



EDB Postgres™ Backup and Recovery User Guide

Release 2.5.2

EDB Postgres™ Backup and Recovery User Guide

Jan 13, 2020

1	Introduction	1
1.1	What's New	2
1.2	Conventions Used in this Guide	3
1.2.1	Restrictions on pg_basebackup	3
2	Overview	4
2.1	Block-Level Incremental Backup	7
2.1.1	Incremental Backup Limitations and Requirements	7
2.1.2	Concept Overview	8
2.1.3	WAL Scanning – Preparation for an Incremental Backup	9
2.1.4	Performing an Incremental Backup	10
2.1.5	Restoring an Incremental Backup	11
2.2	Creating a Backup Chain	13
3	Using BART	14
3.1	BART Management Overview	14
3.1.1	Performing a Restore Operation	16
3.1.2	Point-In-Time Recovery Operation	18
3.2	Managing Backups Using a Retention Policy	20
3.2.1	Overview - Managing Backups Using a Retention Policy	20
3.2.2	Marking the Backup Status	21
3.2.3	Setting the Retention Policy	22
3.2.4	Managing the Backups Based on the Retention Policy	23
3.2.5	Managing Incremental Backups	25
3.3	Basic BART Subcommand Usage	27
3.3.1	CHECK-CONFIG	29
3.3.2	INIT	30
3.3.3	BACKUP	32
3.3.4	SHOW-SERVERS	36
3.3.5	SHOW-BACKUPS	37
3.3.6	VERIFY-CHKSUM	38
3.3.7	MANAGE	39
3.3.8	RESTORE	42
3.3.9	DELETE	44
3.4	Running the BART WAL Scanner	45
4	Using Tablespaces	47

5 Conclusion

49

Index

50

The EDB Backup and Recovery Tool (BART) is an administrative utility that provides simplified backup and recovery management for multiple local or remote EDB Postgres™ Advanced Server and PostgreSQL® database servers.

BART features:

- Support for complete, hot, physical backups of multiple Advanced Servers and PostgreSQL database servers
- Support for two types of backups – full base backups and block-level incremental backups
- Backup and recovery management of database servers on local or remote hosts
- A single, centralized catalog for backup data
- Retention policy support for defining and managing how long backups should be kept
- The capability to store the backup data in compressed format
- Verified backup data with checksums
- Backup information displayed in an easy-to-read format
- A simplified point-in-time recovery process

This guide provides the following information about using BART:

- an *overview* of the BART components and concepts.
- information about the *backup and recovery management process* with BART.
- information about *using tablespaces*.

1.1 What's New

In BART 2.5.2, the following configuration parameters have been added in the `bart.conf` file:

- The `bart_socket_directory` parameter is added in the global section of the BART configuration file. Using this parameter, you can specify the socket directory path where all BART sockets will be stored. The default directory is `/tmp`.
- The `archive_path` parameter is added in the server section of the BART configuration file. Using this parameter, you can specify the path where archived WAL files will be stored. The default location of the archived WAL files is the BART backup catalog (`<backup_path>/<server_name>/archived_wals`).

For more details about these parameters, please see the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

1.2 Conventions Used in this Guide

The following is a list of conventions used throughout this document.

- Much of the information in this document applies interchangeably to the PostgreSQL and EDB Postgres Advanced Server database systems. The term *Advanced Server* is used to refer to EDB Postgres Advanced Server. The term *Postgres* is used to generically refer to both PostgreSQL and Advanced Server. When a distinction needs to be made between these two database systems, the specific names, PostgreSQL or Advanced Server are used.
- The installation directory of the PostgreSQL or Advanced Server products is referred to as `POSTGRES_INSTALL_HOME`:
 - For PostgreSQL Linux installations, this defaults to `/opt/PostgreSQL/x.x` for version 10 and earlier. For later versions, the installation directory is `/var/lib/pgsql/x`.
 - For Advanced Server Linux installations performed using the interactive installer for version 10 and earlier, this defaults to `/opt/PostgresPlus/x.xAS` or `/opt/edb/asx.x`. For Advanced Server Linux installations performed with an RPM package, this defaults to `/usr/ppas-x.x` or `/usr/edb/asx.x`. For Advanced Server Linux installations performed with an RPM package for version 11 or later, this defaults to `/usr/edb/asxx`.

1.2.1 Restrictions on `pg_basebackup`

BART takes full backups using the `pg_basebackup` utility program under the following conditions:

- The backup is taken on a standby server.
- The `--with-pg_basebackup` option is specified with the `BACKUP` subcommand (see *Backup*).
- The number of thread count in effect is 1, and the `with-pg_basebackup` option is not specified with the `BACKUP` subcommand.
- Database servers can only be backed up by using `pg_basebackup` utility program of the same or later version than the database server version. For example, `pg_basebackup` version 9.5 can back up database server version 9.5, but it cannot back up database server version 9.6.

In the global section of the BART configuration file, the `pg_basebackup_path` parameter specifies the complete directory path to the `pg_basebackup` program.

For information about the `pg_basebackup_path` parameter and setting the `thread_count`, see the configuration section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

For information about `pg_basebackup`, see [PostgreSQL Core Documentation](#).

BART provides a simplified interface for the continuous archiving and point-in-time recovery method provided with Postgres database servers. This consists of the following processes:

- Capturing a complete image of a database cluster as a full base backup or referred to simply as a *full backup*.
- Capturing a modified image of a database cluster called a *block-level incremental backup*, which is similar to a full backup, but contains the modified blocks of the relation files that have been changed since a previous backup instead of all, full relation files.
- Archiving the Write-Ahead Log segments (WAL files), which continuously record changes to be made to the database files.
- Performing *Point-In-Time Recovery* (PITR) to a specified transaction ID or timestamp with respect to a timeline using a full backup along with successive, *block-level incremental backups* that reside in the same backup chain, and the WAL files.

Detailed information regarding WAL files and point-in-time recovery is documented in the [PostgreSQL Core Documentation](#).

Block-level incremental backups are referred as *incremental backups*. The general term *backup* refers to both full backups and incremental backups. When a distinction must be made between the two, the complete term *full backup* or *incremental backup* is used.

When taking a full backup of a standby server, BART uses the PostgreSQL `pg_basebackup` utility program. However, it must be noted that for standby servers, you can only take a full backup, but cannot take incremental and parallel backups. For information about standby servers, see the [PostgreSQL Documentation](#).

BART simplifies the management process by use of a centralized backup catalog, a single configuration file, and a command line interface controlling the necessary operations. Reasonable defaults are automatically used for various backup and restore options. BART also performs the necessary recovery file configuration required for point-in-time recovery using its command line interface.

BART also provides the following features to enhance backup management:

- Automation of the WAL archiving command configuration.
- Usage of retention policies to evaluate, categorize, and delete obsolete backups.
- Compression of WAL files to conserve disk space.
- Customizable naming of backups to ease their usage.
- Easy access to comprehensive information about each backup.

The key components of a BART installation are:

- **BART Host.** The host system on which BART is installed. BART operations are invoked from this host system. The database server backups and archived WAL files are stored on this host as well.
- **BART User Account.** Linux operating system user account you choose to run BART. The BART user account owns the BART backup catalog directory.
- **BART Configuration File.** File in editable text format containing the configuration information used by BART.
- **BART Backup Catalog.** File system directory structure containing all of the backups for the database servers managed by BART. It is also the default `archive_path` to store archived WAL files.
- **BART Backupinfo File.** File in text format containing information for a BART backup. A `backupinfo` file resides in each backup subdirectory within the BART backup catalog.
- **BART Command Line Utility Program.** Single, executable file named `bart`, which is used to manage all BART operations.
- **BART WAL Scanner Program.** Single, executable file named `bart-scanner`, which is used to scan WAL files to locate and record the modified blocks for incremental backups.

Other concepts and terms referred to in this document include the following:

- **Postgres Database Cluster.** Also commonly called the *data directory*, this is the file system directory where all of the data files related to a particular Postgres database server instance are stored. (Each specific running instance is identified by its host and port number when connecting to a database.) The database cluster is identified by the `-D` option when it is created, started, stopped, etc. by the `Postgres initdb` and `pg_ctl` commands. Typically by default, the initial database cluster is located in the `POSTGRES_INSTALL_HOME/data` directory. A full backup is a copy of a database cluster.

The terms database cluster and database server are used somewhat interchangeably throughout this document, though a single database server can run multiple database clusters.

- **Postgres User Account.** Linux operating system user account that runs the Advanced Server or PostgreSQL database server and owns the database cluster directory.
 - By default, the database user account is `enterprisedb` when Advanced Server is installed to support compatibility with Oracle databases.
 - By default, the database user account is `postgres` when Advanced Server is installed in PostgreSQL compatible mode. For a PostgreSQL database server, the default database user account is also `postgres`.
The BART configuration parameter `cluster_owner` must be set to the database user account for each database server.
- **Replication Database User.** For each database server managed by BART, a database superuser must be selected to act as the replication database user. This database user is used to connect to the database server when backups are taken. The database superusers created with an initial Postgres database server installation (`enterprisedb` or `postgres`) may be used for this purpose.

The BART configuration parameter `user` must be set to this replication database `user` for each database server.

- **Secure Shell (SSH)/Secure Copy (SCP).** Linux utility programs used to log into hosts (SSH) and copy files (SCP) between hosts. A valid user account must be specified that exists on the target host and in effect, is the user account under which the SSH or SCP operations occur.

For information on how all of these components are configured and used with BART, see the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

2.1 Block-Level Incremental Backup

This section describes the basic concepts of a block-level incremental backup (referred to as an incremental backup). An incremental backup is functionality unique to BART.

An incremental backup provides a number of advantages when compared to using a full backup:

- The amount of time required to produce an incremental backup is generally less than a full backup, as modified relation blocks are saved instead of all full relation files of the database cluster.
- An incremental backup uses less disk space than a full backup.

Generally, all BART features (such as retention policy management) apply to incremental backups as well as full backups. See *Managing Incremental Backups* for information about retention policy management as applied to incremental backups.

2.1.1 Incremental Backup Limitations and Requirements

The following limitations apply to incremental backup:

- If you have restored a full or incremental backup, you must take a full backup before enabling incremental backup.
- If a standby node has been promoted to the role of primary node, you must take a full backup before enabling incremental backup on the cluster.
- An incremental backup cannot be taken on a standby database server.

The following requirements must be met before implementing incremental backup:

- You must create or select an operating system account to be used as the BART user account.
- You must create or select the replication database user, which must be a superuser.
- In the BART configuration file:
 - You must set the `cluster_owner` parameter to the Linux operating system user account that owns the database cluster directory from which incremental backups are to be taken.
 - You must enable the `allow_incremental_backups` parameter.
- A passwordless SSH/SCP connection must be established between the BART user account on the BART host and the `cluster_owner` user account on the database server.

It must be noted that a passwordless SSH/SCP connection must be established even if BART and the database server are running on the same host and the BART user account and the `cluster_owner` user account are the same account.

- In addition to the BART host where the BART backup catalog resides, the BART package must also be installed on every remote database server on which incremental backups are to be restored. To restore an incremental backup, the `bart` program must be executable on the remote host by the remote user specified by the `remote_host` parameter in the BART configuration file or by the `-r` option (when using the `RESTORE` subcommand to restore the incremental backup).
- When *restoring incremental backups on a remote database server*, a passwordless SSH/SCP connection must be established from the BART user account on the BART host to the remote user on the remote host that is specified by the `remote_host` parameter in the BART configuration file or by the `-r` option when using the `RESTORE` subcommand to restore the incremental backup.

