
sqlc

Release 1.10.0

Kyle Conroy

Sep 07, 2021

OVERVIEW

1	Installing sqlc	3
1.1	macOS	3
1.2	Ubuntu	3
1.3	go get	3
1.4	Docker	3
1.5	Downloads	3
2	Getting started with MySQL	5
3	Getting started with PostgreSQL	9
4	Retrieving rows	13
4.1	Selecting columns	15
4.2	Passing a slice as a parameter to a query	16
5	Counting rows	19
6	Inserting rows	21
6.1	Returning columns from inserted rows	22
7	Updating rows	25
8	Deleting rows	27
9	Preparing queries	29
10	Using transactions	31
11	Naming parameters	33
12	Modifying the database schema	35
12.1	Handling SQL migrations	35
13	Configuring generated structs	39
13.1	Naming scheme	39
13.2	JSON tags	39
14	CLI	41
15	Configuration file (version 1)	43
15.1	Type Overrides	44
15.2	Per-Column Type Overrides	45

15.3	Package Level Overrides	45
15.4	Renaming Struct Fields	45
16	Datatypes	47
16.1	Arrays	47
16.2	Dates and Time	47
16.3	Enums	48
16.4	Null	48
16.5	UUIDs	49
17	Query annotations	51
17.1	:exec	51
17.2	:execresult	51
17.3	:execrows	52
17.4	:many	52
17.5	:one	52
18	Database and Language Support	53
18.1	Future Language Support	53
18.2	Future Database Support	53
19	Developing sqlc	55
19.1	Building	55
19.2	Running Tests	55
19.3	Regenerate expected test output	56
20	Privacy and data collection	57
20.1	Our Privacy Pledge	57
20.2	Distribution Channels	57

And lo, the Great One looked down upon the people and proclaimed: “SQL is actually pretty great”

sqlc generates **fully type-safe idiomatic Go code** from SQL. Here’s how it works:

1. You write SQL queries
2. You run sqlc to generate Go code that presents type-safe interfaces to those queries
3. You write application code that calls the methods sqlc generated

Seriously, it’s that easy. You don’t have to write any boilerplate SQL querying code ever again.

INSTALLING SQLC

sqlc is distributed as a single binary with zero dependencies.

1.1 macOS

```
brew install sqlc
```

1.2 Ubuntu

```
sudo snap install sqlc
```

1.3 go get

```
go get github.com/kyleconroy/sqlc/cmd/sqlc
```

1.4 Docker

```
docker pull kjconroy/sqlc
```

Run sqlc using docker run:

```
docker run --rm -v $(pwd):/src -w /src kjconroy/sqlc generate
```

1.5 Downloads

Get pre-built binaries for *v1.10.0*:

- [Linux](#)
- [macOS](#)
- [Windows \(MySQL only\)](#)

GETTING STARTED WITH MYSQL

This tutorial assumes that the latest version of sqlc is [installed](#) and ready to use.

Create a new directory called `sqlc-tutorial` and open it up.

Initialize a new Go module named `tutorial.sqlc.dev/app`

```
go mod init tutorial.sqlc.dev/app
```

sqlc looks for either a `sqlc.yaml` or `sqlc.json` file in the current directory. In our new directory, create a file named `sqlc.yaml` with the following contents:

```
version: 1
packages:
  - path: "tutorial"
    name: "tutorial"
    engine: "mysql"
    schema: "schema.sql"
    queries: "query.sql"
```

sqlc needs to know your database schema and queries. In the same directory, create a file named `schema.sql` with the following contents:

```
CREATE TABLE authors (
  id BIGINT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  name text NOT NULL,
  bio text
);
```

Next, create a `query.sql` file with the following four queries:

```
-- name: GetAuthor :one
SELECT * FROM authors
WHERE id = ? LIMIT 1;

-- name: ListAuthors :many
SELECT * FROM authors
ORDER BY name;

-- name: CreateAuthor :execresult
INSERT INTO authors (
  name, bio
) VALUES (
  ?, ?
);
```

(continues on next page)

(continued from previous page)

```
-- name: DeleteAuthor :exec
DELETE FROM authors
WHERE id = ?;
```

