VDI plugin is compatible with Copy/Migration jobs. Please read the CopyMigrationJobsReplication for more information.

3 Scope

This document will present solutions for **Bacula Enterprise** 8.4 and later, which are not applicable to prior versions. The MSSQL Plugin has been tested and is supported from Windows Server 2003 R2 up to Windows Server 2019 and from Microsoft SQL Server 2005 up to 2019.

4 Installation

You must ensure that the `mssql-driver.dll` is in the Bacula program directory and the `mssql-fd.dll` plugin is in the plugins directory on the FD doing the backup (done by default with the installer), and that the `Plugin Directory` directive line is present and enabled in the FD's configuration file `bacula-fd.conf`. An example configuration file and status output of a client with the MSSQL plugin available is shown below.

```
*status client
Automatically selected Client: win2008-fd
Connecting to Client win2008-fd at win2008-64-r2:9102

win-fd Version: 8.4.3 (30 November 2015) VSS Linux Cross-compile Win64
Daemon started 11-Dec-15 05:13. Jobs: run=1 running=0.
Microsoft Windows Server 2008 R2 Standard Edition (build 7600), 64-bit
Heap: heap=0 smbytes=348,013 max_bytes=481,620 bufs=134 max_bufs=162
Sizes: boffset_t=8 size_t=8 debug=0 trace=1 mode=0,2010 bwlimit=0kB/s
Plugin: alldrives-fd.dll mssql-fd.dll
```

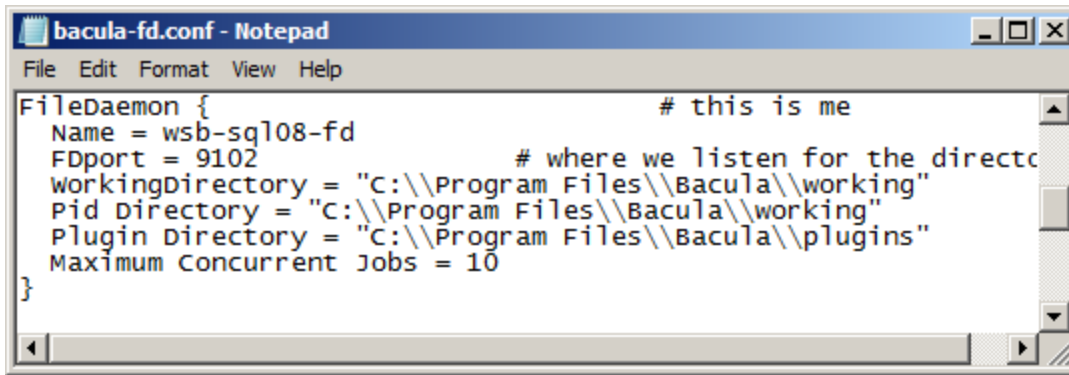


Fig. 1: File Daemon Configuration Excerpt with “Plugin Directory” Line

If the SQL Server database is running under an account that is not NT AUTHORIZED/SYSTEM, it will be mandatory to configure the SQL Server instance to allow the Bacula File Daemon service account to connect and perform backup operations. By default, the Bacula File Daemon service runs under the NT AUTHORIZED/SYSTEM account.

The permission sysadmin can be granted with the following SQL command:

```
ALTER SERVER ROLE [sysadmin] ADD MEMBER [NT AUTHORITY\\SYSTEM]
```

5 Configuration

Note: In case of a clustered MSSQL environment, or AlwaysOn Availability Groups, the Bacula Enterprise MSSQL VDI plugin will not automatically detect the configuration. Precautions need to be taken on both the MSSQL Server and Bacula Plugin side to make sure databases are adequately protected in such an environment.

Please contact Bacula Systems (support@baculasystems.com) for help on configuring the MSSQL VDI plugin if in doubt.

To activate the MSSQL plugin you have to put the following into the Include section of the File Set which will be used to back up the SQL Server data:

```
Plugin = "mssql"
```

This will back up all SQL server databases (tempdb is excluded by default).

The plugin directive must be specified exactly as shown above.

If everything is set up correctly as above then the backup will include the SQL server data. The SQL server data files backed up will appear in a **bconsole** or **bat** restore as follows:

```
/@mssql/MSSQLSERVER/master/data.bak
/@mssql/MSSQLSERVER/production/data.bak
...
etc
```

It is possible to select different instances or databases to be backed up with the following parameters:

- instance
- database

- include
- exclude

Full, Differential and Logs (Incremental) backups are supported.