- Compression of archived WAL files in the BART backup catalog is not permitted for database servers on which incremental backups are to be taken. The `wal_compression` setting in the BART configuration file must be disabled for those database servers.
- The incremental backup must be on the same timeline as the parent backup. The timeline changes after each recovery operation so an incremental backup cannot use a parent backup from an earlier timeline.

For information about setting or configuring all these requirements, see the *Configuration* section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

The following section provides an overview of the basic incremental backup concepts.

2.1.2 Concept Overview

Using incremental backups involves the following sequence of steps:

1. Configure BART (see the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for details about configuring BART).
2. Enable and initiate archiving of WAL files to the `archive_path` in the same manner as done for full backups.

The default `archive_path` is the BART backup catalog (`<backup_path>/<server_name>/archived_wals`). Using the `<archive_path>` parameter in the server section of the BART configuration file, you can specify the location where WAL files will be archived.

For more information about the `archive_path` parameter, see the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

3. Take an initial full backup with the `BACKUP` subcommand. This full backup establishes the parent of the first incremental backup.
4. Scan all WAL files produced by database servers on which incremental backups are to be taken. These WAL files are scanned once they have been archived to the `archive_path`.

Each scanned WAL file results in a modified block map (MBM) file that records the location of modified blocks obtained from the corresponding WAL file. The BART WAL scanner program `bart-scanner` performs this process.

5. Take incremental backups using the `BACKUP` subcommand with the `--parent` option to specify the backup identifier or name of a previous, full backup or an incremental backup. Any previous backup may be chosen as the parent as long as all backups belong to the same timeline.
6. The incremental backup process identifies which WAL files may contain changes from when the parent backup was taken to the starting point of the incremental backup. The corresponding MBM files are used to locate and copy the modified blocks to the incremental backup directory along with other database cluster directories and files. Instead of backing up all, full relation files, only the modified blocks are copied and saved. In addition, the relevant MBM files are condensed into one consolidated block map (CBM) file that is stored with the incremental backup.

Multiple block copier threads can be used to copy the modified blocks to the incremental backup directory. See the *Configuration* section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for information about setting the `thread_count` parameter in the BART configuration file. See *Section - Backup* for using the `--thread-count` option with the `BACKUP` subcommand.

7. Invoke the restore process for an incremental backup using the `RESTORE` subcommand in the same manner as restoring a full backup. The `-i` option specifies the backup identifier or name of the incremental backup to restore. The restore process begins by going back through the chain of past, parent incremental backups until the

initial full backup starting the chain is identified. This full backup provides the initial set of directories and files to be restored to the location specified with the `-p` option. Each subsequent incremental backup in the chain is then restored. Restoration of an incremental backup uses its CBM file to restore the modified blocks from the incremental backup.

The following sections provide some additional information on these procedures.

2.1.3 WAL Scanning – Preparation for an Incremental Backup

The WAL scanner program (`bart-scanner`) scans the WAL files created from the time of the parent backup up to the start of the incremental backup to determine which blocks have modified since the parent backup, and records the information in a file called the *modified block map (MBM) file*. One MBM file is created for each WAL file.

The MBM file is stored in the directory where `archived_wals` will be stored, as specified in the `archive_path` parameter in the `bart.cfg` file. If the `archive_path` is not specified, the default `archived_wals` directory is:

```
<backup_path>/<server_name>/<archived_wals>
```

Where:

`<backup_path>` is the BART backup catalog parent directory specified in the global section of the BART configuration file.

`<server_name>` is the lowercase conversion of the database server name specified in the server section of the BART configuration file.

The following code snippet is the content of the archive path showing the MBM files created for the WAL files. (The user name and group name of the files have been removed from the example to list the WAL files and MBM files in a more comparable manner.)

```
[root@localhost archived_wals]# pwd
/opt/backup/acctg/archived_wals
[root@localhost archived_wals]# ls -l
total 131104
-rw----- 1 ... .. 16777216 Oct 12 09:38 000000010000000100000078
-rw----- 1 ... .. 16777216 Oct 12 09:38 000000010000000100000079
-rw----- 1 ... .. 16777216 Oct 12 09:38 00000001000000010000007A
-rw----- 1 ... .. 16777216 Oct 12 09:35 00000001000000010000007B
-rw----- 1 ... .. 16777216 Oct 12 09:38 00000001000000010000007C
-rw----- 1 ... .. 16777216 Oct 12 09:39 00000001000000010000007D
-rw----- 1 ... .. 16777216 Oct 12 09:42 00000001000000010000007E
-rw----- 1 ... .. 16777216 Oct 12 09:47 00000001000000010000007F
-rw-rw-r-- 1 ... .. 161 Oct 12 09:49 0000000100000001780000280000000179000000.mbm
-rw-rw-r-- 1 ... .. 684 Oct 12 09:49 000000010000000179000028000000017A000000.mbm
-rw-rw-r-- 1 ... .. 161 Oct 12 09:49 00000001000000017A000028000000017B000000.mbm
-rw-rw-r-- 1 ... .. 161 Oct 12 09:49 00000001000000017B000028000000017C000000.mbm
-rw-rw-r-- 1 ... ..1524 Oct 12 09:49 00000001000000017C000028000000017D000000.mbm
-rw-rw-r-- 1 ... .. 161 Oct 12 09:49 00000001000000017D000028000000017E000000.mbm
-rw-rw-r-- 1 ... .. 161 Oct 12 09:49 00000001000000017E000028000000017F000000.mbm
-rw-rw-r-- 1 ... .. 161 Oct 12 09:49 00000001000000017F0000280000000180000000.mbm
```

MBM files have the suffix, `.mbm`.

In preparation for any incremental backup, the WAL files should be scanned as soon as they are copied to the `archive_path`. Thus, the WAL scanner should be running as soon as the WAL files from the database cluster are archived to the `archive_path`.

If the `archive_path` contains WAL files that have not yet been scanned, starting the WAL scanner begins scanning these files. If WAL file fails to be scanned (resulting in a missing MBM file), you can use the WAL scanner to specify an individual WAL file.

Under certain conditions (such as when the `rsync` utility is used to copy WAL files to the `archive_path`), the WAL files may have been missed by the WAL scanner program for scanning and creation of MBM files. Use the `scan_interval` parameter in the BART configuration file to initiate force scanning of WAL files in the `archive_path` to ensure MBM files are generated. See the *Configuration* section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for more information about the `scan_ineterval` parameter.

See *Running the BART WAL Scanner* for information about using the WAL scanner.

2.1.4 Performing an Incremental Backup

The WAL files produced at the time of the parent backup up to the start of the incremental backup contain information about which blocks were modified during that time interval. That information is consolidated into an MBM file for each WAL file by the WAL scanner.

The MBM files for the relevant WAL files are read, and the information is used to copy the modified blocks from the database cluster to the `archived_wals` directory as specified in the `archive_path` parameter in the `bart.cfg` file. When compared to a full backup, the number and sizes of relation files can be significantly less for an incremental backup.

For comparison, the following is an abbreviated list of the files copied to the archived `base` subdirectory of a full backup for one database:

```
[root@localhost 14845]# pwd
/opt/backup/acctg/1476301238969/base/base/14845
[root@localhost 14845]# ls
112          13182_vm    14740  16467  16615      2608_vm    2655  2699      2995      ...
113          13184      14742  16471  174        2609       2656  2701      2995_vm   ...
1247        13186      14745  16473  175        2609_fsm   2657  2702      2996      ...
1247_fsm    13187      14747  16474  2187       2609_vm    2658  2703      2998      ...
1247_vm     13187_fsm  14748  16476  2328       2610       2659  2704      2998_vm   ...
1249        13187_vm    14749  16477  2328_fsm   2610_fsm   2660  2753      2999      ...
1249_fsm    13189      14752  16479  2328_vm    2610_vm    2661  2753_fsm  2999_vm   ...
1249_vm     13191      14754  16488  2336       2611       2662  2753_vm   3079      ...
1255        13192      14755  16490  2336_vm    2611_vm    2663  2754      3079_fsm  ...
.
.
.
13182_fsm   14739      16465  16603  2608_fsm   2654       2696  2893_vm   3501_vm   ...
```

In contrast, the following is the content of the archived `base` subdirectory of the same database from a subsequent incremental backup:

```
[root@localhost 14845]# pwd
/opt/backup/acctg/1476301835391/base/base/14845
[root@localhost 14845]# ls
1247        1249        1259        16384  17006  2608      2610      2658  2663  2678  ...
1247_fsm    1249_fsm    1259_fsm    16387  17009  2608_fsm  2610_fsm  2659  2673  2679  ...
1247_vm     1249_vm     1259_vm     16389  17011  2608_vm   2610_vm   2662  2674  2703  ...
```

The information from the MBM files are consolidated into one file called a *consolidated block map* (CBM) file. During the restore operation for the incremental backup, the CBM file is used to identify the modified blocks to be restored for that backup.

In addition, the incremental backup also stores other required subdirectories and files from the database cluster as is done for full backups.

Before taking an incremental backup, an initial full backup must be taken with the `BACKUP` subcommand. This full backup establishes the parent of the first incremental backup.

The syntax for taking a full backup is:

```
bart BACKUP -s { <server_name> | all } [ -F { p | t } ]
  [ -z ] [ -c <compression_level> ]
  [ --backup-name <backup_name> ]
  [ --thread-count <number_of_threads> ]
  [ { --with-pg_basebackup | --no-pg_basebackup } ]
```

Note: While a BACKUP subcommand is in progress, no other processes must run in parallel.

The syntax for taking an incremental backup is:

```
bart BACKUP -s { <server_name> | all } [ -F p]
  [ --parent { <backup_id> | <backup_name> } ]
  [ --backup-name <backup_name> ]
  [ --thread-count <number_of_threads> ]
  [ --check ]
```

You must specify the following before taking an incremental backup:

- `-F p` option for plain text format as incremental backup can only be taken in the plain text format.
- `--check` option to verify if the required MBM files are present in the `archived_wals` directory. The `--parent` option must be specified when the `--check` option is used.

See the *BACKUP* section for more information about using the BACKUP subcommand.

2.1.5 Restoring an Incremental Backup

Restoring an incremental backup may require additional steps depending upon the host on which the incremental backup is to be restored:

- *Restoring an Incremental Backup on a BART Host* - This section outlines restoring an incremental backup onto the same host where BART has been installed.
- *Restoring an Incremental Backup on a Remote Host* - This section outlines restoring an incremental backup onto a remote host where BART has not been installed.

Ensure the `bart` program is available on the remote host when restoring an incremental backup on a remote host; the invocation of the `RESTORE` subcommand for an incremental backup results in the execution of the `bart` program on the remote host to restore the modified blocks to their proper location.

Restoring an Incremental Backup on a BART Host

Specify a backup identifier or name, and include the `-i` option when invoking the `RESTORE` subcommand to restore an incremental backup. All `RESTORE` options may be used in the same manner as when restoring a full backup.

First, all files from the full backup from the beginning of the backup chain are restored. For each incremental backup, the CBM file is used to identify and restore blocks from the incremental backup. If there are new relations or databases identified in the CBM file, then relevant relation files are copied. If consolidated block map information is found indicating the drop of a relation or a database, then the relevant files are removed from the restore directory. Similarly, if there is any indication of a table truncation, then the related files are truncated.

Also note that you can use the `-w` option of the `RESTORE` subcommand to specify a multiple number of parallel worker processes to stream the modified blocks to the restore host.

Restoring an Incremental Backup on a Remote Host

To restore an incremental backup onto a remote host where BART has not been installed, first follow the restore process steps outlined in *Restoring an Incremental Backup on a BART Host* and then, perform the following steps:

Step 1: Install BART on the remote host to which an incremental backup is to be restored.