You are now ready to generate code. Run the `generate` command. You shouldn't see any errors or output.

```
sqlc generate
```

You should now have a `db` package containing three files.

```
├── go.mod
├── query.sql
├── schema.sql
├── sqlc.yaml
├── tutorial
│   ├── db.go
│   ├── models.go
│   └── query.sql.go
```

You can use your newly generated queries in `app.go`.

```
package main

import (
    "context"
    "database/sql"
    "log"
    "reflect"

    "tutorial.sqlc.dev/app/tutorial"
    _ "github.com/go-sql-driver/mysql"
)

func run() error {
    ctx := context.Background()

    db, err := sql.Open("mysql", "user:password@dbname")
    if err != nil {
        return err
    }

    queries := tutorial.New(db)

    // list all authors
    authors, err := queries.ListAuthors(ctx)
    if err != nil {
        return err
    }
    log.Println(authors)

    // create an author
    result, err := queries.CreateAuthor(ctx, tutorial.CreateAuthorParams{
        Name: "Brian Kernighan",
        Bio:  sql.NullString{String: "Co-author of The C Programming Language, ↵
↵and The Go Programming Language", Valid: true},
```

(continues on next page)

(continued from previous page)

```
    })
    if err != nil {
        return err
    }

    insertedAuthorID, err := result.LastInsertId()
    if err != nil {
        return err
    }
    log.Println(insertedAuthorID)

    // get the author we just inserted
    fetchedAuthor, err := queries.GetAuthor(ctx, insertedAuthorID)
    if err != nil {
        return err
    }

    // prints true
    log.Println(reflect.DeepEqual(insertedAuthorID, fetchedAuthor.ID))
    return nil
}

func main() {
    if err := run(); err != nil {
        log.Fatal(err)
    }
}
```

Before the code will compile, you'll need to add the Go MySQL driver.

```
go get github.com/go-sql-driver/mysql
go build ./...
```

To make that possible, sqlc generates readable, **idiomatic** Go code that you otherwise would have had to write yourself. Take a look in `tutorial/query.sql.go`.

GETTING STARTED WITH POSTGRESQL

This tutorial assumes that the latest version of sqlc is [installed](#) and ready to use.

Create a new directory called `sqlc-tutorial` and open it up.

Initialize a new Go module named `tutorial.sqlc.dev/app`

```
go mod init tutorial.sqlc.dev/app
```

sqlc looks for either a `sqlc.yaml` or `sqlc.json` file in the current directory. In our new directory, create a file named `sqlc.yaml` with the following contents:

```
version: 1
packages:
  - path: "tutorial"
    name: "tutorial"
    engine: "postgresql"
    schema: "schema.sql"
    queries: "query.sql"
```

sqlc needs to know your database schema and queries. In the same directory, create a file named `schema.sql` with the following contents:

```
CREATE TABLE authors (
  id BIGSERIAL PRIMARY KEY,
  name text NOT NULL,
  bio text
);
```

Next, create a `query.sql` file with the following four queries:

```
-- name: GetAuthor :one
SELECT * FROM authors
WHERE id = $1 LIMIT 1;

-- name: ListAuthors :many
SELECT * FROM authors
ORDER BY name;

-- name: CreateAuthor :one
INSERT INTO authors (
  name, bio
) VALUES (
  $1, $2
)
```

(continues on next page)

(continued from previous page)

```
RETURNING *;

-- name: DeleteAuthor :exec
DELETE FROM authors
WHERE id = $1;
```

You are now ready to generate code. Run the `generate` command. You shouldn't see any errors or output.

```
sqlc generate
```

You should now have a `db` package containing three files.

```
├── go.mod
├── query.sql
├── schema.sql
├── sqlc.yaml
├── tutorial
│   ├── db.go
│   ├── models.go
│   └── query.sql.go
```

