### postgresql.org

# **PostgreSQL: Documentation: 9.2: pg\_dump**

### dbname

Specifies the name of the database to be dumped. If this is not specified, the environment variable PGDATABASE is used. If that is not set, the user name specified for the connection is used.

-a --data-only

Dump only the data, not the schema (data definitions). Table data, large objects, and sequence values are dumped.

This option is similar to, but for historical reasons not identical to, specifying --section=data.

-b --blobs

Include large objects in the dump. This is the default behavior except when --schema, --table, or --schema-only is specified. The -b switch is therefore only useful to add large objects to dumps where a specific schema or table has been requested. Note that blobs are considered data and therefore will be included when --data-only is used, but not when --schema-only is.

```
-c
--clean
```

Output commands to clean (drop) database objects prior to outputting the commands for creating them. (Restore might generate some harmless error messages, if any objects were not present in the destination database.)

This option is only meaningful for the plain-text format. For the archive formats, you can specify the option when you call pg\_restore.

```
-C
--create
```

Begin the output with a command to create the database itself and reconnect to the created database. (With a script of this form, it doesn't matter which database in the destination installation you connect to before running the script.) If --clean is also specified, the script drops and recreates the target database before reconnecting to it.

This option is only meaningful for the plain-text format. For the archive formats, you can specify the option when you call pg\_restore.

-E encoding --encoding=encoding

Create the dump in the specified character set encoding. By default, the dump is created in the database encoding. (Another way to get the same result is to set the PGCLIENTENCODING environment variable to the desired dump encoding.)

-f file --file=file

Send output to the specified file. This parameter can be omitted for file based output formats, in which case the standard output is used. It must be given for the directory output format however, where it specifies the target directory instead of a file. In this case the directory is created by pg\_dump and must not exist before.

-F format --format=format

Selects the format of the output. format can be one of the following:

## р

plain

Output a plain-text SQL script file (the default).

### С

custom

Output a custom-format archive suitable for input into pg\_restore. Together with the directory output format, this is the most flexible output format in that it allows manual selection and reordering of archived items during restore. This format is also compressed by default.

### d

directory

Output a directory-format archive suitable for input into pg\_restore. This will create a directory with one file for each table and blob being dumped, plus a so-called Table of Contents file describing the dumped objects in a machine-readable format that pg\_restore can read. A directory format archive can be manipulated with standard Unix tools; for example, files in an uncompressed archive can be compressed with the gzip tool. This format is compressed by default.

#### t tar

Output a tar-format archive suitable for input into pg\_restore. The tar format is compatible with the directory format: extracting a tar-format archive produces a valid directory-format archive. However, the tar format does not support compression. Also, when using tar format the relative order of table data items cannot be changed during restore.

-i --ignore-version

A deprecated option that is now ignored.

-n schema --schema=schema

Dump only schemas matching SChema; this selects both the schema itself, and all its contained objects. When this option is not specified, all non-system schemas in the target database will be dumped. Multiple schemas can be selected by writing multiple - n switches. Also, the SChema parameter is interpreted as a pattern according to the same rules used by psql's \d commands (see <u>Patterns</u>), so multiple schemas can also be selected by writing wildcard characters in the pattern. When using wildcards, be careful to quote the pattern if needed to prevent the shell from expanding the wildcards; see <u>Examples</u>.

**Note:** When -n is specified, pg\_dump makes no attempt to dump any other database objects that the selected schema(s) might depend upon. Therefore, there is no guarantee that the results of a specific-schema dump can be successfully restored by themselves into a clean database.

**Note:** Non-schema objects such as blobs are not dumped when - n is specified. You can add blobs back to the dump with the - -blobs switch.

-N schema

--exclude-schema=schema

Do not dump any schemas matching the SChema pattern. The pattern is interpreted according to the same rules as for -n. -N can be given more than once to exclude schemas matching any of several patterns.

When both -n and -N are given, the behavior is to dump just the schemas that match at least one -n switch but no -N switches. If -N appears without -n, then schemas matching -N are excluded from what is otherwise a normal dump.

- 0

--oids

Dump object identifiers (OIDs) as part of the data for every table. Use this option if your application references the OID columns in some way (e.g., in a foreign key constraint). Otherwise, this option should not be used.