No editing is needed in the `bart.cfg` file installed on the remote host.

Step 2: Determine the Linux operating system user account on the remote host to be used as the remote user. This user is specified by the `remote_host` parameter in the BART configuration file or by the `-r` option when using the `RESTORE` subcommand to restore the incremental backup. The remote user must be the owner of the directory where the incremental backup is to be restored on the remote host. By default, the user account is `enterprisedb` for Advanced Server or `postgres` for PostgreSQL.

Step 3: Ensure a passwordless SSH/SCP connection is established from the BART user on the BART host to the remote user on the remote host.

See the configuration section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

Example

When restoring an incremental backup, specify the `RESTORE` subcommand and the backup identifier or name of the incremental backup that will be restored. For example:

```
-bash-4.2$ bart RESTORE -s acctg -p /opt/restore -i incr_1-b
INFO: restoring incremental backup 'incr_1-b' of server 'acctg'
INFO: base backup restored
INFO: archiving is disabled
INFO: permissions set on $PGDATA
INFO: incremental restore completed successfully
```

Restoring incremental backup `incr_1-b` as shown in the preceding example results in the restoration of full backup `full_1`, then incremental backups `incr_1-a` and finally, `incr_1-b`.

2.2 Creating a Backup Chain

A *backup chain* is the set of backups consisting of a full backup and all of its successive incremental backups. Tracing back on the parent backups of all incremental backups in the chain eventually leads back to that single, full backup.

It is possible to have a *multi-forked* backup chain, which is two or more successive lines of incremental backups, all of which begin with the same, full backup. Thus, within the chain there is a backup that serves as the parent of more than one incremental backup.

Since restoration of an incremental backup is dependent upon first restoring the full backup, then all successive incremental backups up to, and including the incremental backup specified by the `RESTORE` subcommand, it is crucial to note the following:

- Deletion or corruption of the full backup destroys the entire backup chain. It is not possible to restore any of the incremental backups that were part of that chain.
- Deletion or corruption of an incremental backup within the chain results in the inability to restore any incremental backup that was added to that successive line of backups following the deleted or corrupted backup. The full backup and incremental backups prior to the deleted or corrupted backup can still be restored.

The actions of retention policy management are applied to the full backup and all of its successive incremental backups within the chain in an identical manner as if they were one backup. Thus, use of retention policy management does not result in the breakup of a backup chain.

See the *EDB Postgres Backup and Recovery Reference Guide* for examples of creating a backup chain and restoring an incremental backup. The reference guide is available at:

<https://www.enterprisedb.com/edb-docs/>

After installing and configuring the BART host and the database servers, you can start using BART. For detailed information about installation and configuration, see the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs>

This section describes how to perform backup and recovery management operations using BART. Review the sections that follow before proceeding with any BART operation.

3.1 BART Management Overview

After configuring BART, you can begin the BART backup and recovery management process. The following steps will help you get started:

1. Run the `CHECK-CONFIG` subcommand without the `-s` option. When the `CHECK-CONFIG` subcommand is used without specifying the `-s` option, it checks the parameters in the global section of the BART configuration file.
2. Run the `INIT` subcommand (if you have not already done so) to finish creation of the BART backup catalog, which results in the complete directory structure to which backups and WAL files are saved. This step must be done before restarting the database servers with enabled WAL archiving, otherwise the copy operation in the `archive_command` parameter of the `postgresql.conf` file or the `postgresql.auto.conf` file fails due to the absence of the target archive directory. When the directory structure is complete, the `archived_wals` subdirectory should exist for each database server.
3. Start the Postgres database servers with archiving enabled. Verify that the WAL files are appearing in the `archive_path`. For more information about the `archive_path` parameter, see the Configuration section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

The archiving frequency is dependent upon other `postgresql.conf` configuration parameters. Check the Postgres database server log files to ensure there are no archiving errors. Archiving should be operational before taking a backup in order to ensure that the WAL files that may be created during the backup process are archived.

4. Start the WAL scanner if you intend to take incremental backups. Since the WAL scanner processes the WAL files copied to the `archive_path`, it is advantageous to commence the WAL scanning as soon as the WAL files begin to appear in the `archive_path` in order to keep the scanning in pace with the WAL archiving.
5. Run the `BART CHECK-CONFIG` subcommand for each database server with the `-s` option specifying the server name. This ensures the database server is properly configured for taking backups.
6. Create a full backup for each database server. The full backup establishes the starting point of when point-in-time recovery can begin and also establishes the initial parent backup for any incremental backups to be taken.

There are now a number of other BART management processes you may perform:

- Execute the `BACKUP` subcommand to create additional full backups or incremental backups.
- Use the `VERIFY-CHECKSUM` subcommand to verify the checksum of the full backups.
- Display database server information with the `SHOW-SERVERS` subcommand.
- Display backup information with the `SHOW-BACKUPS` subcommand.
- Compress the archived WAL files in the `archive_path` by enabling WAL compression in the BART configuration file and then invoking the `MANAGE` subcommand.
- Determine and set the retention policy for backups in the BART configuration file.
- Establish the procedure for using the `MANAGE` subcommand to enforce the retention policy for backups. This may include using `cron` jobs to schedule the `MANAGE` subcommand.

3.1.1 Performing a Restore Operation

Perform the following steps for restore operation:

Step 1: Stop the Postgres database server on which you will be performing the restore operation.

Step 2: Inspect the `pg_xlog` subdirectory of the data directory and save any WAL files that have not yet been archived to the `archive_path`. If there are files that have not been archived, save them to a temporary location.

Step 3: If you are restoring to:

- the current data directory, delete all files and subdirectories under the data directory.
- a new directory, create the directory on which you want to restore the backed up database cluster. Ensure the data directory can be written to by the BART user account or by the user account specified by the `remote_host` configuration parameter, or by the `--remote-host` option of the `RESTORE` subcommand (if these are to be used).

Step 4: Perform the same process for tablespaces as described in Step 3. The `tablespace_path` parameter in the BART configuration file must contain the tablespace directory paths to which the tablespace data files are to be restored. See the *Configuration* section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for more information.

Step 5: Identify the timeline ID you wish to use to perform the restore operation.

The available timeline IDs can be identified by the first non-zero digit of the WAL file names reading from left to right.

Step 6: Identify the backup to use for the restore operation and obtain the backup ID or backup name.

To use the latest backup, omit the `-i` option; the `RESTORE` subcommand uses that backup by default. The backups can be listed with the `SHOW-BACKUPS` subcommand.

Step 7: Run the BART `RESTORE` subcommand.

- If you have specified the `-t timeline_id`, `-x target_xid`, or `-g target_timestamp` option, then recovery settings will be saved in the `postgresql.auto.conf` file and point-in-time recovery will be performed when you restart the database server.
- If you do not specify `-t timeline_id`, `-x target_xid`, or `-g target_timestamp`, then minimal recovery settings will be saved in the `postgresql.auto.conf` file and archive recovery will proceed only until consistency is reached, with no restoration of files from the archive. See the *Restore* section for detailed information about `Restore` subcommand.
- If the `-c` option is specified or if the enabled setting of the `copy_wals_during_restore` BART configuration parameter is in effect for this database server, then the following actions occur:
 - In addition to restoring the database cluster to the directory specified by the `-p restore_path` option, the archived WAL files of the backup are copied from the BART backup catalog to the subdirectory `restore_path/archived_wals`.
 - If recovery settings are saved in the `postgresql.auto.conf` file, the command string set in the `restore_command` parameter retrieves the WAL files from this `archived_wals` subdirectory relative to the `restore_path` parent directory as: `restore_command = cp archived_wals/%f %p`

You must ensure that valid options are specified when using the `RESTORE` subcommand. BART will not generate an error message if invalid option values or invalid option combinations are provided. BART will accept the invalid options and pass them to the `postgresql.auto.conf` file, which will then be processed by the database server when it is restarted.

Step 8: Copy any saved WAL files from Step 2 to the `restore_path/pg_xlog` subdirectory.

Step 9: Inspect the restored directories and data files of the restored database cluster in directory `restore_path`.

All files and directories must be owned by the user account that you intend to use to start the database server. Recursively change the user and group ownership of the `restore_path` directory, its files, and its subdirectories if necessary. There must only be directory access privileges for the user account that will start the database server. No other groups or users can have access to the directory.

Step 10: Inspect the `postgresql.auto.conf` file (if you are performing a point-in-time recovery) located in the `restore_path` directory to verify if it contains the appropriate parameter settings to recover to the indicated point. Otherwise, the `postgresql.auto.conf` file should be configured to recover only until the cluster reaches consistency. In either case, the settings may be modified as desired.

Step 11: Disable WAL archiving at this point. The BART RESTORE subcommand adds `archive_mode = off` to the end of the `postgresql.conf` file.

- If you want to restart the database server with WAL archiving activated, ensure that this additional parameter is deleted.
- The original `archive_mode` parameter still resides in the `postgresql.conf` file in its initial location with its last setting.

Step 12: Start the database server to initiate recovery. After completion, check the database server log file to ensure the recovery was successful.

If the backup is restored to a different location than where the original database cluster resided, operations dependent upon the database cluster location may fail if supporting service scripts are not updated to reflect the location where the backup has been restored. For information about the use and modification of service scripts, see the *EDB Postgres Advanced Server Installation Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

3.1.2 Point-In-Time Recovery Operation

The following steps outline how to perform a point-in-time recovery operation for a database cluster:

1. Use your system-specific command to shut down the database server.
2. If you want to restore the database cluster and tablespace files to new directories, create the new directories with the appropriate directory ownership and permissions.

If you plan to reuse existing database cluster directories, delete all the files and subdirectories in the existing directories. We strongly recommend that you make a copy of this data before deleting it. Be sure to save any recent WAL files in the `pg_xlog` subdirectory that have not been archived to `archive_path`.

3. Run the `BART SHOW-BACKUPS -s <server_name>` subcommand to list the backup IDs and backup names of the backups for the database server. You will need to provide the appropriate backup ID or backup name with the `BART RESTORE` subcommand, unless you intend to restore the latest backup in which case the `-i` option of the `RESTORE` subcommand for specifying the backup ID or backup name may be omitted.
4. Run the `BART RESTORE` subcommand with the appropriate options.

- The backup is restored to the directory specified by the `-p restore_path` option.
- In addition, if the `RESTORE` subcommand `-c` option is specified or if the enabled setting of the `copy_wals_during_restore` BART configuration parameter is applicable to the database server, then the required archived WAL files from the `archive_path` are copied to the `restore_path/archived_wals` subdirectory.

Ensure the `restore_path` directory and all subdirectories and files in the `restore_path` are owned by the proper Postgres user account (for example, `enterprisedb` or `postgres`). Also ensure that only the Postgres user account has access permission to the `restore_path` directory.

Use the `chown` command to make the appropriate adjustments to file permissions; for example, the following command changes the ownership of `restore_path` to `enterprisedb`:

```
chown -R enterprisedb:enterprisedb restore_path
```

The following command restricts access to `restore_path`:

```
chmod 700 restore_path
```

5. Copy any saved WAL files from Step 2 that were not archived to the BART backup catalog to the `restore_path/pg_xlog` subdirectory.
6. For point-in-time recovery, verify that the `postgresql.auto.conf` file created in the directory specified with the `RESTORE` subcommand's `-p <restore_path>` option generated the correct recovery parameter settings.

If the `RESTORE` subcommand `-c` option is specified or if the enabled setting of the `copy_wals_during_restore` BART configuration parameter is applicable to the database server, then the `restore_command` parameter retrieves the archived WAL files from the `<restore_path>/archived_wals` subdirectory that was created by the `RESTORE` subcommand, otherwise the `restore_command` retrieves the archived WAL files from the BART backup catalog.