You can use your newly generated queries in `app.go`.

```
package main

import (
    "context"
    "database/sql"
    "log"
    "reflect"

    "tutorial.sqlc.dev/app/tutorial"

    _ "github.com/lib/pq"
)

func run() error {
    ctx := context.Background()

    db, err := sql.Open("postgres", "user=pqgotest dbname=pqgotest sslmode=verify-
↪full")
    if err != nil {
        return err
    }

    queries := tutorial.New(db)

    // list all authors
    authors, err := queries.ListAuthors(ctx)
    if err != nil {
        return err
    }
    log.Println(authors)

    // create an author
    insertedAuthor, err := queries.CreateAuthor(ctx, tutorial.CreateAuthorParams{
        Name: "Brian Kernighan",
```

(continues on next page)

(continued from previous page)

```
        Bio: sql.NullString{String: "Co-author of The C Programming Language,
↳and The Go Programming Language", Valid: true},
    })
    if err != nil {
        return err
    }
    log.Println(insertedAuthor)

    // get the author we just inserted
    fetchedAuthor, err := queries.GetAuthor(ctx, insertedAuthor.ID)
    if err != nil {
        return err
    }

    // prints true
    log.Println(reflect.DeepEqual(insertedAuthor, fetchedAuthor))
    return nil
}

func main() {
    if err := run(); err != nil {
        log.Fatal(err)
    }
}
```

Before the code will compile, you'll need to add the Go PostgreSQL driver.

```
go get github.com/lib/pq
go build ./...
```

To make that possible, sqlc generates readable, **idiomatic** Go code that you otherwise would have had to write yourself. Take a look in `tutorial/query.sql.go`.

RETRIEVING ROWS

To generate a database access method, annotate a query with a specific comment.

```
CREATE TABLE authors (  
  id          SERIAL PRIMARY KEY,  
  bio         text    NOT NULL,  
  birth_year int    NOT NULL  
);  
  
-- name: GetAuthor :one  
SELECT * FROM authors  
WHERE id = $1;  
  
-- name: ListAuthors :many  
SELECT * FROM authors  
ORDER BY id;
```

A few new pieces of code are generated beyond the `Author` struct. An interface for the underlying database is generated. The `*sql.DB` and `*sql.Tx` types satisfy this interface.

The database access methods are added to a `Queries` struct, which is created using the `New` method.

Note that the `*` in our query has been replaced with explicit column names. This change ensures that the query will never return unexpected data.

Our query was annotated with `:one`, meaning that it should only return a single row. We scan the data from that one into a `Author` struct.

Since the get query has a single parameter, the `GetAuthor` method takes a single `int` as an argument.

Since the list query has no parameters, the `ListAuthors` method accepts no arguments.

```
package db  
  
import (  
    "context"  
    "database/sql"  
)  
  
type Author struct {  
    ID          int  
    Bio         string  
    BirthYear  int  
}
```

(continues on next page)

```
type DBTX interface {
    QueryContext(context.Context, string, ...interface{}) (*sql.Rows, error)
    QueryRowContext(context.Context, string, ...interface{}) *sql.Row
}

func New(db DBTX) *Queries {
    return &Queries{db: db}
}

type Queries struct {
    db DBTX
}

const getAuthor = `-- name: GetAuthor :one
SELECT id, bio, birth_year FROM authors
WHERE id = $1
`

func (q *Queries) GetAuthor(ctx context.Context, id int) (Author, error) {
    row := q.db.QueryRowContext(ctx, getAuthor, id)
    var i Author
    err := row.Scan(&i.ID, &i.Bio, &i.BirthYear)
    return i, err
}

const listAuthors = `-- name: ListAuthors :many
SELECT id, bio, birth_year FROM authors
ORDER BY id
`

func (q *Queries) ListAuthors(ctx context.Context) ([]Author, error) {
    rows, err := q.db.QueryContext(ctx, listAuthors)
    if err != nil {
        return nil, err
    }
    defer rows.Close()
    var items []Author
    for rows.Next() {
        var i Author
        if err := rows.Scan(&i.ID, &i.Bio, &i.BirthYear); err != nil {
            return nil, err
        }
        items = append(items, i)
    }
    if err := rows.Close(); err != nil {
        return nil, err
    }
    if err := rows.Err(); err != nil {
        return nil, err
    }
    return items, nil
}
```

4.1 Selecting columns

```
CREATE TABLE authors (
  id          SERIAL PRIMARY KEY,
  bio         text    NOT NULL,
  birth_year int    NOT NULL
);

-- name: GetBioForAuthor :one
SELECT bio FROM authors
WHERE id = $1;

-- name: GetInfoForAuthor :one
SELECT bio, birth_year FROM authors
WHERE id = $1;
```

When selecting a single column, only that value that returned. The `GetBioForAuthor` method takes a single `int` as an argument and returns a `string` and an error.

