

# Package: RPostgres (via r-universe)

August 31, 2024

**Title** C++ Interface to PostgreSQL

**Version** 1.4.7.9010

**Date** 2024-08-31

**Description** Fully DBI-compliant C++-backed interface to PostgreSQL  
<<https://www.postgresql.org/>>, an open-source relational  
database.

**License** MIT + file LICENSE

**URL** <https://rpostgres.r-dbi.org>, <https://github.com/r-dbi/RPostgres>

**BugReports** <https://github.com/r-dbi/RPostgres/issues>

**Depends** R (>= 3.1.0)

**Imports** bit64, blob (>= 1.2.0), DBI (>= 1.2.0), hms (>= 1.0.0),  
lubridate, methods, withr

**Suggests** callr, covr, DBItest (>= 1.7.3), knitr, rlang, rmarkdown,  
testthat (>= 3.0.0)

**LinkingTo** cpp11, plogr (>= 0.2.0)

**Config/Needs/website** r-dbi/dbitemplate

**VignetteBuilder** knitr

**Config/Needs/build** decor

**Config/autostyle/scope** line\_breaks

**Config/autostyle/strict** false

**Config/testthat/edition** 3

**Encoding** UTF-8

**LazyLoad** true

**Roxygen** list(markdown = TRUE)

**RoxygenNote** 7.3.2.9000

**SystemRequirements** libpq >= 9.0: libpq-dev (deb) or postgresql-devel  
(rpm)

**Collate** 'PqDriver.R' 'PqConnection.R' 'PqResult.R' 'RPostgres-pkg.R'  
 'Redshift.R' 'cpp11.R' 'dbAppendTable\_PqConnection.R'  
 'dbBegin\_PqConnection.R' 'dbBind\_PqResult.R'  
 'dbClearResult\_PqResult.R' 'dbColumnInfo\_PqResult.R'  
 'dbCommit\_PqConnection.R' 'dbConnect\_PqDriver.R'  
 'dbConnect\_RedshiftDriver.R' 'dbDataType\_PqConnection.R'  
 'dbDataType\_PqDriver.R' 'dbDisconnect\_PqConnection.R'  
 'dbExistsTable\_PqConnection\_Id.R'  
 'dbExistsTable\_PqConnection\_character.R' 'dbFetch\_PqResult.R'  
 'dbGetInfo\_PqConnection.R' 'dbGetInfo\_PqDriver.R'  
 'dbGetRowCount\_PqResult.R' 'dbGetRowsAffected\_PqResult.R'  
 'dbGetStatement\_PqResult.R' 'dbHasCompleted\_PqResult.R'  
 'dbIsValid\_PqConnection.R' 'dbIsValid\_PqDriver.R'  
 'dbIsValid\_PqResult.R' 'dbListFields\_PqConnection\_Id.R'  
 'dbListFields\_PqConnection\_character.R'  
 'dbListObjects\_PqConnection\_ANY.R'  
 'dbListTables\_PqConnection.R'  
 'dbQuoteIdentifier\_PqConnection\_Id.R'  
 'dbQuoteIdentifier\_PqConnection\_SQL.R'  
 'dbQuoteIdentifier\_PqConnection\_character.R'  
 'dbQuoteLiteral\_PqConnection.R'  
 'dbQuoteString\_PqConnection\_SQL.R'  
 'dbQuoteString\_PqConnection\_character.R'  
 'dbReadTable\_PqConnection\_character.R'  
 'dbRemoveTable\_PqConnection\_character.R'  
 'dbRollback\_PqConnection.R' 'dbSendQuery\_PqConnection.R'  
 'dbUnloadDriver\_PqDriver.R'  
 'dbUnquoteIdentifier\_PqConnection\_SQL.R'  
 'dbWriteTable\_PqConnection\_character\_data.frame.R' 'default.R'  
 'export.R' 'quote.R' 'show\_PqConnection.R'  
 'sqlData\_PqConnection.R' 'tables.R' 'transactions.R' 'utils.R'

**Repository** <https://r-dbi.r-universe.dev>

**RemoteUrl** <https://github.com/r-dbi/rpostgres>

**RemoteRef** HEAD

**RemoteSha** db9b80bde242a2b63d17af39a056148788fcc9a4

## Contents

RPostgres-package	3
Postgres	4
postgres-query	5
postgres-tables	7
postgres-transactions	9
postgresHasDefault	10
postgresIsTransacting	11
postgresWaitForNotify	11
quote	12

<i>RPostgres-package</i>	3
Redshift . . . . .	13
<b>Index</b>	<b>15</b>

---

*RPostgres-package*      *RPostgres: C++ Interface to PostgreSQL*

---

## Description

Fully DBI-compliant C++-backed interface to PostgreSQL <https://www.postgresql.org/>, an open-source relational database.