7. The `BART RESTORE` subcommand disables WAL archiving in the restored database cluster. If you want to immediately enable WAL archiving, modify the `postgresql.conf` file by deleting the `archive_mode = off` parameter that BART appends to the end of the file.
8. Start the database server, which will then perform the point-in-time recovery operation if recovery settings are saved in the `postgresql.auto.conf` file.

A detailed description of the `RESTORE` subcommand is available in the *Basic Bart Subcommand Usage* section.

Example of a Point-in-Time recovery operation is documented in the *EDB Postgres Backup and Recovery Reference Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

3.2 Managing Backups Using a Retention Policy

Over the course of time when using BART, the number of backups can grow significantly. This ultimately leads to a large consumption of disk space unless an administrator periodically performs the process of deleting the oldest backups that are no longer needed. This process of determining when a backup is old enough to be deleted and then actually deleting such backups can be done and automated with the following basic steps:

1. Determine and set a retention policy in the BART configuration file. A *retention policy* is a rule that determines when a backup is considered obsolete. The retention policy can be applied globally to all servers, but each server can override the global retention policy with its own (see the configuration section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for details).
2. Use the `MANAGE` subcommand to categorize and manage backups according to the retention policy.
3. Create a cron job to periodically run the `MANAGE` subcommand to evaluate the backups and then list and/or delete the obsolete backups.

There is a difference on how retention policy management applies to incremental backups as compared to full backups. See *Section - Managing Incremental Backups* for information about how retention policy management applied to full backups affects incremental backups.

The following sections describe how retention policy management generally applies to backups, and its specific usage and effect on full backups:

- *Section - Overview - Managing Backups Using a Retention Policy* provides an overview of the terminology and types of retention policies.
- *Section - Marking the Backup Status* describes the concept of marking the status of backups according to the retention policy.
- *Section - Setting the Retention Policy* describes setting the different types of retention policies.
- *Section - Managing the Backups Based on the Retention Policy* describes the process of managing the backups such as marking the backup status, keeping selected backups indefinitely, listing obsolete backups, and deleting obsolete backups.

3.2.1 Overview - Managing Backups Using a Retention Policy

The BART retention policy results in the categorization of each backup in one of three statuses –*active*, *obsolete*, and *keep*:

- **Active.** The backup satisfies the retention policy applicable to its server. Such backups would be considered necessary to ensure the recovery safety for the server and thus should be retained.
- **Obsolete.** The backup does not satisfy the retention policy applicable to its server. The backup is no longer considered necessary for the recovery safety of the server and thus can be deleted.
- **Keep.** The backup is to be retained regardless of the retention policy applicable to its server. The backup is considered vital to the recovery safety for the server and thus should not be deleted for an indefinite period of time.

There are two types of retention policies - redundancy retention policy and recovery window retention policy.

- **Redundancy Retention Policy** - The *redundancy retention policy* relies on a specified, maximum number of most recent backups to retain for a given server. When the number of backups exceeds that maximum number, the oldest backups are considered obsolete (except for backups marked as keep).
- **Recovery Window Retention Policy** - The *recovery window retention policy* relies on a time frame (the recovery window) for when a backup should be considered active. The boundaries defining the recovery window are

the current date/time (the ending boundary of the recovery window) and the date/time going back in the past for a specified length of time (the starting boundary of the recovery window).

- If the date/time the backup was taken is within the recovery window (that is, the backup date/time is on or after the starting date/time of the recovery window), then the backup is considered active, otherwise it is considered obsolete (except for backups marked as keep).
- Thus, for the recovery window retention policy, the recovery window time frame dynamically shifts, so the end of the recovery window is always the current date/time when the `MANAGE` subcommand is run. As you run the `MANAGE` subcommand at future points in time, the starting boundary of the recovery window moves forward in time. At some future point, the date/time of when a backup was taken will be earlier than the starting boundary of the recovery window. This is when an active backup's status will then be considered obsolete.
- You can see the starting boundary of the recovery window at any point in time by running the `SHOW-SERVERS` subcommand. The `RETENTION_POLICY` field of the `SHOW-SERVERS` subcommand displays the starting boundary of the recovery window.

3.2.2 Marking the Backup Status

When a backup is initially created with the `BACKUP` subcommand, it is always recorded with active status. Use the `MANAGE` subcommand to evaluate if the backup status should be changed to obsolete in accordance with the retention policy. See *MANAGE* for usage information of the `MANAGE` subcommand.

Active backups are evaluated and also marked (that is, internally recorded by BART) as obsolete only when the `MANAGE` subcommand is invoked either with no options or with only the `-s` option.

Once a backup has been marked as obsolete, you cannot change it back to active unless you perform the following steps:

- Use the `MANAGE` subcommand with the `-c` option along with the backup identifier or name with the `-i` option. To keep this particular backup indefinitely, use `-c keep`, otherwise use `-c nokeep`.
- If you use the `-c nokeep` option, the backup status is changed back to active. When the `MANAGE` subcommand is used the next time, the backup is re-evaluated to determine if its status needs to be changed back to obsolete based on the current retention policy in the BART configuration file.

After setting the `retention_policy` parameter and running the `MANAGE` subcommand if you change the `retention_policy` parameter, the current, marked status of the backups are probably inconsistent with the new `retention_policy` setting. To modify the backup status to be consistent with the new `retention_policy` setting, you need to run the `MANAGE` subcommand with:

- the `-c nokeep` option to change the obsolete status to active status if there are backups currently marked as obsolete that would no longer be considered obsolete under a new retention policy. You can also specify the `-i all` option to change all backups back to active status, including those currently marked as keep.
- no options or with only the `-s` option to reset the marked status based on the new `retention_policy` setting in the BART configuration file.

You can review the current marking (the status) of your backups with the `SHOW-BACKUPS` subcommand.

3.2.3 Setting the Retention Policy

The retention policy is determined by the `retention_policy` parameter in the BART configuration file. It can be applied globally to all servers, but each server can override the global retention policy with its own. For information about creating a global retention policy and an individual database server retention policy, see the *Configuration* section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

There are two types of retention policies - redundancy retention policy and the recovery window retention policy as described in the following sections.

Redundancy Retention Policy

To use the redundancy retention policy, set `retention_policy = max_number BACKUPS` where `max_number` is a positive integer designating the maximum number of most recent backups.

Additional Restrictions:

- The keyword `BACKUPS` must always be specified in plural form (for example, `1 BACKUPS`).
- BART will accept a maximum integer value of 2,147,483,647 for `max_number`; however, you should use a realistic, practical value based on your system environment.

The redundancy retention policy is the default type of retention policy if all keywords `BACKUPS`, `DAYS`, `WEEKS`, and `MONTHS` following the `max_number` integer are omitted as shown by the following example:

```
retention_policy = 3
```

In the following example the redundancy retention policy setting considers the three most recent backups as the active backups. Any older backups, except those marked as `keep`, are considered obsolete.

```
[ACCTG]
host = 127.0.0.1
port = 5444
user = enterprisedb
archive_command = 'cp %p %a/%f'
retention_policy = 3 BACKUPS
description = "Accounting"
```

The `SHOW-SERVERS` subcommand displays the 3 Backups redundancy retention policy in the `RETENTION POLICY` field:

```
-bash-4.1$ bart SHOW-SERVERS -s acctg
SERVER NAME      : acctg
HOST NAME       : 127.0.0.1
USER NAME       : enterprisedb
PORT            : 5444
REMOTE HOST     :
RETENTION POLICY : 3 Backups
DISK UTILIZATION : 627.04 MB
NUMBER OF ARCHIVES : 25
ARCHIVE PATH    : /opt/backup/acctg/archived_wals
ARCHIVE COMMAND : cp %p /opt/backup/acctg/archived_wals/%f
XLOG METHOD     : fetch
WAL COMPRESSION : disabled
TABLESPACE PATH(s) :
DESCRIPTION    : "Accounting"
```

Recovery Window Retention Policy

To use the recovery window retention policy, set the `retention_policy` parameter to the desired length of time for the recovery window in one of the following ways:

- Set to `max_number DAYS` to define the start date/time recovery window boundary as the number of days specified by `max_number` going back in time from the current date/time.
- Set to `max_number WEEKS` to define the start date/time recovery window boundary as the number of weeks specified by `max_number` going back in time from the current date/time.
- Set to `max_number MONTHS` to define the start date/time recovery window boundary as the number of months specified by `max_number` going back in time from the current date/time.

Additional Restrictions:

- The keywords `DAYS`, `WEEKS`, and `MONTHS` must always be specified in plural form (for example, `1 DAYS`, `1 WEEKS`, or `1 MONTHS`).
- BART will accept a maximum integer value of `2,147,483,647` for `max_number`, however, a realistic, practical value based on your system environment must always be used.

A backup is considered active if the date/time of the backup is equal to or greater than the start of the recovery window date/time.

You can view the actual, calculated recovery window by:

- Invoking the `MANAGE` subcommand in debug mode, along with the `-n` option.
- Using the `SHOW-SERVERS` subcommand.

See the *EDB Postgres Backup and Recovery Reference Guide* for examples of how to view the actual, calculated recovery window.

3.2.4 Managing the Backups Based on the Retention Policy

The `MANAGE` subcommand is used to evaluate and categorize backups according to the retention policy set in the BART configuration file. When a backup is first created with the `BACKUP` subcommand, it is always marked as active. By using the `MANAGE` subcommand, an active backup may be marked as obsolete. Obsolete backups can then be deleted.

This section covers following aspects of backup management:

- The rules for *deleting backups* depending upon the backup status and the subcommand used.
- The process to retain a backup indefinitely by *marking it as keep*. This section also provides information about resetting backups status (that are marked as `obsolete` and `keep`) back to active status.
- The general process for *evaluating, marking, and then deleting obsolete backups*.

Deletions Permitted Under a Retention Policy

This section describes how and under what conditions backups may be deleted under a retention policy.

You must use the `MANAGE` subcommand to delete obsolete backups. Use the `DELETE` subcommand only for special administrative purposes.

The deletion behavior of the `MANAGE` subcommand and the `DELETE` subcommand are based on different aspects of the retention policy.

- The `MANAGE` subcommand deletion relies solely upon how a backup status is currently marked (that is, internally recorded by BART). The current setting of the `retention_policy` parameter in the BART configuration file is ignored.
- The `DELETE` subcommand relies solely upon the current setting of the `retention_policy` parameter in the BART configuration file. The current active, obsolete, or keep status of a backup is ignored.

The specific deletion rules for the `MANAGE` and `DELETE` subcommands are as follows:

- `MANAGE` subcommand: The `MANAGE` subcommand with the `-d` option can only delete backups marked as obsolete. This deletion occurs regardless of the current `retention_policy` setting in the BART configuration file. The deletion of backups relies on the last occasion when the backups have been marked.
- `DELETE` subcommand:
 - Under a redundancy retention policy currently set with the `retention_policy` parameter in the BART configuration file, any backup regardless of its marked status, can be deleted with the `DELETE` subcommand when the backup identifier or name is specified with the `-i` option and if the current total number of backups for the specified database server is greater than the maximum number of redundancy backups currently specified with the `retention_policy` parameter.

If the total number of backups is less than or equal to the specified, maximum number of redundancy backups, then no additional backups can be deleted using `DELETE` with the `-i` backup option.

- Under a recovery window retention policy currently set with the `retention_policy` parameter in the BART configuration file, any backup regardless of its marked status, can be deleted with the `DELETE` subcommand when the backup identifier or name is specified with the `-i` option, and if the backup date/time is not within the recovery window currently specified with the `retention_policy` parameter. If the backup date/time is within the recovery window, then it cannot be deleted using `DELETE` with the `-i` backup option.
- Invoking the `DELETE` subcommand with the `-i all` option results in the deletion of all backups regardless of the retention policy and regardless of whether the status is marked as active, obsolete, or keep.