When selecting multiple columns, a row record (method-specific struct) is returned. In this case, `GetInfoForAuthor` returns a struct with two fields: `Bio` and `BirthYear`.

```
package db

import (
    "context"
    "database/sql"
)

type DBTX interface {
    QueryRowContext(context.Context, string, ...interface{}) *sql.Row
}

func New(db DBTX) *Queries {
    return &Queries{db: db}
}

type Queries struct {
    db DBTX
}

const getBioForAuthor = `-- name: GetBioForAuthor :one
SELECT bio FROM authors
WHERE id = $1
`

func (q *Queries) GetBioForAuthor(ctx context.Context, id int) (string, error) {
    row := q.db.QueryRowContext(ctx, getBioForAuthor, id)
    var i string
    err := row.Scan(&i)
    return i, err
}

const getInfoForAuthor = `-- name: GetInfoForAuthor :one
SELECT bio, birth_year FROM authors
WHERE id = $1
`
```

(continues on next page)

(continued from previous page)

```

type GetInfoForAuthorRow struct {
    Bio          string
    BirthYear   int
}

func (q *Queries) GetInfoForAuthor(ctx context.Context, id int) (GetInfoForAuthorRow, error) {
    row := q.db.QueryRowContext(ctx, getInfoForAuthor, id)
    var i GetInfoForAuthorRow
    err := row.Scan(&i.Bio, &i.BirthYear)
    return i, err
}

```

4.2 Passing a slice as a parameter to a query

In PostgreSQL, `ANY` allows you to check if a value exists in an array expression. Queries using `ANY` with a single parameter will generate method signatures with slices as arguments. Use the postgres data types, eg: `int`, `varchar`, etc.

```

CREATE TABLE authors (
    id          SERIAL PRIMARY KEY,
    bio         text    NOT NULL,
    birth_year int    NOT NULL
);

-- name: ListAuthorsByIDs :many
SELECT * FROM authors
WHERE id = ANY($1::int[]);

```

The above SQL will generate the following code:

```

package db

import (
    "context"
    "database/sql"

    "github.com/lib/pq"
)

type Author struct {
    ID          int
    Bio         string
    BirthYear  int
}

type DBTX interface {
    QueryContext(context.Context, string, ...interface{}) (*sql.Rows, error)
    QueryRowContext(context.Context, string, ...interface{}) *sql.Row
}

func New(db DBTX) *Queries {
    return &Queries{db: db}
}

```

(continues on next page)

(continued from previous page)

```
type Queries struct {
    db DBTX
}

const listAuthors = `-- name: ListAuthorsByIDs :many
SELECT id, bio, birth_year FROM authors
WHERE id = ANY($1::int[])
`

func (q *Queries) ListAuthorsByIDs(ctx context.Context, ids []int) ([]Author, error) {
    rows, err := q.db.QueryContext(ctx, listAuthors, pq.Array(ids))
    if err != nil {
        return nil, err
    }
    defer rows.Close()
    var items []Author
    for rows.Next() {
        var i Author
        if err := rows.Scan(&i.ID, &i.Bio, &i.BirthYear); err != nil {
            return nil, err
        }
        items = append(items, i)
    }
    if err := rows.Close(); err != nil {
        return nil, err
    }
    if err := rows.Err(); err != nil {
        return nil, err
    }
    return items, nil
}
```