## Author(s)

**Maintainer:** Kirill Müller <[kirill@cynkra.com](mailto:kirill@cynkra.com)> ([ORCID](#))

Authors:

- Hadley Wickham
- Jeroen Ooms

Other contributors:

- RStudio [copyright holder]
- R Consortium [funder]
- Tomoaki Nishiyama (Code for encoding vectors into strings derived from RPostgreSQL) [contributor]

## See Also

Useful links:

- <https://rpostgres.r-dbi.org>
- <https://github.com/r-dbi/RPostgres>
- Report bugs at <https://github.com/r-dbi/RPostgres/issues>

Postgres

*Postgres driver***Description**

`DBI::dbConnect()` establishes a connection to a database. Set `drv = Postgres()` to connect to a PostgreSQL(-ish) database. Use `drv = Redshift()` instead to connect to an AWS Redshift cluster.

Manually disconnecting a connection is not necessary with **RPostgres**, but still recommended; if you delete the object containing the connection, it will be automatically disconnected during the next GC with a warning.

**Usage**

```
Postgres()
```

```
## S4 method for signature 'PqDriver'
```

```
dbConnect(
  drv,
  dbname = NULL,
  host = NULL,
  port = NULL,
  password = NULL,
  user = NULL,
  service = NULL,
  ...,
  bigint = c("integer64", "integer", "numeric", "character"),
  check_interrupts = FALSE,
  timezone = "UTC",
  timezone_out = NULL
)
```

```
## S4 method for signature 'PqConnection'
```

```
dbDisconnect(conn, ...)
```

**Arguments**

<code>drv</code>	<a href="#">DBI::DBIDriver</a> . Use <a href="#">Postgres()</a> to connect to a PostgreSQL(-ish) database or <a href="#">Redshift()</a> to connect to an AWS Redshift cluster. Use an existing <a href="#">DBI::DBIConnection</a> object to clone an existing connection.
<code>dbname</code>	Database name. If NULL, defaults to the user name. Note that this argument can only contain the database name, it will not be parsed as a connection string (internally, <code>expand_dbname</code> is set to <code>false</code> in the call to <a href="#">PQconnectdbParams()</a> ).
<code>host, port</code>	Host and port. If NULL, will be retrieved from <code>PGHOST</code> and <code>PGPORT</code> env vars.
<code>user, password</code>	User name and password. If NULL, will be retrieved from <code>PGUSER</code> and <code>PGPASSWORD</code> envvars, or from the appropriate line in <code>~/ .pgpass</code> . See <a href="https://www.postgresql.org/docs/current/libpq-pgpass.html">https://www.postgresql.org/docs/current/libpq-pgpass.html</a> for more details.

service	Name of service to connect as. If NULL, will be ignored. Otherwise, connection parameters will be loaded from the pg_service.conf file and used. See <a href="https://www.postgresql.org/docs/current/libpq-pgservice.html">https://www.postgresql.org/docs/current/libpq-pgservice.html</a> for details on this file and syntax.
...	Other name-value pairs that describe additional connection options as described at <a href="https://www.postgresql.org/docs/current/libpq-connect.html#LIBPQ-PARAMKEYWORDS">https://www.postgresql.org/docs/current/libpq-connect.html#LIBPQ-PARAMKEYWORDS</a>
bigint	The R type that 64-bit integer types should be mapped to, default is <code>bit64::integer64</code> , which allows the full range of 64 bit integers.
check_interrupts	Should user interrupts be checked during the query execution (before first row of data is available)? Setting to TRUE allows interruption of queries running too long.
timezone	Sets the timezone for the connection. The default is "UTC". If NULL then no timezone is set, which defaults to the server's time zone.
timezone_out	The time zone returned to R, defaults to <code>timezone</code> . If you want to display date-time values in the local timezone, set to <code>Sys.timezone()</code> or <code>""</code> . This setting does not change the time values returned, only their display.
conn	Connection to disconnect.