-0

--no-owner

Do not output commands to set ownership of objects to match the original database. By default, pg\_dump issues ALTER OWNER or SET SESSION AUTHORIZATION statements to set ownership of created database objects. These statements will fail when the script is run unless it is started by a superuser (or the same user that owns all of the objects in the script). To make a script that can be restored by any user, but will give that user ownership of all the objects, specify -**O**.

This option is only meaningful for the plain-text format. For the archive formats, you can specify the option when you call pg\_restore.

-R --no-reconnect

This option is obsolete but still accepted for backwards compatibility.

-s --schema-only

Dump only the object definitions (schema), not data.

This option is the inverse of --data-only. It is similar to, but for historical reasons not identical to, specifying --section=pre-data --section=post-data.

(Do not confuse this with the --Schema option, which uses the word "schema" in a different meaning.)

To exclude table data for only a subset of tables in the database, see --exclude-table-data.

-S username --superuser=username

Specify the superuser user name to use when disabling triggers. This is only relevant if --disable-triggers is used. (Usually, it's better to leave this out, and instead start the resulting script as superuser.)

```
-t table
--table=table
```

Dump only tables (or views or sequences or foreign tables) matching table. Multiple tables can be selected by writing multiple -t switches. Also, the table parameter is interpreted as a pattern according to the same rules used by psql's \d commands (see <u>Patterns</u>), so multiple tables can also be

selected by writing wildcard characters in the pattern. When using wildcards, be careful to quote the pattern if needed to prevent the shell from expanding the wildcards; see <u>Examples</u>.

The -n and -N switches have no effect when -t is used, because tables selected by -t will be dumped regardless of those switches, and non-table objects will not be dumped.

**Note:** When -t is specified, pg\_dump makes no attempt to dump any other database objects that the selected table(s) might depend upon. Therefore, there is no guarantee that the results of a specific-table dump can be successfully restored by themselves into a clean database.

**Note:** The behavior of the -t switch is not entirely upward compatible with pre-8.2 PostgreSQL versions. Formerly, writing -t tab would dump all tables named tab, but now it just dumps whichever one is visible in your default search path. To get the old behavior you can write -t '\*.tab'. Also, you must write something like -t sch.tab to select a table in a particular schema, rather than the old locution of -n sch -t tab.

```
-T table
--exclude-table=table
```

Do not dump any tables matching the table pattern. The pattern is interpreted according to the same rules as for -t. -T can be given more than once to exclude tables matching any of several patterns.

When both -t and -T are given, the behavior is to dump just the tables that match at least one -t switch but no -T switches. If -T appears without -t, then tables matching -T are excluded from what is otherwise a normal dump.

-v --verbose

Specifies verbose mode. This will cause pg\_dump to output detailed object comments and start/stop times to the dump file, and progress messages to standard error.

-V --version

Print the pg\_dump version and exit.

```
-x
--no-privileges
--no-acl
```

Prevent dumping of access privileges (grant/revoke commands).

-Z 0..9 --compress=0..9 Specify the compression level to use. Zero means no compression. For the custom archive format, this specifies compression of individual table-data segments, and the default is to compress at a moderate level. For plain text output, setting a nonzero compression level causes the entire output file to be compressed, as though it had been fed through gzip; but the default is not to compress. The tar archive format currently does not support compression at all.

### --binary-upgrade

This option is for use by in-place upgrade utilities. Its use for other purposes is not recommended or supported. The behavior of the option may change in future releases without notice.

--column-inserts --attribute-inserts

Dump data as INSERT commands with explicit column names (INSERT INTO table (column, ...) VALUES ...). This will make restoration very slow; it is mainly useful for making dumps that can be loaded into non-PostgreSQL databases. However, since this option generates a separate command for each row, an error in reloading a row causes only that row to be lost rather than the entire table contents.

--disable-dollar-quoting

This option disables the use of dollar quoting for function bodies, and forces them to be quoted using SQL standard string syntax.

### --disable-triggers

This option is only relevant when creating a data-only dump. It instructs pg\_dump to include commands to temporarily disable triggers on the target tables while the data is reloaded. Use this if you have referential integrity checks or other triggers on the tables that you do not want to invoke during data reload.