COUNTING ROWS

```
CREATE TABLE authors (  
  id SERIAL PRIMARY KEY,  
  hometown text NOT NULL  
);  
  
-- name: CountAuthors :one  
SELECT count(*) FROM authors;  
  
-- name: CountAuthorsByTown :many  
SELECT hometown, count(*) FROM authors  
GROUP BY 1  
ORDER BY 1;
```

```
package db  
  
import (  
  "context"  
  "database/sql"  
)  
  
type DBTX interface {  
  QueryContext(context.Context, string, ...interface{}) (*sql.Rows, error)  
  QueryRowContext(context.Context, string, ...interface{}) *sql.Row  
}  
  
func New(db DBTX) *Queries {  
  return &Queries{db: db}  
}  
  
type Queries struct {  
  db DBTX  
}  
  
const countAuthors = `-- name: CountAuthors :one  
SELECT count(*) FROM authors  
`  
  
func (q *Queries) CountAuthors(ctx context.Context) (int, error) {  
  row := q.db.QueryRowContext(ctx, countAuthors)  
  var i int  
  err := row.Scan(&i)  
  return i, err  
}
```

(continues on next page)

```
const countAuthorsByTown = `-- name: CountAuthorsByTown :many
SELECT hometown, count(*) FROM authors
GROUP BY 1
ORDER BY 1
`

type CountAuthorsByTownRow struct {
    Hometown string
    Count    int
}

func (q *Queries) CountAuthorsByTown(ctx context.Context) ([]CountAuthorsByTownRow, error) {
    rows, err := q.db.QueryContext(ctx, countAuthorsByTown)
    if err != nil {
        return nil, err
    }
    defer rows.Close()
    items := []CountAuthorsByTownRow{}
    for rows.Next() {
        var i CountAuthorsByTownRow
        if err := rows.Scan(&i.Hometown, &i.Count); err != nil {
            return nil, err
        }
        items = append(items, i)
    }
    if err := rows.Close(); err != nil {
        return nil, err
    }
    if err := rows.Err(); err != nil {
        return nil, err
    }
    return items, nil
}
```


INSERTING ROWS

```
CREATE TABLE authors (  
  id          SERIAL PRIMARY KEY,  
  bio        text    NOT NULL  
);  
  
-- name: CreateAuthor :exec  
INSERT INTO authors (bio) VALUES ($1);
```

```
package db  
  
import (  
  "context"  
  "database/sql"  
)  
  
type DBTX interface {  
  ExecContext(context.Context, string, ...interface{}) error  
}  
  
func New(db DBTX) *Queries {  
  return &Queries{db: db}  
}  
  
type Queries struct {  
  db DBTX  
}  
  
const createAuthor = `-- name: CreateAuthor :exec  
INSERT INTO authors (bio) VALUES ($1)  
`  
  
func (q *Queries) CreateAuthor(ctx context.Context, bio string) error {  
  _, err := q.db.ExecContext(ctx, createAuthor, bio)  
  return err  
}
```

6.1 Returning columns from inserted rows

sqlc has full support for the RETURNING statement.

```
CREATE TABLE authors (
  id          SERIAL PRIMARY KEY,
  bio        text    NOT NULL
);

-- name: Delete :exec
DELETE FROM authors WHERE id = $1;

-- name: DeleteAffected :execrows
DELETE FROM authors WHERE id = $1;

-- name: DeleteID :one
DELETE FROM authors WHERE id = $1
RETURNING id;

-- name: DeleteAuthor :one
DELETE FROM authors WHERE id = $1
RETURNING *;
```

```
package db

import (
    "context"
    "database/sql"
)

type Author struct {
    ID int
    Bio string
}

type DBTX interface {
    ExecContext(context.Context, string, ...interface{}) error
    QueryRowContext(context.Context, string, ...interface{}) error
}

func New(db DBTX) *Queries {
    return &Queries{db: db}
}

type Queries struct {
    db DBTX
}

const delete = `-- name: Delete :exec
DELETE FROM authors WHERE id = $1
`

func (q *Queries) Delete(ctx context.Context, id int) error {
    _, err := q.db.ExecContext(ctx, delete, id)
    return err
}
```

(continues on next page)