Following the creation of a new database, you should run a Full backup of the SQL server data. A new database will not be backed up at a different level, and a Differential backup will automatically be upgraded to a Full backup on this specific new database.

The MSSQL Plugin does not use VSS snapshots to perform the backup so, unless some disk folder is present in the fileset, **Enable VSS** can be set to “no”.

A complete example of the Job setup for MSSQL Server data is shown below:

```
FileSet {
  Name = MSSQL
  Enable VSS = no      # VSS is not required
  Include {
    Options {
      Signature = SHA1
    }
    Plugin = "mssql: database=production"
  }
}

Job {
  Name = MSSQL08
  File Set = MSSQL
  Client = wsb-sql08-fd
  Job Defs = DefaultJob
  Level = Differential
}
```

```
FileSet {
  Name = MSSQL09
  Enable VSS = no      # VSS is not required
  Include {
    Options {
      Signature = SHA1
    }
    Plugin = "mssql: abort_on_error exclude=test1 exclude=test2 copyonly"
  }
}

Job {
  Name = MSSQL09
  File Set = MSSQL09
  Client = wsb-sql08-fd
  Job Defs = DefaultJob
  Level = Full
}
```

```
FileSet {
  Name = MSSQL10
```

(continues on next page)

```

Enable VSS = no      # VSS is not required
Include {
  Options {
    Signature = SHA1
  }
  Plugin = "mssql: include=test2 include=prod1 include=r7*"
}
}
Job {
  Name = MSSQL10
  File Set = MSSQL10
  Client = wsb-sql08-fd
  Job Defs = DefaultJob
  Level = Full
}

```

```

FileSet {
  Name = SQLExpress
  Enable VSS = no      # VSS is not required
  Include {
    Options {
      Signature = SHA1
    }
    Plugin = "mssql: instance=SQLExpress hostname=."
  }
}
Job {
  Name = SQLExpress
  File Set = SQLExpress
  Client = wsb-sql08-fd
  Job Defs = DefaultJob
  Level = Differential
}

```

Attention: New in Bacula Enterprise 12.8.0

MS SQL Server Multi Instances Backup Job}footnote{ Available in Bacula Enterprise 12.8.0}

```

FileSet {
  Name = SQLAllInstances
  Enable VSS = no      # VSS is not required
  Include {
    Options {
      Signature = SHA1
    }
    Plugin = "mssql: all_instances"
  }
}
Job {
  Name = SQLAllInstances

```

```

File Set = SQLAllInstances
Client = wsb-sql08-fd
Job Defs = DefaultJob
Level = Incremental
}

```

Attention: New in Bacula Enterprise 12.8.0

MS SQL Server 2 Instances Backup Job

```

FileSet {
  Name = SQL2Instances
  Enable VSS = no      # VSS is not required
  Include {
    Options {
      Signature = SHA1
    }
    Plugin = "mssql: instance=MSSQLSERVER instance=PRODUCTION"
  }
}

Job {
  Name = SQL2Instances
  File Set = SQL2Instances
  Client = wsb-sql08-fd
  Job Defs = DefaultJob
  Level = Incremental
}

```

6 Plugin Options

- **database=<glob>** Specifies the name of the databases to backup. This parameter is optional. By default, all databases (except tempdb) will be backed up.
- **include=<glob>** Specifies the names of databases to backup. It is possible to specify the include parameter multiple times on the plugin command line.
If a database that was explicitly specified is not found, a warning or error message will be printed to the Job report. The **abort_on_error** option is adhered to.
- **exclude=<glob>** Specifies the names of databases to exclude from the backup. It is possible to specify the exclude parameter multiple times on the plugin command line.
- **user=<str>** The username used to connect to the MSSQL instance. If the user is part of a domain, please also specify the domain parameter below. If no domain is specified the user parameter will refer to a local account. This parameter is optional, and if not set, the bacula-fd service account will be used to connect to the MSSQL instance.
- **domain=<str>** The domain of the user used to connect to the MSSQL instance. This parameter is optional. If not set and no user is set, the bacula-fd service account will be used to connect to the MSSQL instance. If not set and a user is set, the user is assumed to be a local account.

- **password=<str>** The password used to connect to the MSSQL instance. This parameter is optional, and if set, the password might be printed in various places such as `status client` output, job log or debug messages.
- **passfile=<file>** The file where the password can be found. This parameter is optional. The plugin will use the first line (limited to 127 characters) as the connection password.
- **authtype=[windows | server]** The authentication type used to connect to the MSSQL instance. This parameter is optional. By default the Windows authentication type is used.
- **instance=<str>** The instance name used to connect to the MSSQL instance. This parameter is optional. By default, the instance name is “MSSQLSERVER”.