### Examples

```
library(DBI)
# Pass more arguments as necessary to dbConnect()
con <- dbConnect(RPostgres::Postgres())
dbDisconnect(con)
```

---

postgres-query

*Execute a SQL statement on a database connection*

---

### Description

To retrieve results a chunk at a time, use `dbSendQuery()`, `dbFetch()`, then `dbClearResult()`. Alternatively, if you want all the results (and they'll fit in memory) use `dbGetQuery()` which sends, fetches and clears for you.

### Usage

```
## S4 method for signature 'PqResult'
dbBind(res, params, ...)

## S4 method for signature 'PqResult'
dbClearResult(res, ...)

## S4 method for signature 'PqResult'
dbFetch(res, n = -1, ..., row.names = FALSE)
```

```
## S4 method for signature 'PqResult'
dbHasCompleted(res, ...)

## S4 method for signature 'PqConnection'
dbSendQuery(conn, statement, params = NULL, ..., immediate = FALSE)
```

### Arguments

<code>res</code>	Code a <a href="#">PqResult</a> produced by <code>DBI::dbSendQuery()</code> .
<code>params</code>	A list of query parameters to be substituted into a parameterised query. Query parameters are sent as strings, and the correct type is imputed by PostgreSQL. If this fails, you can manually cast the parameter with e.g. <code>"\$1::bigint"</code> .
<code>...</code>	Other arguments needed for compatibility with generic (currently ignored).
<code>n</code>	Number of rows to return. If less than zero returns all rows.
<code>row.names</code>	Either TRUE, FALSE, NA or a string. If TRUE, always translate row names to a column called "row_names". If FALSE, never translate row names. If NA, translate rownames only if they're a character vector. A string is equivalent to TRUE, but allows you to override the default name. For backward compatibility, NULL is equivalent to FALSE.
<code>conn</code>	A <a href="#">PqConnection</a> created by <code>DBI::dbConnect()</code> .
<code>statement</code>	An SQL string to execute.
<code>immediate</code>	If TRUE, uses the <code>PGsendQuery()</code> API instead of <code>PGprepare()</code> . This allows to pass multiple statements and turns off the ability to pass parameters.

### Multiple queries and statements

With `immediate = TRUE`, it is possible to pass multiple queries or statements, separated by semi-colons. For multiple statements, the resulting value of `DBI::dbGetRowsAffected()` corresponds to the total number of affected rows. If multiple queries are used, all queries must return data with the same column names and types. Queries and statements can be mixed.

### Examples

```
library(DBI)
db <- dbConnect(RPostgres::Postgres())
dbWriteTable(db, "usarrests", datasets::USArrests, temporary = TRUE)

# Run query to get results as dataframe
dbGetQuery(db, "SELECT * FROM usarrests LIMIT 3")

# Send query to pull requests in batches
res <- dbSendQuery(db, "SELECT * FROM usarrests")
dbFetch(res, n = 2)
dbFetch(res, n = 2)
dbHasCompleted(res)
dbClearResult(res)
```

```
dbRemoveTable(db, "usarrests")  
  
dbDisconnect(db)
```

---

postgres-tables

*Convenience functions for reading/writing DBMS tables*

---

## Description

`DBI::dbAppendTable()` is overridden because **RPostgres** uses placeholders of the form \$1, \$2 etc. instead of ?.

`DBI::dbWriteTable()` executes several SQL statements that create/overwrite a table and fill it with values. **RPostgres** does not use parameterised queries to insert rows because benchmarks revealed that this was considerably slower than using a single SQL string.

## Usage

```
## S4 method for signature 'PqConnection'  
dbAppendTable(conn, name, value, copy = NULL, ..., row.names = NULL)  
  
## S4 method for signature 'PqConnection,Id'  
dbExistsTable(conn, name, ...)  
  
## S4 method for signature 'PqConnection,character'  
dbExistsTable(conn, name, ...)  
  
## S4 method for signature 'PqConnection,Id'  
dbListFields(conn, name, ...)  
  
## S4 method for signature 'PqConnection,character'  
dbListFields(conn, name, ...)  
  
## S4 method for signature 'PqConnection'  
dbListObjects(conn, prefix = NULL, ...)  
  
## S4 method for signature 'PqConnection'  
dbListTables(conn, ...)  
  
## S4 method for signature 'PqConnection,character'  
dbReadTable(conn, name, ..., check.names = TRUE, row.names = FALSE)  
  
## S4 method for signature 'PqConnection,character'  
dbRemoveTable(conn, name, ..., temporary = FALSE, fail_if_missing = TRUE)  
  
## S4 method for signature 'PqConnection,character,data.frame'
```

```

dbWriteTable(
  conn,
  name,
  value,
  ...,
  row.names = FALSE,
  overwrite = FALSE,
  append = FALSE,
  field.types = NULL,
  temporary = FALSE,
  copy = NULL
)

## S4 method for signature 'PqConnection'
sqlData(con, value, row.names = FALSE, ...)