The following table summarizes the deletion rules of backups according to their marked status. An entry of `Yes` indicates the backup may be deleted under the specified circumstances. An entry of `No` indicates that the backup may not be deleted.

Table 3-1 Allowable Backup Deletion by Status

Operation	Redundancy Retention Policy			Recovery Window Retention Policy		
	Active	Obsolete	Keep	Active	Obsolete	Keep
<code>MANAGE -d</code>	No	Yes	No	No	Yes	No
<code>DELETE -i backup</code>	Yes (see <i>Note1</i>)	Yes (see <i>Note1</i>)	Yes (see <i>Note1</i>)	Yes (see <i>Note2</i>)	Yes (see <i>Note2</i>)	Yes (see <i>Note2</i>)
<code>DELETE -i all</code>	Yes	Yes	Yes	Yes	Yes	Yes

Note: Redundancy Retention Policy (Note1) : Deletion occurs only if the total number of backups for the specified database server is greater than the specified, maximum number of redundancy backups currently set with the `redundancy_policy` parameter in the BART configuration file.

Note: Recovery Window Retention Policy (Note2): Deletion occurs only if the backup is not within the recovery window currently set with the `redundancy_policy` parameter in the BART configuration file.

Marking Backups for Indefinite Keep Status

There may be certain backups that you wish to keep for an indefinite period of time and do not wish to delete based upon the retention policy applied to the database server. Such backups can be marked as `keep` to exclude them from being marked as obsolete. Use the `MANAGE` subcommand with the `-c keep` option to retain such backups indefinitely.

Evaluating, Marking, and Deleting Obsolete Backups

When the `MANAGE` subcommand is invoked, BART evaluates active backups:

- If you include the `-s` option when invoking the `MANAGE` subcommand, BART evaluates backups for the database server.
- If you include the `-s all` option when invoking the `MANAGE` subcommand, BART evaluates backups for all database servers.
- If the `-s` option is omitted, the command evaluates the current number of backups for the database server based on the redundancy retention policy or the current date/time for a recovery window retention policy.

Note: The status of backups currently marked as `obsolete` or `keep` is not changed. To re-evaluate such backups and then classify them, their status must first be reset to `active` with the `MANAGE -c nokeep` option. See *Marking the Backup Status* for more information.

See the *EDB Postgres Backup and Recovery Reference Guide* to view an example of how to evaluate, mark, and delete backups using redundancy retention policy and recovery window retention policy, as well as view examples of `MANAGE` subcommand.

3.2.5 Managing Incremental Backups

The following section summarizes how retention policy management affects incremental backups.

- The retention policy rules are applied to full backups.
 - A redundancy retention policy uses the number of full backups to determine if a backup is obsolete. Incremental backups are excluded from the comparison count against the `retention_policy` setting for the maximum number of backups.
 - A recovery window retention policy uses the backup date/time of any full backups to determine if a backup is obsolete. The backup date/time of any successive incremental backups in the chain are ignored when comparing with the recovery window.
- The retention status of all incremental backups in a chain is set to the same status applied to the full backup of the chain.
- The actions applied by the `MANAGE` and `DELETE` subcommands on a full backup are applied to all incremental backups in the chain in the same manner.
- Thus, a backup chain (that is, the full backup and all its successive incremental backups) are treated by retention policy management as if they are all one, single backup.
 - The status setting applied to the full backup is also applied to all incremental backups in its chain.

- If a full backup is marked as obsolete and then deleted according to the retention policy, all incremental backups in the chain are also marked obsolete and then deleted as well.

The following are some specific points regarding the `MANAGE` and `DELETE` subcommands on incremental backups.

- `MANAGE` subcommand:
 - When the `MANAGE` subcommand is invoked, the status applied to the full backup is also applied to all successive incremental backups in the chain.
 - The `MANAGE` subcommand with the `-c { keep | nokeep }` option cannot specify the backup identifier or backup name of an incremental backup with `-i` backup option. The `-i` backup option can only specify the backup identifier or backup name of a full backup.
 - You can also use the `-i all` option to take a backup of all backups. When the `MANAGE` subcommand with the `-c { keep | nokeep }` option is applied to a full backup, the same status change is made to all incremental backups in the chain.
- `DELETE` subcommand:
 - The `DELETE` subcommand with the `-s server -i` backup option specifies the backup identifier or backup name of an incremental backup in which case that incremental backup along with all its successive incremental backups are deleted, thus shortening that backup chain.

Using a Redundancy Retention Policy with Incremental Backups

When a *redundancy retention policy* is used and the `MANAGE` subcommand is invoked, the status of the oldest `active` full backup is changed to `obsolete` if the number of full backups exceeds the maximum number specified by the `retention_policy` parameter in the BART configuration file.

Note: When a full backup is changed from `active` to `obsolete`, all successive incremental backups in the chain of the full backup are also changed from `active` to `obsolete`.

When determining the number of backups that exceeds the number specified by the `retention_policy` parameter, only full backups are counted for the comparison. Incremental backups are not included in the count for the comparison against the `retention_policy` parameter setting.

See the *EDB Postgres Backup and Recovery Reference Guide* for examples demonstrating use of the `MANAGE` and `DELETE` subcommands when a redundancy retention policy is in effect.

Using a Recovery Window Retention Policy with Incremental Backups

If the `MANAGE` command is invoked when BART is configured to use a *recovery window retention policy*, the status of `active` full backups are changed to `obsolete` if the date/time of the full backup is outside of the recovery window.

Note: If a full backup is changed from `active` to `obsolete`, all successive incremental backups in the chain of the full backup are also changed from `active` to `obsolete`.

The status of an incremental backup is changed to `obsolete` regardless of whether or not the date/time of when the incremental backup was taken still lies within the recovery window.

See the *EDB Postgres Backup and Recovery Reference Guide* for examples demonstrating use of the `MANAGE` and `DELETE` subcommands when a recovery window retention policy is in effect.

3.3 Basic BART Subcommand Usage

This section briefly describes the BART subcommands and options. You can invoke the `bart` program (located in the `<BART_HOME>/bin` directory) with the desired options and subcommands to manage your BART installation.

Syntax for invoking BART:

```
bart [ general_option ]... [ subcommand ] [subcommand_option ]...
```

Please note:

- When invoking a subcommand, the subcommand name is not case-sensitive (that is, the subcommand can be specified in uppercase, lowercase, or mixed case).
- Each subcommand has a number of its own applicable options that are specified following the subcommand. All options are available in both single-character and multi-character forms.
- Keywords are case-sensitive; options are generally specified in lowercase unless specified otherwise in this section.
- When invoking BART, the current user must be the BART user account (operating system user account used to run the BART command line program). For example, `enterprisedb` or `postgres` can be selected as the BART user account when the managed database servers are Advanced Server or PostgreSQL respectively.
- The chosen operating system user account must have the following capabilities:
 - Own the BART backup catalog directory.
 - Be able to run the `bart` program and the `bart scanner` program.
 - Have a passwordless SSH/SCP connection established between database servers managed by BART.

You can specify one or more of the following general options:

Options

- `-h` or `-help`
Displays general syntax and information on BART usage.
All subcommands support a help option (`-h`, `--help`). If the help option is specified, information is displayed regarding that particular subcommand. The subcommand, itself, is not executed.
- `-v` or `-version`
Displays the BART version information.
- `-d` or `-debug`
Displays debugging output while executing BART subcommands.
- `-c` or `--config-path config_file_path`
Specifies `config_file_path` as the full directory path to a BART configuration file. Use this option if you do not want to use the default BART configuration file `<BART_HOME>/etc/bart.cfg`.

Troubleshooting: Setting Path Environment Variable

If execution of BART subcommands fails with the following error message, then you need to set the `LD_LIBRARY_PATH` to include the `libpq` library directory:

```
./bart: symbol lookup error: ./bart: undefined symbol: PQping
```

Workaround: Set the `LD_LIBRARY_PATH` environment variable for the BART user account to include the directory containing the `libpq` library. This directory is `POSTGRES_INSTALL_HOME/lib` as shown by the following example:

```
export LD_LIBRARY_PATH=/opt/PostgresPlus/9.5AS/lib/:$LD_LIBRARY_PATH
```

It is suggested that the `PATH` and the `LD_LIBRARY_PATH` environment variable settings be placed in the BART user account's profile. See the configuration section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for details available at:

<https://www.enterprisedb.com/edb-docs/>

In the following sections, the help option is omitted from the syntax diagrams for the purpose of providing readability for the subcommand options.

3.3.1 CHECK-CONFIG

The CHECK-CONFIG subcommand checks the parameter settings in the BART configuration file as well as the database server configuration for which the `-s` option is specified. To view examples of CHECK-CONFIG subcommand, refer to the *EDB Postgres Backup and Recovery Reference Guide*.

Syntax:

```
bart CHECK-CONFIG [ -s server_name ]
```

Options

`-s` or `--server` <server_name>

`server_name` is the name of the database server to be checked for proper configuration. If the option is omitted, the settings of the global section of the BART configuration file are checked.

Please note:

- When the `-s` option is omitted, the global section [BART], which contains parameters including `bart_host`, `backup_path`, and `pg_basebackup_path` is checked.
- When the `-s` option is specified, the parameters in the specified server section are checked. In addition, certain `postgresql.conf` parameters for the database server must be properly set and the database server must be activated for certain processes. These requirements include the following:
 - The `cluster_owner` parameter must be set to the user account owning the database cluster directory.
 - A passwordless SSH/SCP connection must be set between the BART user and the user account specified by the `cluster_owner` parameter.
 - A database superuser must be specified by the BART `user` parameter.
 - The `pg_hba.conf` file must contain a replication entry for the database superuser specified by the BART `user` parameter.
 - The `archive_mode` parameter in the `postgresql.conf` file must be enabled.
 - The `archive_command` parameter in the `postgresql.auto.conf` or the `postgresql.conf` file must be set.
 - The `allow_incremental_backups` parameter in the BART configuration file must be enabled for database servers for which incremental backups are to be taken.
 - Archiving of WAL files to the `archive_path` must be in process.
 - The WAL scanner program must be running.

The CHECK-CONFIG subcommand displays an error message if the required configuration is not properly set.

3.3.2 INIT

The INIT subcommand is used to perform the following actions:

- Create the BART backup catalog directory.
- Rebuild the BART backupinfo file.
- Set the `archive_command` in the PostgreSQL server based on the `archive_command` setting in the `bart.cfg` file.

Note: If the `archive_mode` configuration parameter is set to `off`, then the `-o` option must be used to set the Postgres `archive_command` using the BART `archive_command` parameter in the BART configuration file even if the `archive_command` is not currently set in `postgresql.conf` nor in `postgresql.auto.conf` file.

To view examples of INIT subcommand, see the *EDB Postgres Backup and Recovery Reference Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

Syntax:

```
bart INIT [ -s { <server_name> | all } ] [ -o ]
  [ -r [ -i { <backup_id> | <backup_name> | all } ] ]
  [--no-configure]
```

Note that subcommand options are generally specified in lowercase.