Attention: New in Bacula Enterprise 12.8.0

Multiple instance parameters are allowed.

Attention: New in Bacula Enterprise 12.8.0

all_instances This parameter is optional. If set, the MSSQL plugin will list and backup all instances defined on the SQL Server. The same authentication settings should be available on each instance.

- **hostname=<str>** This parameter is optional. If set, the `hostname` string will be used in the `Server` connection string of the ODBC driver. If not set, the plugin will configure automatically the ODBC driver to use `(localdb)` or `(local)`. This advanced option might be used in conjunction with the `instance` option, often, it is necessary to specify “.” as `hostname` parameter. Ex: `instance=MYINSTANCE hostname=.`
- **abort_on_error** By default, if the plugin is not able to reach the MSSQL instance, or to find an explicitly named database, an error will be generated and the job will continue. Some users might prefer to abort the Job, which will happen if this option is set.
- **copyonly=[incremental|all]** A copy-only backup is a MSSQL backup that is independent of the sequence of conventional MSSQL backups. The next Differential or Transaction Log (Incremental) backup will not use it. If the `incremental` option is set, only Transaction Log (Incremental) backups will use the option. The last Incremental job will contain all the information necessary to perform a point in time recovery.
- **fullbackup** Force the level of the MSSQL backup. If specified in the plugin command line, a job can run with the incremental level for standard files, and the databases backed up with the plugin will always be backed up like with a Full backup. (Available since 8.4.11)
- **skipreadonly** Skip read only databases when the backup level is incremental (BACKUP LOG) and the last modification/creation time for tables and views is prior to the last backup. By default, read only databases are upgraded to Full. (Available since 8.6.20)
- **dblayout** Store each database layout as a RestoreObject in the Bacula Catalog with a Full backup. The layout can be displayed in `bconsole` with `list restoreobjects` command. (Available since 8.6.20)
- **lock_timeout=<int>** Use the `SET LOCK_TIMEOUT` before issuing queries. (Available since 8.6.20)
- **buffercount=<int>** Specifies the total number of I/O buffers to be used for the backup operation. You can specify any positive integer; however, large numbers of buffers might cause “out of memory” errors because of inadequate virtual address space in the `Sqlservr.exe` process. This parameter is optional and the default value is 10.
- **blocksize=<bytes>** Specifies the physical block size, in bytes. The supported sizes are 512, 1024, 2048, 4096, 8192, 16384, 32768, and 65536 (64 KB) bytes. This parameter is optional and the default is 65536.

- **maxtransfersize=<bytes>** Specifies the largest unit of transfer in bytes to be used between the MSSQL Server and the backup media. The possible values are multiples of 65536 bytes (64 KB) ranging up to 4194304 bytes (4 MB). This parameter is optional and the default is 65536.
- **connection_string=<str>** Specifies the ODBC connection string. This parameter is optional.
- **checksum=<No/Yes>**

Available since Bacula Enterprise 8.11.6

Use the WITH CHECKSUM option in the backup query. This parameter is optional and the default is no.

- **driver=<str>**

Available since Bacula Enterprise 10.2.3

The driver name used in the ODBC connection string. Default is SQL Server. If the `connection_string` option is set, it will overwrite this option.

- **target_backup_recovery_models=<str>**

Available since Bacula Enterprise 10.2.4

Only the databases with the specified recovery model are backed up. The other databases are skipped with warning. The parameter should be one of the following:

- SIMPLE
 - FULL
 - BULK_LOGGED
-

- **simple_recovery_models_incremental_action=<str>**

Available since Bacula Enterprise 10.2.4

Specifies the behavior when incremental backup is requested over a simple recovery model database from the following parameters:

- `make_full` : The incremental backup is automatically upgraded to full.
 - `ignore_with_error` : The backup is skipped with an error (the error translated to a warning at the job level).
 - `ignore` : The backup is skipped silently.
-

7 Backup

8 Full Backup

The Full backup saves the database files and MSSQL to provide complete protection against media failure. If one or more data files are damaged, media recovery can restore all committed transactions. In-process transactions are rolled back. The master and the mddb databases are always backed up in this mode.

9 Differential Backup

A differential backup is based on the most recent, previous full database backup. A differential backup captures only the data that has changed since that full backup. When using the Differential backup feature, the backup chain is very critical. If for some reason, the Full backup used as referenced by MSSQL is not available, the Differential data will not be usable. The plugin uses different techniques to avoid this problem, so if a problem is detected, the Differential database backup might be automatically upgraded to a Full backup.