```

### Arguments

conn	a <a href="#">PqConnection</a> object, produced by <a href="#">DBI::dbConnect()</a>
name	a character string specifying a table name. Names will be automatically quoted so you can use any sequence of characters, not just any valid bare table name. Alternatively, pass a name quoted with <a href="#">DBI::dbQuoteIdentifier()</a> , an <a href="#">Id()</a> object, or a string escaped with <a href="#">DBI::SQL()</a> .
value	A data.frame to write to the database.
copy	If TRUE, serializes the data frame to a single string and uses COPY name FROM stdin. This is fast, but not supported by all postgres servers (e.g. Amazon's Redshift). If FALSE, generates a single SQL string. This is slower, but always supported. The default maps to TRUE on connections established via <a href="#">Postgres()</a> and FALSE on connections established via <a href="#">Redshift()</a> .
...	Ignored.
row.names	Either TRUE, FALSE, NA or a string. If TRUE, always translate row names to a column called "row_names". If FALSE, never translate row names. If NA, translate rownames only if they're a character vector. A string is equivalent to TRUE, but allows you to override the default name. For backward compatibility, NULL is equivalent to FALSE.
prefix	A fully qualified path in the database's namespace, or NULL. This argument will be processed with <a href="#">dbUnquoteIdentifier()</a> . If given the method will return all objects accessible through this prefix.
check.names	If TRUE, the default, column names will be converted to valid R identifiers.
temporary	If TRUE, only temporary tables are considered.
fail_if_missing	If FALSE, <a href="#">dbRemoveTable()</a> succeeds if the table doesn't exist.
overwrite	a logical specifying whether to overwrite an existing table or not. Its default is FALSE.

append	a logical specifying whether to append to an existing table in the DBMS. Its default is FALSE.
field.types	character vector of named SQL field types where the names are the names of new table's columns. If missing, types are inferred with <code>DBI::dbDataType()</code> . The types can only be specified with <code>append = FALSE</code> .
con	A database connection.

### Schemas, catalogs, tablespaces

Pass an identifier created with `Id()` as the name argument to specify the schema or catalog, e.g. `name = Id(catalog = "my_catalog", schema = "my_schema", table = "my_table")`. To specify the tablespace, use `dbExecute(conn, "SET default_tablespace TO my_tablespace")` before creating the table.

### Examples

```
library(DBI)
con <- dbConnect(RPostgres::Postgres())
dbListTables(con)
dbWriteTable(con, "mtcars", mtcars, temporary = TRUE)
dbReadTable(con, "mtcars")

dbListTables(con)
dbExistsTable(con, "mtcars")

# A zero row data frame just creates a table definition.
dbWriteTable(con, "mtcars2", mtcars[0, ], temporary = TRUE)
dbReadTable(con, "mtcars2")

dbDisconnect(con)
```

---

postgres-transactions *Transaction management.*

---

### Description

`dbBegin()` starts a transaction. `dbCommit()` and `dbRollback()` end the transaction by either committing or rolling back the changes.

### Usage

```
## S4 method for signature 'PqConnection'
dbBegin(conn, ..., name = NULL)

## S4 method for signature 'PqConnection'
dbCommit(conn, ..., name = NULL)

## S4 method for signature 'PqConnection'
dbRollback(conn, ..., name = NULL)
```

**Arguments**

conn	a <a href="#">PqConnection</a> object, produced by <code>DBI::dbConnect()</code>
...	Unused, for extensibility.
name	If provided, uses the SAVEPOINT SQL syntax to establish, remove (commit) or undo a <code>savepoint</code> .

**Value**

A boolean, indicating success or failure.