Options

- `-s` or `--server` {<server_name> | all }
`server_name` is the name of the database server to which the INIT actions are to be applied. If `all` is specified or if the option is omitted, the actions are applied to all servers.
- `-o` or `-override`
 Overrides the existing, active Postgres `archive_command` configuration parameter setting in the `postgresql.conf` file or the `postgresql.auto.conf` file using the BART `archive_command` parameter in the BART configuration file. The INIT generated archive command string is written to the `postgresql.auto.conf` file.
- `-r` or `-rebuild`
 Rebuilds the backupinfo file (a text file named `backupinfo`) located in each backup subdirectory.
- `-i` or `--backupid` { <backup_id> | <backup_name> | all }
`<backup_id>` is an integer, backup identifier and `<backup_name>` is the user-defined alphanumeric name for the backup. If `all` is specified or if the option is omitted, the backupinfo files of all backups for the database servers specified by the `-s` option are recreated. The `-i` option can only be used with the `-r` option.

Archive Command Setting

After the `archive_command` is set, you need to either restart the PostgreSQL server or reload the configuration file in the PostgreSQL server based on the following conditions. See the *EDB Postgres Backup and Recovery Reference Guide* for examples of each of the following.

- If the `archive_mode` is set to `off` and `archive_command` is not set in the PostgreSQL server, the `archive_command` is set based on the `archive_command` setting in the `bart.cfg` and also sets the `archive_mode` to `on`. In this case, you need to restart the PostgreSQL server using `pg_ctl restart`

- If the `archive_mode` is set to `on` and `archive_command` is not set in the PostgreSQL server, the `archive_command` is set based on the `archive_command` setting in the `bart.cfg`. In this case, you need to reload the configuration in the PostgreSQL server using `pg_reload_conf()` or `pg_ctl reload`.
- If the `archive_mode` is set to `off` and `archive_command` is already set in the PostgreSQL server, the `archive_mode` is set to `on`. In this case, you need to restart the PostgreSQL server using `pg_ctl restart`.
- If the `archive_mode` is set to `on` and `archive_command` is already set in the PostgreSQL server, then the `archive_command` is not set unless `-o` option is specified.

3.3.3 BACKUP

The BACKUP subcommand is used to create a full backup or an incremental backup. To view examples of a full backup and an incremental backup, see the *EDB Postgres Backup and Recovery Reference Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

Syntax for full backup:

```
bart BACKUP -s { <server_name> | all } [ -F { p | t } ]
[ -z ] [ -c <compression_level> ]
[ --backup-name <backup_name> ]
[ --thread-count <number_of_threads> ]
[ { --with-pg_basebackup | --no-pg_basebackup } ]
```

Note: While taking a backup, if a file (for example, database server log file) exceeding 1 GB size is stored in the \$PGDATA directory, the backup will fail. To avoid such backup failure, you need to store large files (exceeding 1 GB) outside the \$PGDATA directory.

Syntax for incremental Backup:

```
bart BACKUP -s { <server_name> | all } [ -F p]
[ --parent { <backup_id> | <backup_name> } ]
[ --backup-name <backup_name> ]
[ --thread-count <number_of_threads> ]
[ --check ]
```

Note: To take an incremental backup, you must take a full backup first followed by incremental backup.

For more details about incremental backup, see *Block-Level Incremental Backup*.

Please Note:

- While a BACKUP subcommand is in progress, no other subcommands (INIT, DELETE, MANAGE, SHOW BACKUPS, VERIFY-CHKSUM) must be invoked. Any subcommands invoked while a backup is in progress will skip and ignore the backups.

For full backup, the target default format is a tar file. For incremental backup, only plain format must be specified.

- The backup is saved in the <backup_path>/<server_name>/<backup_id> directory.
 - where:
 - <backup_path> is the value assigned to the <backup_path> parameter in the BART configuration file.
 - <server_name> is the lowercase name of the database server as listed in the configuration file.
 - <backup_id> is a backup identifier assigned by BART to the particular backup.
 - MD5 checksums of the full backup and any user-defined tablespaces are saved as well for tar backups.
 - When you use BART to take a backup of PostgreSQL server versions:
 - 9.5 or prior, only one backup per server may be in progress at any given time and if a backup is interrupted, you must manually run the `pg_stop_backup()` command to terminate the backup mode.

- 9.6 or higher, multiple backups can be taken simultaneously and if a backup is interrupted, the backup mode is terminated automatically without the need to run `pg_stop_backup()` command manually to terminate the backup.

- Before performing the backup, BART checks to ensure if there is enough disk space to completely store the backup in the BART backup catalog.
- In the `postgresql.conf` file, ensure the `wal_keep_segments` configuration parameter is set to a sufficiently large value.

To view more information about the error messages displayed due to insufficient disk space, failure of a database server connection, and if `wal_keep_segments` configuration parameter is not set to a sufficiently large value, see the `BACKUP` section of the *EDB Postgres Backup and Recovery Reference Guide*.

See the error message table in the `BACKUP` section of the *EDB Postgres Backup and Recovery Reference Guide* for the error message that will be displayed.

A low setting of the `wal_keep_segments` configuration parameter may result in the deletion of some WAL files before the BART `BACKUP` subcommand saves them to the `archive_path`. For information about the `wal_keep_segments` parameter, see [PostgreSQL Core Documentation](#).

- If in the BART configuration file, parameter setting `xlog_method=stream` applies to a given database server, streaming of the transaction log in parallel with creation of the backup is performed for that database server, otherwise the transaction log files are collected upon completion of the backup. See the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for details about database server setting.

Note: If the transaction log streaming method is used, the `-F p` option for a plain text backup format must be specified with the `BACKUP` subcommand.

Options

Along with the `BACKUP` subcommand, you can specify the following options:

- `-s { server_name | all }` or `--server { server_name | all }` `server_name` is the database server name to be backed up as specified in the BART configuration file. If `all` is specified, all servers are backed up. This option is mandatory.

Note: If `all` is specified, and a connection to a database server listed in the BART configuration file cannot be opened, the backup for that database server is skipped, but the backup operation continues for the other database servers.

Specify the following options only if required. If you do not specify any of the following options, backup is created with the default setting.

- `-F { p | t }` or `--format { p | t }`

Specifies the backup file format. Use `p` for plain text or `t` for tar. If the option is omitted, the default is tar format.

Note: For taking incremental backups, the option `-F p` must be specified.

- `-z` or `--gzip` (applicable only for full backup)

Specifies usage of `gzip` compression on the tar file output using the default compression level. This option is applicable only for the tar format.

- `-c compression_level` or `--compress-level compression_level` (applicable only for full backup)

Specifies the gzip compression level on the tar file output. `compression_level` is a digit from 1 through 9, with 9 being the best compression. This option is applicable only for the tar format.

- `--parent { backup_id | backup_name }`
 - Specify this option if you want to take an incremental backup. `<backup_id>` is the backup identifier of a parent backup. `<backup_name>` is the user-defined alphanumeric name of a parent backup.
 - The parent is a backup taken prior to the incremental backup. The parent backup can be either a full backup or an incremental backup.
 - The option `-F p` must be specified since an incremental backup can only be taken in plain text format.

Note: An incremental backup cannot be taken on a standby database server. See [Section - Block-Level Incremental Backup](#) for additional information on incremental backups.

- `--backup-name <backup_name>`

User-defined, alphanumeric friendly name to be assigned to the backup. The maximum permitted length of backup name is 49 characters.

- The backup name may include the following variables to be substituted by the timestamp values when the backup is taken: 1) `%year` – 4-digit year, 2) `%month` – 2-digit month, 3) `%day` – 2-digit day, 4) `%hour` – 2-digit hour, 5) `%minute` – 2-digit minute, and 6) `%second` – 2-digit second.

For example,

```
./bart backup -s ppas12 -Ft --backup-name "YEAR = %year
MONTH = %month DAY = %day"
```

- To include the percent sign (`%`) as a character in the backup name, specify `%%` in the alphanumeric string.

For example,

```
./bart backup -s ppas12 -Ft --backup-name "YEAR = %year
MONTH = %month DAY = %day %%"
```

- If the backup name contains space characters (i.e. more than one word) or when referenced with the option `-i` by other subcommands (such as `restore`), enclose the string in single quotes (`'`) or double quotes (`"`).

For example,

```
./bart show-backups -s ppas12 -i "test backup"
```

- If the option `--backup-name` is not specified, and the `backup_name` parameter is not set for this database server in the BART configuration file, then the backup can only be referenced in other BART subcommands by the BART assigned backup identifier.
- `--thread-count number_of_threads`
 - If the option `--thread-count` is specified, `number_of_threads` is the number of worker threads to run in parallel to copy blocks for a backup.

- If the option `--thread-count` is omitted, then the `thread_count` parameter in the BART configuration file applicable to this database server is used.
- If the option `--thread-count` is not enabled for this database server, then the `thread_count` setting in the global section of the BART configuration file is used.
- If the option `--thread-count` is not set in the global section as well, the default number of threads is 1.

Note:

- If parallel backup is run with N number of worker threads, then it will initiate N+ 1 concurrent connections with the server.
 - Thread count will not be effective if backup is taken on a standby server.
-

For more information about the `--thread-count` parameter, see the configuration section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

- `--with-pg_basebackup` (applicable only for full backup)

Specifies that `pg_basebackup` is to be used to take a full backup. The number of thread counts in effect is ignored as given by the `thread_count` parameter in the BART configuration file.

Note: When taking a full backup, if the thread count in effect is greater than 1, then the `pg_basebackup` utility is not used to take the full backup (parallel worker threads are used) unless the option `--with-pg_basebackup` is specified with the `BACKUP` subcommand.

- `--no-pg_basebackup` (applicable only for full backup)

Specifies that `pg_basebackup` is not to be used to take a full backup.

Note: When taking a full backup, if the thread count in effect is only 1, then the `pg_basebackup` utility is used to take the full backup unless the option `--no-pg_basebackup` is specified with the `BACKUP` subcommand.

- `--check` (applicable only for incremental backup)

Verifies that the required MBM files are present in the `archived_wals` directory as specified in the `archive_path` parameter in the `bart.cfg` file before taking an incremental backup. The option `--parent` must be specified when the option `--check` is used. An actual incremental backup is not taken when the option `--check` is specified.

3.3.4 SHOW-SERVERS

The SHOW-SERVERS subcommand displays the information for the managed database servers listed in the BART configuration file. To view examples of SHOW-SERVERS subcommand, see the *EDB Postgres Backup and Recovery Reference Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

Syntax:

```
bart SHOW-SERVERS [ -s { <server_name> | all } ]
```

Options

-s { <server_name> | all } or --server { <server_name> | all }

<server_name> is the name of the database server whose BART configuration information is to be displayed. If all is specified or if the option is omitted, information for all database servers is displayed.

3.3.5 SHOW-BACKUPS

The SHOW-BACKUPS subcommand displays the backup information for the managed database servers. To view examples of SHOW-BACKUPS subcommand, see the *EDB Postgres Backup and Recovery Reference Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

Syntax:

```
bart SHOW-BACKUPS [ -s { <server_name> | all } ]
  [ -i { <backup_id> | <backup_name> | all } ]
  [ -t ]
```

Options

- `-s` or `--server` { <server_name> | all }
 - <server_name> is the name of the database server whose backup information is to be displayed.
 - If `all` is specified or if the option is omitted, the backup information for all database servers is displayed with the exception as described by the following note:

Note: If SHOW-BACKUPS is invoked while the BART BACKUP subcommand is in progress, backups affected by the backup process are shown in progress status in the displayed backup information.

- `-i` or `--backupid` { <backup_id> | <backup_name> | all }
 - <backup_id> is a backup identifier.
 - <backup_name> is the user-defined alphanumeric name for the backup.
 - If `all` is specified or if the option is omitted, all backup information for the relevant database server is displayed.
- `-t` or `-toggle`
 - Displays more backup information in a list format. If the option is omitted, the default is a tabular format.