10 Transaction Log Backup

The “Transaction Log Backup” MSSQL feature is implemented as the “Incremental” level with Bacula. The database must be configured with the full recovery model or bulk-logged recovery model. If the database uses the simple recovery model, the MSSQL file will be truncated after each checkpoint. The full restore will be possible, but not the restore to a point in time. For more information, see <https://msdn.microsoft.com/en-us/library/ms189275.aspx>.

11 Point in Time Restore (PITR)

Point In Time Restore (PITR) requires the database to be configured with the full recovery model. If the database uses the simple recovery model, the MSSQL file will be truncated after each checkpoint. For more information, see <https://msdn.microsoft.com/en-us/library/ms189275.aspx>.

12 MSSQL Database Configuration

The master database must be backed up. If master is damaged in some way, for example because of media failure, an instance of MSSQL may not be able to start. In this event, it is necessary to rebuild master, and then restore the database from a backup. Only full database backups of master can be created. See <https://technet.microsoft.com/en-us/library/aa213839%28v=sql.80%29.aspx> for more information.

13 Restore

You can use all the regular ways to start a restore. However, you must make sure that if restoring differential data, the previous full backup is also restored. This happens automatically if you start the restore, in `bconsole`, using the restore options 5 or 12. In the file tree generated, you should mark either complete databases or databases instances.

14 Restore Options

It is possible to restore the data with different scenarios using the following options:

- The common restore “Where” parameter (**where=<path>**)
- The common restore “Replace” parameter (**replace=<never|always>**)
- The Plugin restore option accessible in the “Plugin Options” menu at the restore prompt (Item 13).

```
Run Restore job
JobName:      RestoreFiles
Bootstrap:    /tmp/regress/working/127.0.0.1-dir.restore.9.bsr
Where:        c:/tmp
Replace:      Never
FileSet:      Full Set
Backup Client: win2008-fd
Restore Client: win2008-fd
Storage:      File
When:         2016-02-22 12:02:56
Catalog:     MyCatalog
Priority:      10
Plugin Options: *None*
OK to run? (yes/mod/no): mod          <-----
Parameters to modify:
  1: Level
  2: Storage
  3: Job
  4: FileSet
  5: Restore Client
  6: When
  7: Priority
  8: Bootstrap
  9: Where
 10: File Relocation
 11: Replace
 12: JobId
 13: Plugin Options
Select parameter to modify (1-13): 13          <-----
Automatically selected : mssql: database=db29187
Plugin Restore Options
instance:      *None*
database:     *None*
username:     *None*
password:     *None*
domain:       *None*
hostname:     *None*
```

(continues on next page)

```

authtype:          *None*
recovery:          *None*          (yes)
stop_before_mark:  *None*
stop_at_mark:     *None*
stop_at:          *None*
restricted_user:   *None*          (no)
verify:           *None*          (no)
connection_string: *None*
checksum:         *None*          (no)
driver:           *None*          (SQL Server)

Use above plugin configuration? (yes/mod/no): mod <-----
You have the following choices:
You have the following choices:
  1: instance (Instance used to restore)
  2: database (New database name)
  3: username (Username used for restore)
  4: password (Password used for restore)
  5: domain (Domain name of user (default to local))
  6: hostname (Server ODBC parameter (default to (local/localdb)))
  7: authtype (Authentication type (server or windows))
  8: recovery (Start Recovery)
  9: stop_before_mark (Stop the recovery before a mark (STOPBEFOREMARK). {lsn:lsn_
↵number | mark_name})
 10: stop_at_mark (Stop the recovery at a mark (STOPATMARK). {lsn:lsn_number | mark_
↵name})
 11: stop_at (Stop at (STOPAT). {datetime})
 12: restricted_user (Restrict access to the restored database)
 13: verify (Verify backup integrity)
 14: connection_string (ODBC Connection string)
 15: checksum (Restore using the WITH CHECKSUM option)
 16: driver (ODBC Driver Name)

Select parameter to modify (1-15):

```

The Plugin restore options are:

- **instance=<str>** The instance name used to connect to the MSSQL instance. This parameter is optional, and if not set, the restore will use the value set during at the backup time. By default, the instance name is “MSSQLSERVER”.
- **database=<name>** Specifies the name of the databases to restore. This parameter is optional. By default, the plugin will use the where option to determine the name of the new database. If both where and database are set to a valid database name, database will be used. A valid database name can contain the following characters: A-Za-z0-9#_
- **username=<str>** The username used to connect to the MSSQL instance. This parameter is optional, and if not set, the restore will use the value set during at the backup time.
- **password=<str>** The password used to connect to the MSSQL instance. This parameter is optional and if not set, the restore will use the value set during at the backup time.
- **domain=<filestr>** The domain used to connect to the MSSQL instance. This parameter is optional and if not set, the restore will use the value set during at the backup time.
- **hostname=<str>** The MSSQL server host name.