**Examples**

```
library(DBI)
con <- dbConnect(RPostgres::Postgres())
dbWriteTable(con, "USarrests", datasets::USArrests, temporary = TRUE)
dbGetQuery(con, 'SELECT count(*) from "USarrests"')

dbBegin(con)
dbExecute(con, 'DELETE from "USarrests" WHERE "Murder" > 1')
dbGetQuery(con, 'SELECT count(*) from "USarrests"')
dbRollback(con)

# Rolling back changes leads to original count
dbGetQuery(con, 'SELECT count(*) from "USarrests"')

dbRemoveTable(con, "USarrests")
dbDisconnect(con)
```

---

`postgresHasDefault`      *Check if default database is available.*

---

**Description**

`RPostgres` examples and tests connect to a default database via `dbConnect(Postgres())`. This function checks if that database is available, and if not, displays an informative message.

`postgresDefault()` works similarly but returns a connection on success and throws a test that skip condition on failure, making it suitable for use in tests.

**Usage**

```
postgresHasDefault(...)
```

```
postgresDefault(...)
```

**Arguments**

...      Additional arguments passed on to `DBI::dbConnect()`

**Examples**

```
if (postgresHasDefault()) {
  db <- postgresDefault()
  print(dbListTables(db))
  dbDisconnect(db)
} else {
  message("No database connection.")
}
```

---

postgresIsTransacting *Return whether a transaction is ongoing*

---

**Description**

Detect whether the transaction is active for the given connection. A transaction might be started with `DBI::dbBegin()` or wrapped within `DBI::dbWithTransaction()`.

**Usage**

```
postgresIsTransacting(conn)
```

**Arguments**

conn            a `PqConnection` object, produced by `DBI::dbConnect()`

**Value**

A boolean, indicating if a transaction is ongoing.

---

postgresWaitForNotify *Wait for and return any notifications that return within timeout*

---

**Description**

Once you subscribe to notifications with `LISTEN`, use this to wait for responses on each channel.

**Usage**

```
postgresWaitForNotify(conn, timeout = 1)
```

**Arguments**

conn            a `PqConnection` object, produced by `DBI::dbConnect()`  
timeout        How long to wait, in seconds. Default 1

**Value**

If a notification was available, a list of:

**channel** Name of channel

**pid** PID of notifying server process

**payload** Content of notification

If no notifications are available, return NULL

**Examples**

```
library(DBI)
library(callr)

# listen for messages on the grapevine
db_listen <- dbConnect(RPostgres::Postgres())
dbExecute(db_listen, "LISTEN grapevine")

# Start another process, which sends a message after a delay
rp <- r_bg(function() {
  library(DBI)
  Sys.sleep(0.3)
  db_notify <- dbConnect(RPostgres::Postgres())
  dbExecute(db_notify, "NOTIFY grapevine, 'psst'")
  dbDisconnect(db_notify)
})

# Sleep until we get the message
n <- NULL
while (is.null(n)) {
  n <- RPostgres::postgresWaitForNotify(db_listen, 60)
}
stopifnot(n$payload == 'psst')

# Tidy up
rp$wait()
dbDisconnect(db_listen)
```

---

quote

*Quote postgres strings, identifiers, and literals*

---

**Description**

If an object of class `Id` is used for `dbQuoteIdentifier()`, it needs at most one table component and at most one schema component.

**Usage**

```
## S4 method for signature 'PqConnection,Id'
dbQuoteIdentifier(conn, x, ...)

## S4 method for signature 'PqConnection,SQL'
dbQuoteIdentifier(conn, x, ...)

## S4 method for signature 'PqConnection,character'
dbQuoteIdentifier(conn, x, ...)

## S4 method for signature 'PqConnection'
dbQuoteLiteral(conn, x, ...)

## S4 method for signature 'PqConnection,SQL'
dbQuoteString(conn, x, ...)

## S4 method for signature 'PqConnection,character'
dbQuoteString(conn, x, ...)

## S4 method for signature 'PqConnection,SQL'
dbUnquoteIdentifier(conn, x, ...)
```

**Arguments**

conn	A <a href="#">PqConnection</a> created by <code>dbConnect()</code>
x	A character vector to be quoted.
...	Other arguments needed for compatibility with generic (currently ignored).

**Examples**

```
library(DBI)
con <- dbConnect(RPostgres::Postgres())

x <- c("a", "b c", "d'e", "\\f")
dbQuoteString(con, x)
dbQuoteIdentifier(con, x)
dbDisconnect(con)
```

---

Redshift

*Redshift driver/connection*


---

**Description**

Use `drv = Redshift()` instead of `drv = Postgres()` to connect to an AWS Redshift cluster. All methods in **RPostgres** and downstream packages can be called on such connections. Some have different behavior for Redshift connections, to ensure better interoperability.