3.3.6 VERIFY-CHKSUM

The VERIFY-CHKSUM subcommand verifies the MD5 checksums of the full backups and any user-defined tablespaces for the specified database server or for all database servers. The checksum is verified by comparing the current checksum of the backup against the checksum when the backup was taken.

Note: The VERIFY-CHKSUM subcommand is only used for tar format backups. It is not applicable to plain format backups.

Syntax:

```
bart VERIFY-CHKSUM
  [ -s { <server_name> | all } ]
  [ -i { <backup_id> | <backup_name> | all } ]
```

Options

- `-s` or `--server` { <server_name> | all }
 <server_name> is the name of the database server whose tar backup checksums are to be verified. If all is specified or if the `-s` option is omitted, the checksums are verified for all database servers.
- `-i` or `--backupid` { <backup_id> | <backup_name> | all }
 <backup_id> is the backup identifier of a tar format full backup whose checksum is to be verified along with any user-defined tablespaces.
 <backup_name> is the user-defined alphanumeric name for the full backup.
 If all is specified or if the `-i` option is omitted, the checksums of all tar backups for the relevant database server are verified.

Example

The following example verifies the checksum of all tar format backups of the specified database server.

```
-bash-4.1$ bart VERIFY-CHKSUM -s acctg -i all
SERVER NAME    BACKUP ID      VERIFY
acctg          1430239348243  OK
acctg          1430232284202  OK
acctg          1430232016284  OK
acctg          1430231949065  OK
acctg          1429821844271  OK
```

3.3.7 MANAGE

The `MANAGE` subcommand can be invoked to:

- Evaluate backups, mark their status, and delete obsolete backups based on the `retention_policy` parameter in the BART configuration file (See *Section - Managing Backups Using a Retention Policy* for information about retention policy management).
- Compress the archived WAL files based on the `wal_compression` parameter in the BART configuration file (See the configuration section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for information about setting this parameter).

Syntax:

```
bart MANAGE [ -s { <server_name> | all } ]
  [ -l ] [ -d ]
  [ -c { keep | nokeep }
    -i { <backup_id> | <backup_name> | all } ]
  [ -n ]
```

The following summarizes the actions performed under certain conditions and options when the `MANAGE` subcommand is invoked.

- When the `MANAGE` subcommand is invoked with no options or with only the `-s <server_name>` or `-s all` option, the following actions are performed:
 - For the server specified by the `-s` option, or for all servers (if `-s all` is specified or the `-s` option is omitted), active backups are marked as `obsolete` in accordance with the retention policy.
 - All backups that were marked `obsolete` or `keep` prior to invoking the `MANAGE` subcommand remain marked with the same prior status.
 - If WAL compression is enabled for the database server, then any uncompressed, archived WAL files in the BART backup catalog of the database server are compressed with `gzip`.
- When the `MANAGE` subcommand is invoked with any other option besides the `-s` option, the following actions are performed:
 - For the server specified by the `-s` option, or for all servers, the action performed is determined by the other specified options (that is, `-l` to list obsolete backups, `-d` to delete obsolete backups, `-c` to keep or to return backups to `active` status, or `-n` to perform a dry run of any action).
 - No marking of `active` backups to `obsolete` status is performed regardless of the retention policy.
 - All backups that were marked `obsolete` or `keep` prior to invoking the `MANAGE` subcommand remain marked with the same prior status unless the `-c` option (without the `-n` option) is specified to change the backup status of the particular backup or all backups referenced with the `-i` option.
 - No compression is applied to any uncompressed, archived WAL file in the BART backup catalog regardless of whether or not WAL compression is enabled.

The following are additional considerations when using WAL compression:

- Compression of archived WAL files is not permitted for database servers on which incremental backups are to be taken.
- The `gzip` compression program must be installed on the BART host and be accessible in the `PATH` of the BART user account.
- When the `RESTORE` subcommand is invoked, if the `-c` option is specified or if the `copy_wals_during_restore` BART configuration parameter is enabled for the database server, then the following actions occur:

- If compressed, archived WAL files are stored in the BART backup catalog and the location to which the WAL files are to be restored is on a remote host relative to the BART host:
 - * the archived WAL files are transmitted across the network to the remote host in compressed format only if the `gzip` compression program is accessible in the `PATH` of the remote user account that is used to log into the remote host when performing the `RESTORE` operation.
 - * This remote user is specified with either the `remote_host` parameter in the BART configuration file or the `RESTORE -r` option (see *Section - RESTORE*).
 - * Transmission of compressed WAL files results in less network traffic. After the compressed WAL files are transmitted across the network, the `RESTORE` subcommand uncompresses the files for the point-in-time recovery operation.
 - * If the `gzip` program is not accessible on the remote host in the manner described in the previous bullet point, then the compressed, archived WAL files are first uncompressed while on the BART host, then transmitted across the network to the remote host in uncompressed format.
- When the `RESTORE` subcommand is invoked without the `-c` option and the `copy_wals_during_restore` BART configuration parameter is disabled for the database server, then any compressed, archived WAL files needed for the `RESTORE` operation are uncompressed in the BART backup catalog. The uncompressed WAL files can then be saved to the remote host by the `restore_command` in the `postgresql.auto.conf` file when the database server archive recovery begins.

Options

- `-s` or `--server { <server_name> | all }`
`<server_name>` is the name of the database server to which the actions are to be applied. If `all` is specified or if the `-s` option is omitted, the actions are applied to all database servers.
- `-l` or `--list-obsolete`
 List the backups marked as obsolete.
- `-d` or `--delete-obsolete`
 Delete the backups marked as `obsolete`. This action physically deletes the backup along with its archived WAL files and any MBM files for incremental backups.
- `-c` or `--change-status { keep | nokeep }`
 - Specify `keep` to change the status of a backup to `keep` to retain it indefinitely.
 - Specify `nokeep` to change the status of any backup back to active status. The backup can then be re-evaluated and possibly be marked to `obsolete` according to the retention policy by subsequent usage of the `MANAGE` subcommand.

Note: The `-i` option must be included when using the `-c` option.

- `-i` or `--backupid { <backup_id> | <backup_name> | all }`
`<backup_id>` is a backup identifier.
`<backup_name>` is the user-defined alphanumeric name for the backup.
 If `all` is specified, then actions are applied to all backups.

Note: The `-c` option must be included when using the `-i` option.

- `-n, --dry-run`

Performs the test run and displays the results prior to actually implementing the actions as if the operation was performed, however, no changes are actually made.

If `-n` is specified:

- with the `-d` option, it displays which backups would be deleted, but does not actually delete the backups.
- with the `-c` option, it displays the keep or nokeep action, but does not actually change the backup from its current status.
- alone with no other options, or with only the `-s` option, it displays which active backups would be marked as obsolete, but does not actually change the backup status. In addition, no compression is performed on uncompressed, archived WAL files even if WAL compression is enabled for the database server.

3.3.8 RESTORE

The RESTORE subcommand restores the backup and its archived WAL files for the designated database server to the specified directory location. If the appropriate RESTORE options are specified, all recovery settings will be saved in the `postgresql.auto.conf` file. To view examples of how to restore the database server, see the *EDB Postgres Backup and Recovery Reference Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

Syntax:

```
bart RESTORE -s <server_name> -p <restore_path>
[ -i { <backup_id> | <backup_name> } ]
[ -r <remote_user@remote_host_address> ]
[ -w <number_of_workers> ]
[ -t <timeline_id> ]
[ { -x <target_xid> | -g <target_timestamp> } ]
[ -c ]
```

For information about using a continuous archive backup for recovery, refer [PostgreSQL Core Documentation](#). This reference material provides detailed information about the underlying point-in-time recovery process and the meaning and usage of the restore options that are generated into the `postgresql.auto.conf` file by BART.

Please note:

- For special requirements when restoring an incremental backup to a remote database server, see *Restoring an Incremental Backup on a Remote Host*.
- Check to ensure that the host where the backup is to be restored contains enough disk space for the backup and its archived WAL files. The RESTORE subcommand may result in an error while copying files if there is not enough disk space available.
- See *Performing a Restore Operation* to view steps on how to perform a restore operation and see *Point-In-Time Recovery Operation* to view steps on how to perform a point-in-time recovery operation.
- If the backup is restored to a different database cluster directory than where the original database cluster resided, certain operations dependent upon the database cluster location may fail. This happens if their supporting service scripts are not updated to reflect the new directory location of restored backup. For information about the usage and modification of service scripts, see the *EDB Postgres Advanced Server Installation Guide*.

To view the table listing the service scripts for RHEL 6/CentOS 6 and RHEL 7/CentOS 7, see the *EDB Postgres Backup and Recovery Reference Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

Options

- `-s` or `--server <server_name>`
`<server_name>` is the name of the database server to be restored.
- `-p` or `--restore-path <restore_path>`
`<restore_path>` is the directory path where the backup of the database server is to be restored. The directory must be empty and have the proper ownership and privileges assigned to it.
- `-i` or `--backupid { <backup_id> | <backup_name> }`
`<backup_id>` is the backup identifier of the backup to be used for the restoration.
`<backup_name>` is the user-defined alphanumeric name for the backup.
 If the option is omitted, the default is to use the latest backup.

- `-r` or `--remote-host <remote_user@remote_host_address>`

`<remote_user>` is the user account on the remote database server host that accepts a passwordless SSH/SCP login connection and is the owner of the directory where the backup is to be restored.

`<remote_host_address>` is the IP address of the remote host to which the backup is to be restored. This option must be specified if the `<remote_host>` parameter for this database server is not set in the BART configuration file.

If the BART user account is not the same as the operating system account owning the `<restore_path>` directory given with the `-p` option, use the `<remote_host>` BART configuration parameter or the RESTORE subcommand `-r` option to specify the `<restore_path>` directory owner even when restoring to a directory on the same host as the BART host.

See the configuration section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for information about the `<remote_host>` parameter.

- `-w` or `--workers <number_of_workers>`

`<number_of_workers>` is the specification of the number of worker processes to run in parallel to stream the modified blocks of an incremental backup to the restore location.

For example, if 4 worker processes are specified, 4 receiver processes on the restore host and 4 streamer processes on the BART host are used. The output of each streamer process is connected to the input of a receiver process. When the receiver gets to the point where it needs a modified block file, it obtains those modified blocks from its input. With this method, the modified block files are never written to the restore host disk. If the `-w` option is omitted, the default is 1 worker process.

- `-t` or `--target-tli <timeline_id>`

`<timeline_id>` is the integer identifier of the timeline to be used for replaying the archived WAL files for point-in-time recovery.

- `-x` or `--target-xid <target_xid>`

`<target_xid>` is the integer identifier of the transaction ID that determines the transaction up to and including, which point-in-time recovery encompasses. Include either the `-x <target_xid>` or the `--target-xid <target_xid>` option if point-in-time recovery is desired.

- `-g` or `--target-timestamp <target_timestamp>`

`<target_timestamp>` is the timestamp that determines the point in time up to and including, which point-in-time recovery encompasses. Include either the `--target-timestamp <target_timestamp>` or the `-g <target_timestamp>` option if point-in-time recovery is desired.

- `-c` or `--copy-wals`

Specify this option to copy archived WAL files from the BART backup catalog to `<restore_path>/archived_wals` directory.

- If recovery settings are saved in the `postgresql.auto.conf` file for point-in-time recovery, the `restore_command` retrieves the WAL files from `<restore_path>/archived_wals` for the database server archive recovery.
- If the `-c` option is omitted and the `copy_wals_during_restore` parameter in the BART configuration file is not enabled in a manner applicable to this database server, the `restore_command` in the `postgresql.auto.conf` file is generated by default to retrieve the archived WAL files directly from the BART backup catalog. See the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for information about the `copy_wals_during_restore` parameter.