- **authtype**=[windows | server] The authentication type. Windows is default.
- **recovery**=<Yes/no> Specifies if the database will use the RECOVERY or the NORECOVERY option during the restore. By default, the restored database will be recovered.
- **stop_before_mark**=<markname> Use the WITH STOPBEFOREMARK = '<point>' clause to specify that the log record that is immediately before the mark is the recovery point. The point can be a LSN number or a mark_name.
- **stop_at_mark**=<markname> Use the WITH STOPATMARK = '<point>' clause to specify that the marked transaction is the recovery point. STOPATMARK rolls forward to the mark and includes the marked transaction in the roll forward. The point can be a LSN number or a mark_name.
- **stop_at**=<datetime> Use the WITH STOPAT = '<datetime>' clause to specify that the date time is the recovery point.
- **restricted_user**=<No/yes> Use the WITH RESTRICT_USER clause to restrict access to the restored database. The default is no.
- **connection_string**=<str> Specifies the ODBC connection string. This parameter is optional.
- **checksum**=<No/yes>

Available since Bacula Enterprise 8.11.6

Use the WITH CHECKSUM clause to the restored database. The default is no.

- **driver**=<str>

Available since Bacula Enterprise 10.2.3

The driver name used in the ODBC connection string. Default is SQL Server. If the connection_string option is set, it will overwrite this option.

On BWeb Management Suite, the Plugin Options are available in the restore tab.

15 Point In Time Restore

This topic is relevant only for SQL Server databases that use the full or bulk-logged recovery models. Under the bulk-logged recovery model, if a log backup contains bulk-logged changes, point-in-time recovery is not possible to a point within that backup. The database must be recovered to the end of the MSSQL backup.

More information can be found on <https://msdn.microsoft.com/en-us/library/ms179451.aspx>

It is possible to do Point In Time Restore of a MSSQL database directly from the MSSQL Plugin. It is also possible to restore files locally and do the operation from the Microsoft SQL Server Mangement Console to have more options.

15.1 LSN Information

LSNs are used internally during a RESTORE sequence to track the point in time to which data has been restored. When a backup is restored, the data is restored to the LSN corresponding to the point in time at which the backup was taken. More information can be found on <https://msdn.microsoft.com/en-us/library/ms190925.aspx>.

The LSN of a log record at which a given backup and restore event occurred is viewable using one or more of the following:

- Bacula Backup job output
- Log file names
- msdb.backupset table
- msdb.backupfile table

During a backup job with MSSQL Plugin, the following information about LSN numbers will be displayed in the Job output:

```
win-fd JobId 3: LSN for "db29187": First: 42000146037, Last: 44000172001
```

The **First LSN** number corresponds to the last LSN of the last transaction logs backup. It can be the very first Full backup, or the last transactional backup (Incremental). The **Last LSN** number corresponds to the last transaction recorded in the log.

With a MSSQL backup (Incremental), the file name associated with this database in the Incremental job will be:

```
/@mssql/MSSQLSERVER/db29187/log-42000162001.trn
```

The number in the name, here **42000162001** corresponds to the last LSN of the previous job (Full or Incremental).




Fig. 2: First, Last and Filename LSNs

In the example shown above, if the administrator needs to restore the database at the state that corresponds to LSN 14, it can be done with the following actions:

- Use restore menu option 5
- Browse the database directory “/@mssql/db29187”
- Select last Full backup file “data.bak” (LSN: 10)
- Select incremental backup “log-10.trn”
- Specify the stop_at_mark option to “LSN:14”
- Run the restore job

or if the last full backup is not available but the previous full backup is.