**Usage**

```

Redshift()

## S4 method for signature 'RedshiftDriver'
dbConnect(
  drv,
  dbname = NULL,
  host = NULL,
  port = NULL,
  password = NULL,
  user = NULL,
  service = NULL,
  ...,
  bigint = c("integer64", "integer", "numeric", "character"),
  check_interrupts = FALSE,
  timezone = "UTC"
)

```

**Arguments**

drv	<a href="#">DBI::DBIDriver</a> . Use <a href="#">Postgres()</a> to connect to a PostgreSQL(-ish) database or <a href="#">Redshift()</a> to connect to an AWS Redshift cluster. Use an existing <a href="#">DBI::DBIConnection</a> object to clone an existing connection.
dbname	Database name. If NULL, defaults to the user name. Note that this argument can only contain the database name, it will not be parsed as a connection string (internally, <code>expand_dbname</code> is set to <code>false</code> in the call to <a href="#">PQconnectdbParams()</a> ).
host, port	Host and port. If NULL, will be retrieved from PGHOST and PGPORT env vars.
user, password	User name and password. If NULL, will be retrieved from PGUSER and PGPASSWORD envvars, or from the appropriate line in <code>~/ .pgpass</code> . See <a href="https://www.postgresql.org/docs/current/libpq-pgpass.html">https://www.postgresql.org/docs/current/libpq-pgpass.html</a> for more details.
service	Name of service to connect as. If NULL, will be ignored. Otherwise, connection parameters will be loaded from the <code>pg_service.conf</code> file and used. See <a href="https://www.postgresql.org/docs/current/libpq-pgservice.html">https://www.postgresql.org/docs/current/libpq-pgservice.html</a> for details on this file and syntax.
...	Other name-value pairs that describe additional connection options as described at <a href="https://www.postgresql.org/docs/current/libpq-connect.html#LIBPQ-PARAMKEYWORDS">https://www.postgresql.org/docs/current/libpq-connect.html#LIBPQ-PARAMKEYWORDS</a>
bigint	The R type that 64-bit integer types should be mapped to, default is <code>bit64::integer64</code> , which allows the full range of 64 bit integers.
check_interrupts	Should user interrupts be checked during the query execution (before first row of data is available)? Setting to TRUE allows interruption of queries running too long.
timezone	Sets the timezone for the connection. The default is "UTC". If NULL then no timezone is set, which defaults to the server's time zone.

# Index

bit64::integer64, [5](#), [14](#)

dbAppendTable, PqConnection-method  
(postgres-tables), [7](#)

dbAppendTable\_PqConnection  
(postgres-tables), [7](#)

dbBegin, PqConnection-method  
(postgres-transactions), [9](#)

dbBegin\_PqConnection  
(postgres-transactions), [9](#)

dbBind, PqResult-method  
(postgres-query), [5](#)

dbBind\_PqResult (postgres-query), [5](#)

dbClearResult, PqResult-method  
(postgres-query), [5](#)

dbClearResult\_PqResult  
(postgres-query), [5](#)

dbCommit, PqConnection-method  
(postgres-transactions), [9](#)

dbCommit\_PqConnection  
(postgres-transactions), [9](#)

dbConnect, PqDriver-method (Postgres), [4](#)

dbConnect, RedshiftDriver-method  
(Redshift), [13](#)

dbConnect\_PqDriver (Postgres), [4](#)

dbConnect\_RedshiftDriver (Redshift), [13](#)

dbDisconnect, PqConnection-method  
(Postgres), [4](#)

dbDisconnect\_PqConnection (Postgres), [4](#)

dbExistsTable, PqConnection, character-method  
(postgres-tables), [7](#)

dbExistsTable, PqConnection, Id-method  
(postgres-tables), [7](#)

dbExistsTable\_PqConnection\_character  
(postgres-tables), [7](#)

dbExistsTable\_PqConnection\_Id  
(postgres-tables), [7](#)

dbFetch, PqResult-method  
(postgres-query), [5](#)

dbFetch\_PqResult (postgres-query), [5](#)

dbHasCompleted, PqResult-method  
(postgres-query), [5](#)

dbHasCompleted\_PqResult  
(postgres-query), [5](#)