3.3.9 DELETE

The DELETE subcommand removes the subdirectory and data files from the BART backup catalog for the specified backups along with its archived WAL files. To view examples of DELETE subcommand, see the *EDB Postgres Backup and Recovery Reference Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

Syntax:

```
bart DELETE -s <server_name>
  -i { all |
      [']{ <backup_id> | <backup_name> },... }[']
  }
  [ -n ]
```

Note: While invoking the DELETE subcommand, you must specify a specific database server.

For database servers under a retention policy, there are conditions where certain backups may not be deleted. See the *Section - Deletions Permitted Under a Retention Policy* for information regarding permitted backup deletions.

Options

- `-s` or `--server <server_name>`
`<server_name>` is the name of the database server whose backups are to be deleted.
- `-i` or `--backupid { all | [']{ <backup_id> | <backup_name> },... }['] }`
 - `<backup_id>` is the backup identifier of the backup to be deleted.
`<backup_name>` is the user-defined alphanumeric name for the backup.
 - Multiple backup identifiers and backup names may be specified in a comma-separated list. The list must be enclosed within single quotes if there is any white space appearing before or after each comma.
 - If `all` is specified, all of the backups and their archived WAL files for the specified database server are deleted.
- `-n` or `--dry-run`

Displays the results as if the deletions were done, however, no physical removal of the files are actually performed. In other words, a test run is performed so that you can see the potential results prior to actually initiating the action.

After the deletion, the BART backup catalog for the database server no longer contains the corresponding directory for the deleted backup ID. The `archived_wals` subdirectory no longer contains the WAL files of the backup.

3.4 Running the BART WAL Scanner

Use the BART WAL scanner to invoke the `bart-scanner` program located in the `BART_HOME/bin` directory.

Syntax:

```
bart-scanner
[ -d ]
[ -c <config_file_path> ]
{ -h |
  -v |
  --daemon |
  -p mbm_file |
  wal_file |
  RELOAD |
  STOP }
```

Note: For clarity, the syntax diagram shows only the single-character option form (for example, `-d`), but the multi-character option form (for example, `--debug`) is supported as well.

The WAL scanner processes each WAL file to find and record modified blocks in a corresponding modified block map (MBM) file. The default approach is that the WAL scanner gets notified whenever a new WAL file is added to the `archived_wals` directory specified in the `archive_path` parameter of the configuration file. It then scans the WAL file and produces the MBM file.

The default approach does not work in some cases; for example when the WAL files are shipped to the `archive_path` using the `rsync` utility and also in case of some specific platforms. This results in the WAL files being copied to the `archived_wals` directory, but the WAL scanner does not scan them (as WAL scanner is not aware of WAL file) and produce the MBM files. This results in the failure of an incremental backup. This can be avoided by using the timer-based WAL scanning approach, which is done by using the `scan_interval` parameter in the BART configuration file. The value for `scan_interval` is the number of seconds after which the WAL scanner will scan the new WAL files.

When the `bart-scanner` program is invoked, it forks a separate process for each database server enabled with the `allow_incremental_backups` parameter.

The WAL scanner processes can run in either the foreground or background depending upon usage of the `--daemon` option:

- If the `--daemon` option is specified, the WAL scanner process runs in the background. All output messages can be viewed in the BART log file.
- If the `--daemon` option is omitted, the WAL scanner process runs in the foreground. All output messages can be viewed from the terminal running the program as well as in the BART log file.

See the configuration section of the *EDB Postgres Backup and Recovery Installation and Upgrade Guide* for additional information about WAL scanning, `scan_interval`, `allow_incremental_backups`, and logfile parameters.

When invoking the WAL scanner, the current user must be the BART user account.

Note: The BART user account's `LD_LIBRARY_PATH` environment variable may need to be set to include the directory containing the `libpq` library if invocation of the WAL scanner program fails. See [Section - Basic BART Subcommand Usage](#) for information about setting the `LD_LIBRARY_PATH` environment variable.

Options

- `-h` or `--help`
Displays general syntax and information on WAL scanner usage.
- `-v` or `--version`
Displays the WAL scanner version information.
- `-d` or `-debug`
Displays debugging output while executing the WAL scanner with any of its options.
- `-c` or `--config-path config_file_path`
Specifies `config_file_path` as the full directory path to a BART configuration file. Use this option if you do not want to use the default BART configuration file `BART_HOME/etc/bart.cfg`.
- `--daemon`
Runs the WAL scanner as a background process.
- `-p` or `--print mbm_file`
Specifies the full directory path to an MBM file whose content is to be printed. The directory specified in the `archive_path` parameter in the `bart.cfg` file contains the MBM files.
- `wal_file`
Specifies the full directory path to a WAL file to be scanned. The directory specified in the `archive_path` parameter in the `bart.cfg` file contains the WAL files. Use it if a WAL file in the archive path is missing its MBM file.

Note: This option is to be used for assisting the EnterpriseDB support team for debugging problems that may have been encountered.

- `RELOAD`
 - Reloads the BART configuration file. The keyword `RELOAD` is not case-sensitive.
 - The `RELOAD` option is useful if you make changes to the configuration file after the WAL scanner has been started. It will reload the configuration file and adjust the WAL scanners accordingly. For example, if a server section allowing incremental backups is removed from the BART configuration file, then the process attached to that server will stop. Similarly, if a server allowing incremental backups is added, a new WAL scanner process will be launched to scan the WAL files of that server.
- `STOP`
Stops the WAL scanner. The keyword `STOP` is not case-sensitive.

See the *EDB Postgres Backup and Recovery Reference Guide* for usage examples of the BART WAL scanner and its options:

<https://www.enterprisedb.com/edb-docs/>

Using Tablespaces

If the database cluster contains user-defined tablespaces (that is, tablespaces created with the `CREATE TABLESPACE` command):

- You can take full backups with the `BACKUP` subcommand in either tar (`-F t`) or plain text (`-F p`) backup file format.
- You must take incremental backups in the plain text (`-F p`) backup file format.
- You can take full backups using the transaction log streaming method (`xlog_method = stream` in the BART configuration file) `--with-pg_basebackup` and the `BACKUP` subcommand in either tar (`-F t`) or plain text (`-F p`) backup file format.

Note: If the particular database cluster you plan to back up contains tablespaces created by the `CREATE TABLESPACE` command, then you must set the `tablespace_path` parameter in the BART configuration file before you perform a BART `RESTORE` operation.

The `tablespace_path` parameter specifies the directory paths to which you want the tablespaces to be restored. It takes the following format:

```
OID_1=tablespace_path_1;OID_2=tablespace_path_2 ...
```

Where `OID_1`, `OID_2`, ... are the Object Identifiers of the tablespaces. You can find the OIDs of the tablespaces and their corresponding soft links to the directories by listing the contents of the `POSTGRES_INSTALL_HOME/data/pg_tblspc` subdirectory as shown in the following example:

```
[root@localhost pg_tblspc ]# pwd
/opt/PostgresPlus/9.5AS/data/pg_tblspc
[root@localhost pg_tblspc]# ls -l
total 0
lrwxrwxrwx 1 enterprisedb enterprisedb 17 Aug 22 16:38 16644 -> /mnt/tablespace_1
lrwxrwxrwx 1 enterprisedb enterprisedb 17 Aug 22 16:38 16645 -> /mnt/tablespace_2
```

The OIDs are 16644 and 16645 to directories `/mnt/tablespace_1` and `/mnt/tablespace_2`, respectively.

- If you later wish to restore the tablespaces to the same locations as indicated in the preceding example, the BART configuration file must contain the following entry:

```
[ACCTG]
host = 127.0.0.1
port = 5444
user = enterprisedb
cluster_owner = enterprisedb
tablespace_path = 16644=/mnt/tablespace_1;16645=/mnt/tablespace_2
description = "Accounting"
```

- If you later wish to restore the tablespaces to different locations, specify the new directory locations in the `tablespace_path` parameter.

In either case, the directories specified in the `tablespace_path` parameter must exist and be empty at the time you perform the BART `RESTORE` operation.

If the database server is running on a remote host (in other words you are also using the `remote_host` configuration parameter or will specify the `--remote-host` option with the `RESTORE` subcommand), the specified tablespace directories must exist on the specified remote host.

To view example of backing up and restoring a database cluster on a remote host containing tablespaces, see the *EDB Postgres Backup and Recovery Reference Guide* available at:

<https://www.enterprisedb.com/edb-docs/>

The directories must be owned by the user account with which you intend to start the database server (typically the Postgres user account) with no access by other users or groups as is required for the directory path to which the main full backup is to be restored.

To view a sample BART managed backup and recovery system consisting of both local and remote database servers, see the *EDB Postgres Backup and Recovery Reference Guide*. For detailed information about steps to configure and operate the system, see the *EDB Postgres Backup and Recovery Installation and Upgrade Guide*.

EDB Postgres™ Backup and Recovery User Guide

Copyright © 2014 - 2020 EnterpriseDB Corporation.

All rights reserved.

EnterpriseDB® Corporation

34 Crosby Drive, Suite 201, Bedford, MA 01730, USA

T +1 781 357 3390 F +1 978 467 1307 E

info@enterprisedb.com

www.enterprisedb.com

- EDB designs, establishes coding best practices, reviews, and verifies input validation for the logon UI for EDB Postgres product where present. EDB follows the same approach for additional input components, however the nature of the product may require that it accepts freeform SQL, WMI or other strings to be entered and submitted by trusted users for which limited validation is possible. In such cases it is not possible to prevent users from entering incorrect or otherwise dangerous inputs.
- EDB reserves the right to add features to products that accept freeform SQL, WMI or other potentially dangerous inputs from authenticated, trusted users in the future, but will ensure all such features are designed and tested to ensure they provide the minimum possible risk, and where possible, require superuser or equivalent privileges.
- EDB does not warrant that we can or will anticipate all potential threats and therefore our process cannot fully guarantee that all potential vulnerabilities have been addressed or considered.

B

BACKUP Subcommand, 32
BART Management Overview, 14
Basic BART Subcommand Usage, 27
Block-Level Incremental Backup, 7

C

CHECK-CONFIG Subcommand, 29
Concept Overview, 8
Conclusion, 49
Conventions Used in this Guide, 3
Creating a Backup Chain, 13

D

DELETE Subcommand, 44
Deleting obsolete backups, 25
Deletions permitted under retention policy, 23

E

Evaluating obsolete backups, 25

I

Incremental Backup Limitations and Requirements, 7
INIT Subcommand, 30
Introduction, 1

M

MANAGE Subcommand, 39
Managing Backups Using a Retention Policy, 20
Managing Incremental Backups, 25
Managing the Backups Based on the Retention Policy, 23
Marking backups for indefinite keep status, 25
Marking obsolete backups, 25
Marking the Backup Status, 21

O

Overview, 4
Overview - Managing Backups Using a Retention Policy,
20

P

Performing a Restore Operation, 16
Performing an Incremental Backup, 10
Point-In-Time Recovery Operation, 18

R

Recovery window retention policy, 23
Recovery window retention with incremental backups, 26
Redundancy retention policy, 22
Redundancy retention with incremental backups, 26
RESTORE Subcommand, 42
Restoring an Incremental Backup, 11
Restoring incremental backup on bart host, 11
Restoring incremental backup on remote host, 12
Restrictions on pg_basebackup, 3
Running the BART WAL Scanner, 45

S

Setting the Retention Policy, 22
SHOW-BACKUPS Subcommand, 37
SHOW-SERVERS Subcommand, 36

U

Using BART, 14
Using Tablespaces, 47

V

VERIFY-CHKSUM Subcommand, 38

W

WAL Scanning – Preparation for an Incremental Backup,
9
What's New, 2