- Use restore menu option 3, select the relevant jobids
- Browse the database directory “/@mssql/db29187”
- Select Full backup file “data.bak” (LSN: 2)
- Select incremental backups “log-2.trn”, “log-3.trn”, “log-10.trn”
- Specify the stop_at_mark option to “LSN:14”
- Run the restore job

16 Restore Scenarios Overview

Table 1: Restore Scenarios

Description	where	regexwhere	database	Example
Restore files to disk	Path			where=c:/tmp
Restore original database to MSSQL				where=/
Restore with a new name to MSSQL	Name			where=newdb
Restore with a new name to MSSQL			Name	database=newdb
Restore with a new name and file relocation to MSSQL	Path		Name	where=c:/tmp database=newdb
Restore with a new name and individual file relocation to MSSQL		Regex	Name	regexwhere=!CLUSTER! MSSQLSERVER! database=newdb

16.1 Restore Using Advanced Relocation

In some cases, a database can use different disks to store the data.

```
c:\data\db1.MDF
e:\logs\db1_log.LDF
```

Available since Bacula Enterprise 8.6.17

Using the Bacula `RegexWhere` function, it is possible to manipulate each file and generate new filenames. The destination directories must exist prior to the restore.

In the following example, the database will be renamed and the path will be adapted to the new server.

```
db1          -> db1-old          !db1!db1-old!i
c:\data\db1.MDF -> f:\data\db1-old.MDF    !c:!f:!i
e:\logs\db1_log.LDF -> g:\data\db1-old_log.MDF !e:!g:!i
```

The resulting `RegexWhere` will be something like:

```
*restore regexwhere="!db1!db1-old!i,!c:!f:!i,!e:!g:!i"
...
database:          db1-old
```


Note that the **database** name must be specified in the Plugin Options menu, else, files will be stored on disk and the SQL restore commands will have to be executed manually.

To convert a path, the windows path separator must be escaped correctly in the console:

```
*restore regexwhere="!c:\\\\data\\\\db1!c:\\data2\\db1!i"
or
*restore regexwhere="!c:.data.db1!c:/data2/db1!i"
or
*restore regexwhere="!c:\\\\data\\\\db1!c:/data2/db1!i"
```

16.2 Restore With Same Name

To restore a database with the same name, the `where` parameter should be empty or `/` and the `replace=` flag should be set to `always` or the original database should be dropped first.

```
* restore where=/ replace=always
...
Using Catalog "MyCatalog"
Run Restore job
JobName:      RestoreFiles
Bootstrap:    /opt/bacula/working/127.0.0.1-dir.restore.1.bsr
Where:        /
Replace:      Always
FileSet:      Full Set
Backup Client: win2008-fd
Restore Client: win2008-fd
Storage:      File
When:         2015-12-14 09:53:36
Catalog:      MyCatalog
Priority:      10
Plugin Options: *None*
OK to run? (yes/mod/no):
```

16.3 Restore Database With a New Name

To restore a database with a new name, it might be required to relocate database files on disk. It depends if the original database is still present.

If the original database is no longer available, the `where` parameter or the “Plugin Options” database can contain the new database name, and the plugin will automatically handle the database creation with the new name.

If the original database is still required, the `where` parameter is used to relocate files on disk, and the new database name should be set with the “Plugin Options” menu with the database option. The `layout.dat` must be selected in the restore tree.

```
* restore where=c:/tmp replace=always
...
Using Catalog "MyCatalog"
Run Restore job
JobName:      RestoreFiles
Bootstrap:    /opt/bacula/working/127.0.0.1-dir.restore.1.bsr
Where:        c:/tmp
```

(continues on next page)

```

Replace:          Always
FileSet:          Full Set
Backup Client:    win2008-fd
Restore Client:   win2008-fd
Storage:          File
When:             2015-12-14 09:53:36
Catalog:          MyCatalog
Priority:          10
Plugin Options:   *None*
OK to run? (yes/mod/no): mod <-----
Parameters to modify:
  1: Level
  2: Storage
  3: Job
  4: FileSet
  5: Restore Client
  6: When
  7: Priority
  8: Bootstrap
  9: Where
 10: File Relocation
 11: Replace
 12: JobId
 13: Plugin Options
Select parameter to modify (1-13): 13 <-----
Automatically selected : mssql: database=db29187
Plugin Restore Options
instance:          *None*
database:          *None*
username:          *None*
password:          *None*
domain:            *None*
recovery:          *None*          (yes)
stop_before_mark: *None*
stop_at_mark:     *None*
stop_at:           *None*
Use above plugin configuration? (yes/mod/no): mod <-----
You have the following choices:
  1: instance (Instance used to restore)
  2: database (New database name)
  3: username (Username used for restore)
  4: password (Password used for restore)
  5: domain (Domain name of user (default to local))
  6: recovery (Start Recovery)
  7: stop_before_mark (Stop the recovery before a mark (STOPBEFOREMARK).
  8: stop_at_mark (Stop the recovery at a mark (STOPATMARK).
  9: stop_at (Stop at (STOPAT). {datetime})
Select parameter to modify (1-9): 2 <-----
Please enter a value for database: newdb <-----
Use above plugin configuration? (yes/mod/no): yes <-----

Using Catalog "MyCatalog"