(continued from previous page)

```
const deleteAffected = `-- name: DeleteAffected :exec
DELETE FROM authors WHERE id = $1
`

func (q *Queries) DeleteAffected(ctx context.Context, id int) (int64, error) {
    result, err := q.db.ExecContext(ctx, deleteAffected, id)
    if err != nil {
        return 0, err
    }
    return result.RowsAffected()
}

const deleteID = `-- name: DeleteID :one
DELETE FROM authors WHERE id = $1
RETURNING id
`

func (q *Queries) DeleteID(ctx context.Context, id int) (int, error) {
    row := q.db.QueryRowContext(ctx, deleteID, id)
    var i int
    err := row.Scan(&i)
    return i, err
}

const deleteAuthor = `-- name: DeleteAuthor :one
DELETE FROM authors WHERE id = $1
RETURNING id, bio
`

func (q *Queries) DeleteAuthor(ctx context.Context, id int) (Author, error) {
    row := q.db.QueryRowContext(ctx, deleteAuthor, id)
    var i Author
    err := row.Scan(&i.ID, &i.Bio)
    return i, err
}
```


UPDATING ROWS

```
CREATE TABLE authors (  
  id          SERIAL PRIMARY KEY,  
  bio        text    NOT NULL  
);  
  
-- name: UpdateAuthor :exec  
UPDATE authors SET bio = $2  
WHERE id = $1;
```

```
package db  
  
import (  
    "context"  
    "database/sql"  
)  
  
type DBTX interface {  
    ExecContext(context.Context, string, ...interface{}) error  
}  
  
func New(db DBTX) *Queries {  
    return &Queries{db: db}  
}  
  
type Queries struct {  
    db DBTX  
}  
  
const updateAuthor = `-- name: UpdateAuthor :exec  
UPDATE authors SET bio = $2  
WHERE id = $1  
`  
  
func (q *Queries) UpdateAuthor(ctx context.Context, id int, bio string) error {  
    _, err := q.db.ExecContext(ctx, updateAuthor, id, bio)  
    return err  
}
```


DELETING ROWS

```
CREATE TABLE authors (  
  id          SERIAL PRIMARY KEY,  
  bio        text    NOT NULL  
);  
  
-- name: DeleteAuthor :exec  
DELETE FROM authors WHERE id = $1;
```

```
package db  
  
import (  
  "context"  
  "database/sql"  
)  
  
type DBTX interface {  
  ExecContext(context.Context, string, ...interface{}) error  
}  
  
func New(db DBTX) *Queries {  
  return &Queries{db: db}  
}  
  
type Queries struct {  
  db DBTX  
}  
  
const deleteAuthor = `-- name: DeleteAuthor :exec  
DELETE FROM authors WHERE id = $1  
`  
  
func (q *Queries) DeleteAuthor(ctx context.Context, id int) error {  
  _, err := q.db.ExecContext(ctx, deleteAuthor, id)  
  return err  
}
```


PREPARING QUERIES

```
CREATE TABLE records (  
  id SERIAL PRIMARY KEY  
);  
  
-- name: GetRecord :one  
SELECT * FROM records  
WHERE id = $1;
```

sqlc has an option to use prepared queries. These prepared queries also work with transactions.

```
package db  
  
import (  
  "context"  
  "database/sql"  
)  
  
type Record struct {  
  ID int  
}  
  
type DBTX interface {  
  PrepareContext(context.Context, string) (*sql.Stmt, error)  
  QueryRowContext(context.Context, string, ...interface{}) *sql.Row  
}  
  
func New(db DBTX) *Queries {  
  return &Queries{db: db}  
}  
  
func Prepare(ctx context.Context, db DBTX) (*Queries, error) {  
  q := Queries{db: db}  
  var err error  
  if q.getRecordStmt, err = db.PrepareContext(ctx, getRecord); err != nil {  
    return nil, err  
  }  
  return &q, nil  
}  
  
func (q *Queries) queryRow(ctx context.Context, stmt *sql.Stmt, query string, args ...  
  interface{}) (*sql.Row) {  
  switch {  
  case stmt != nil && q.tx != nil:  
    return q.tx.StmtContext(ctx, stmt).QueryRowContext(ctx, args...)
```