Presently, the commands emitted for --disable-triggers must be done as superuser. So, you should also specify a superuser name with -S, or preferably be careful to start the resulting script as a superuser.

This option is only meaningful for the plain-text format. For the archive formats, you can specify the option when you call pg\_restore.

--exclude-table-data=table

Do not dump data for any tables matching the table pattern. The pattern is interpreted according to the same rules as for -t. --exclude-table-data can be given more than once to exclude tables matching any of several patterns. This option is useful when you need the definition of a particular table even though you do not need the data in it.

To exclude data for all tables in the database, see --schema-only.

--inserts

Dump data as INSERT commands (rather than COPY). This will make restoration very slow; it is mainly useful for making dumps that can be loaded into non-PostgreSQL databases. However, since this option generates a separate command for each row, an error in reloading a row causes only that row to be lost rather than the entire table contents. Note that the restore might fail altogether if you have rearranged column order. The --column-inserts option is safe against column order changes, though even slower.

--lock-wait-timeout=timeout

Do not wait forever to acquire shared table locks at the beginning of the dump. Instead fail if unable to lock a table within the specified timeout. The timeout may be specified in any of the formats accepted by SET statement\_timeout. (Allowed values vary depending on the server version you are dumping from, but an integer number of milliseconds is accepted by all versions since 7.3. This option is ignored when dumping from a pre-7.3 server.)

--no-security-labels

Do not dump security labels.

--no-tablespaces

Do not output commands to select tablespaces. With this option, all objects will be created in whichever tablespace is the default during restore.

This option is only meaningful for the plain-text format. For the archive formats, you can specify the option when you call pg\_restore.

--no-unlogged-table-data

Do not dump the contents of unlogged tables. This option has no effect on whether or not the table definitions (schema) are dumped; it only suppresses dumping the table data. Data in unlogged tables is always excluded when dumping from a standby server.

--quote-all-identifiers

Force quoting of all identifiers. This option is recommended when dumping a database from a server whose PostgreSQL major version is different from pg\_dump's, or when the output is intended to be loaded into a server of a different major version. By default, pg\_dump quotes only identifiers that are reserved words in its own major version. This sometimes results in compatibility issues when dealing with servers of other versions that may have slightly different sets of reserved words. Using --quote-all-identifiers prevents such issues, at the price of a harder-to-read dump script.

### --section=sectionname

Only dump the named section. The section name can be pre-data, data, or post-data. This option can be specified more than once to select multiple sections. The default is to dump all sections.

The data section contains actual table data, large-object contents, and sequence values. Post-data items include definitions of indexes, triggers, rules, and constraints other than validated check constraints. Pre-data items include all other data definition items.

--serializable-deferrable

Use a serializable transaction for the dump, to ensure that the snapshot used is consistent with later database states; but do this by waiting for a point in the transaction stream at which no anomalies can be present, so that there isn't a risk of the dump failing or causing other transactions to roll back with a serialization\_failure. See <u>Chapter 13</u> for more information about transaction isolation and concurrency control.

This option is not beneficial for a dump which is intended only for disaster recovery. It could be useful for a dump used to load a copy of the database for reporting or other read-only load sharing while the original database continues to be updated. Without it the dump may reflect a state which is not consistent with any serial execution of the transactions eventually committed. For example, if batch processing techniques are used, a batch may show as closed in the dump without all of the items which are in the batch appearing.

This option will make no difference if there are no read-write transactions active when pg\_dump is started. If read-write transactions are active, the start of the dump may be delayed for an indeterminate length of time. Once running, performance with or without the switch is the same.

--use-set-session-authorization

Output SQL-standard SET SESSION AUTHORIZATION commands instead of ALTER OWNER commands to determine object ownership. This makes the dump more standards-compatible, but depending on the history of the objects in the dump, might not restore properly. Also, a dump using SET SESSION AUTHORIZATION will certainly require superuser privileges to restore correctly, whereas ALTER OWNER requires lesser privileges.

```
-?
--he
```

--help

Show help about pg\_dump command line arguments, and exit.