```

```

Run Restore job
JobName:      RestoreFiles
Bootstrap:    /opt/bacula/working/127.0.0.1-dir.restore.1.bsr
Where:        c:/tmp
Replace:      Always
FileSet:      Full Set
Backup Client: win2008-fd
Restore Client: win2008-fd
Storage:      File
When:         2015-12-14 09:53:36
Catalog:     MyCatalog
Priority:      10
Plugin Options: User Specified
OK to run? (yes/mod/no): yes          <-----

```

16.4 Restore to Local Disk

When specifying `where=c:/path/`, files will be restored to the local filesystem and the MSSQL administrator can use a TSQL or the Microsoft SQL Server Management Console to restore the database. SQL commands needed to restore the database are printed in the Job output as showed in the next example.

```
* restore where=c:/tmp
```

First you select one or more JobIds that contain files to be restored. You will be presented several methods of specifying the JobIds. Then you will be allowed to select which files from those JobIds are to be restored.

To select the JobIds, you have the following choices:

- 1: List last 20 Jobs run
- 2: List Jobs where a given File is saved
- 3: Enter list of comma separated JobIds to select
- 4: Enter SQL list command
- 5: Select the most recent backup for a client
- 6: Select backup for a client before a specified time
- 7: Enter a list of files to restore
- 8: Enter a list of files to restore before a specified time
- 9: Find the JobIds of the most recent backup for a client
- 10: Find the JobIds for a backup for a client before a specified time
- 11: Enter a list of directories to restore for found JobIds
- 12: Select full restore to a specified Job date
- 13: Cancel

Select item: (1-13): 5

Automatically selected Client: win2008-fd

jobid	level	jobfiles	jobbytes	starttime	volumename
1	F	3	65,771	2015-12-14 09:52:31	TestVolume001
2	I	2	65,771	2015-12-14 09:52:42	TestVolume001
3	I	2	65,771	2015-12-14 09:52:52	TestVolume001

(continues on next page)

```
You have selected the following JobIds: 1,2,3

Building directory tree for JobId(s) 1,2,3 ...
6 files inserted into the tree.

You are now entering file selection mode where you add (mark) and
remove (unmark) files to be restored. No files are initially added, unless
you used the "all" keyword on the command line.
Enter "done" to leave this mode.

cwd is: /
$ cd @mssql
cwd is: /@mssql/
$ cd MSSQLSERVER
cwd is: /@mssql/MSSQLSERVER/
$ m db1684
6 files marked.
$ done
Bootstrap records written to /opt/bacula/working/127.0.0.1-dir.restore.1.bsr

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

Volumes marked with "*" are in the Autochanger.

2 files selected to be restored.

Using Catalog "MyCatalog"
Run Restore job
JobName:      RestoreFiles
Bootstrap:    /opt/bacula/working/127.0.0.1-dir.restore.1.bsr
Where:        /tmp
Replace:      Always
FileSet:      Full Set
Backup Client: win2008-fd
Restore Client: win2008-fd
Storage:      File
When:         2015-12-14 09:53:36
Catalog:      MyCatalog
Priority:      10
Plugin Options: *None*
OK to run? (yes/mod/no): yes
Job queued. JobId=6
wait
You have messages.
* messages

$ done
```

```

17:18 dir JobId 6: Start Restore Job RestoreFiles.2015-12-14_17.18.18_14
17:18 dir JobId 6: Using Device "FileStorage" to read.
17:18 sd JobId 6: Ready to read from volume "TestVolume001" on file device "FileStorage"
↳ (/tmp/regress/tmp).
17:18 sd JobId 6: Forward spacing Volume "TestVolume001" to file:block 0:224.
17:18 fd JobId 6: RESTORE DATABASE [db1684] FROM DISK='c:/tmp/@mssql/MSSQLSERVER/db1684/
↳ data.bak'
                WITH BLOCKSIZE=65536, BUFFERCOUNT=10, MAXTRANSFERSIZE=65536,
↳ NORECOVERY , REPLACE
17:18 fd JobId 6: RESTORE DATABASE [db1684]
                FROM DISK='c:/tmp/@mssql/MSSQLSERVER/db1684/log-34000000014400001.
↳ bak'
                WITH BLOCKSIZE=65536, BUFFERCOUNT=10, MAXTRANSFERSIZE=65536,
↳ NORECOVERY
17:18 fd JobId 6: RESTORE DATABASE [db1684]
                FROM DISK='c:/tmp/@mssql/MSSQLSERVER/db1684/log-34000000018400001.
↳ bak'
                WITH BLOCKSIZE=65536, BUFFERCOUNT=10, MAXTRANSFERSIZE=65536,
↳ NORECOVERY
17:18 fd JobId 6: RESTORE DATABASE [db1684]
                FROM DISK='c:/tmp/@mssql/MSSQLSERVER/db1684/log-34000000029100001.
↳ bak'
                WITH BLOCKSIZE=65536, BUFFERCOUNT=10, MAXTRANSFERSIZE=65536,
↳ NORECOVERY
17:18 sd JobId 6: End of Volume at file 0 on device "FileStorage" (/tmp/regress/tmp),
↳ Volume "TestVolume001"
17:18 fd JobId 6: RESTORE DATABASE [db1684]
                FROM DISK='c:/tmp/@mssql/MSSQLSERVER/db1684/log-36000000017200001.
↳ bak'
                WITH BLOCKSIZE=65536, BUFFERCOUNT=10, MAXTRANSFERSIZE=65536,
↳ NORECOVERY
17:18 sd JobId 6: Elapsed time=00:00:01, Transfer rate=9.372 M Bytes/second
17:18 fd JobId 6: RESTORE DATABASE [db1684]
17:18 dir JobId 6: Bacula dir 8.4.8 (22Feb16):
  Build OS:          x86_64-unknown-linux-gnu archlinux
  JobId:             6
  Job:               RestoreFiles.2015-12-11_17.18.18_14
  Restore Client:   win2008-fd
  Start time:       14-Dec-2015 17:18:20
  End time:         14-Dec-2015 17:18:22
  Files Expected:   6
  Files Restored:   6
  Bytes Restored:   9,371,785
  Rate:             4685.9 KB/s
  FD Errors:        0
  FD termination status: OK
  SD termination status: OK
  Termination:     Restore OK