DBI::dbAppendTable(), [7](#)

DBI::dbBegin(), [11](#)

DBI::dbConnect(), [6](#), [8](#), [10](#), [11](#)

DBI::dbDataType(), [9](#)

DBI::dbGetRowsAffected(), [6](#)

DBI::DBIConnection, [4](#), [14](#)

DBI::DBIDriver, [4](#), [14](#)

DBI::dbQuoteIdentifier(), [8](#)

DBI::dbSendQuery(), [6](#)

DBI::dbWithTransaction(), [11](#)

DBI::dbWriteTable(), [7](#)

DBI::SQL(), [8](#)

dbListFields, PqConnection, character-method  
(postgres-tables), [7](#)

dbListFields, PqConnection, Id-method  
(postgres-tables), [7](#)

dbListFields\_PqConnection\_character  
(postgres-tables), [7](#)

dbListFields\_PqConnection\_Id  
(postgres-tables), [7](#)

dbListObjects, PqConnection-method  
(postgres-tables), [7](#)

dbListObjects\_PqConnection\_ANY  
(postgres-tables), [7](#)

dbListTables, PqConnection-method  
(postgres-tables), [7](#)

dbListTables\_PqConnection  
(postgres-tables), [7](#)

dbQuoteIdentifier, PqConnection, character-method  
(quote), [12](#)

dbQuoteIdentifier, PqConnection, Id-method  
(quote), [12](#)

dbQuoteIdentifier, PqConnection, SQL-method  
(quote), [12](#)

dbQuoteIdentifier\_PqConnection\_character

- (quote), [12](#)
- dbQuoteIdentifier\_PqConnection\_Id (quote), [12](#)
- dbQuoteIdentifier\_PqConnection\_SQL (quote), [12](#)
- dbQuoteLiteral, PqConnection-method (quote), [12](#)
- dbQuoteLiteral\_PqConnection (quote), [12](#)
- dbQuoteString, PqConnection, character-method (quote), [12](#)
- dbQuoteString, PqConnection, SQL-method (quote), [12](#)
- dbQuoteString\_PqConnection\_character (quote), [12](#)
- dbQuoteString\_PqConnection\_SQL (quote), [12](#)
- dbReadTable, PqConnection, character-method (postgres-tables), [7](#)
- dbReadTable\_PqConnection\_character (postgres-tables), [7](#)
- dbRemoveTable, PqConnection, character-method (postgres-tables), [7](#)
- dbRemoveTable\_PqConnection\_character (postgres-tables), [7](#)
- dbRollback, PqConnection-method (postgres-transactions), [9](#)
- dbRollback\_PqConnection (postgres-transactions), [9](#)
- dbSendQuery, PqConnection-method (postgres-query), [5](#)
- dbSendQuery\_PqConnection (postgres-query), [5](#)
- dbUnquoteIdentifier(), [8](#)
- dbUnquoteIdentifier, PqConnection, SQL-method (quote), [12](#)
- dbUnquoteIdentifier\_PqConnection\_SQL (quote), [12](#)
- dbWriteTable, PqConnection, character, data.frame-method (postgres-tables), [7](#)
- dbWriteTable\_PqConnection\_character\_data.frame (postgres-tables), [7](#)
  
- Id, [12](#)
- Id(), [8](#), [9](#)
  
- Postgres, [4](#)
- Postgres(), [4](#), [8](#), [10](#), [14](#)
- postgres-query, [5](#)
- postgres-tables, [7](#)
- postgres-transactions, [9](#)
- postgresDefault (postgresHasDefault), [10](#)
- postgresHasDefault, [10](#)
- postgresIsTransacting, [11](#)
- postgresWaitForNotify, [11](#)
- PqConnection, [6](#), [8](#), [10](#), [11](#), [13](#)
- PqResult, [6](#)
  
- quote, [12](#)
  
- Redshift, [13](#)
- Redshift(), [4](#), [8](#), [14](#)
- RedshiftConnection-class (Redshift), [13](#)
- RedshiftDriver-class (Redshift), [13](#)
- RPostgres (RPostgres-package), [3](#)
- RPostgres-package, [3](#)
  
- sqlData, PqConnection-method (postgres-tables), [7](#)
- sqlData\_PqConnection (postgres-tables), [7](#)
- Sys.timezone(), [5](#)