```

17 Restore the “master” Database

Instructions on how to restore the “master” database are detailed in this article: <https://technet.microsoft.com/en-us/library/aa213839%28v=sql.80%29.aspx>

18 Database in *restoring* State

At the end of a restore, if the plugin option `recovery` was set to `no`, the restored database will be in the “restoring” state. To end the restore process, the recovery process must be run. It can be done with the following SQL command:

```
RESTORE [yourdatabase] WITH RECOVERY;
```

19 Always On Availability Groups

Availability groups can be created in MSSQL to provide high data availability over WSFC cluster nodes, as detailed in this article: <https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/overview-of-always-on-availability-groups-sql-server?view=sql-server-2017>

It can be interesting depending on the system scale to perform backups on secondary replicas to free up time and resource costs from the primary replica. The Bacula MSSQL Plugin is compatible with this type of architecture, with some restrictions.

20 Backup

When availability group(s) are defined, automated backup preference should also be specified to indicate the type of backup strategy (prefer secondary, primary only, etc.). The MSSQL plugin enforces this preference, only allowing backup of the preferred replica.

Note that depending on replica, not all backup type are allowed. Secondary will support copy-only full database backups and log backups while differential is not supported. This needs to be taken in account when creating Bacula’s backup Jobs. Refer to this article for more details: <https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/active-secondaries-backup-on-secondary-replicas-always-on-availability-groups?view=sql-server-2017>

21 Restore

Database restores can not be performed when they are part of one or many availability groups. This is detected by the plugin and the database restore will be skipped when this is the case. To workaround this, you need to remove the database from the availability group(s), then restore the database and eventually put it back in its original availability group(s). More details can be found here: <https://social.technet.microsoft.com/wiki/contents/articles/28148-restoring-a-backup-database-to-an-availability-group-sql-server.aspx>

22 Limitations

The current version of the plugin has the following limitations:

- The plugin can have only one Plugin command line in the FileSet. To backup multiple databases in the same Job, it is possible to use the **include** parameter multiple times on the plugin command string.
- To restore a database with a new name and a new location, the `layout.dat` file must be selected in the restore process. The relocation function will not work correctly if data files were added after the last Full backup job.
- It is not possible to restore multiple databases with a new name in a single restore job. The restore should be performed in different restore jobs.
- Database snapshots are automatically excluded from the backup.
- 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.

These limitations will be addressed in a future version.