(continues on next page)

```
    case stmt != nil:
        return stmt.QueryRowContext(ctx, args...)
    default:
        return q.db.QueryRowContext(ctx, query, args...)
    }
}

type Queries struct {
    db          DBTX
    tx          *sql.Tx
    getRecordStmt *sql.Stmt
}

func (q *Queries) WithTx(tx *sql.Tx) *Queries {
    return &Queries{
        db:          tx,
        tx:          tx,
        getRecordStmt: q.getRecordStmt,
    }
}

const getRecord = `-- name: GetRecord :one
SELECT id FROM records
WHERE id = $1
`

func (q *Queries) GetRecord(ctx context.Context, id int) (Record, error) {
    row := q.queryRow(ctx, q.getRecordStmt, getRecord, id)
    var i Record
    err := row.Scan(&i.ID)
    return i, err
}
```

USING TRANSACTIONS

```
CREATE TABLE records (  
  id SERIAL PRIMARY KEY  
)  
;  
  
-- name: GetRecord :one  
SELECT * FROM records  
WHERE id = $1;
```

The `WithTx` method allows a `Queries` instance to be associated with a transaction.

```
package db  
  
import (  
  "context"  
  "database/sql"  
)  
  
type Record struct {  
  ID int  
}  
  
type DBTX interface {  
  QueryRowContext(context.Context, string, ...interface{}) *sql.Row  
}  
  
func New(db DBTX) *Queries {  
  return &Queries{db: db}  
}  
  
type Queries struct {  
  db DBTX  
}  
  
func (*Queries) WithTx(tx *sql.Tx) *Queries {  
  return &Queries{db: tx}  
}  
  
const getRecord = `-- name: GetRecord :one  
SELECT id FROM records  
WHERE id = $1  
`  
  
func (q *Queries) GetRecord(ctx context.Context, id int) (Record, error) {  
  row := q.db.QueryRowContext(ctx, getRecord, id)
```

(continues on next page)

(continued from previous page)

```
var i Record
err := row.Scan(&i.ID)
return i, err
}
```

NAMING PARAMETERS

sqlc tried to generate good names for positional parameters, but sometimes it lacks enough context. The following SQL generates parameters with less than ideal names:

```
-- name: UpsertAuthorName :one
UPDATE author
SET
  name = CASE WHEN $1::bool
    THEN $2::text
    ELSE name
  END
RETURNING *;
```

```
type UpdateAuthorNameParams struct {
  Column1      bool      `json:""`
  Column2_2    string    `json:"_2"`
}
```

In these cases, named parameters give you the control over field names on the Params struct.

```
-- name: UpsertAuthorName :one
UPDATE author
SET
  name = CASE WHEN sqlc.arg(set_name)::bool
    THEN sqlc.arg(name)::text
    ELSE name
  END
RETURNING *;
```

```
type UpdateAuthorNameParams struct {
  SetName      bool      `json:"set_name"`
  Name         string    `json:"name"`
}
```

If the `sqlc.arg()` syntax is too verbose for your taste, you can use the `@` operator as a shortcut.

```
-- name: UpsertAuthorName :one
UPDATE author
SET
  name = CASE WHEN @set_name::bool
    THEN @name::text
    ELSE name
  END
RETURNING *;
```


MODIFYING THE DATABASE SCHEMA

sqlc understands ALTER TABLE statements when parsing SQL.

```
CREATE TABLE authors (  
  id SERIAL PRIMARY KEY,  
  birth_year int NOT NULL  
);  
  
ALTER TABLE authors ADD COLUMN bio text NOT NULL;  
ALTER TABLE authors DROP COLUMN birth_year;  
ALTER TABLE authors RENAME TO writers;
```

```
package db  
  
type Writer struct {  
  ID int  
  Bio string  
}
```

12.1 Handling SQL migrations

sqlc will ignore rollback statements when parsing migration SQL files. The following tools are current supported:

- dbmate
- golang-migrate
- goose
- sql-migrate
- tern

12.1.1 goose

```
-- +goose Up  
CREATE TABLE post (  
  id int NOT NULL,  
  title text,  
  body text,  
  PRIMARY KEY(id)  
);
```

(continues on next page)

(continued from previous page)

```
-- +goose Down
DROP TABLE post;
```

```
package db

type Post struct {
    ID      int
    Title  sql.NullString
    Body   sql.NullString
}
```

12.1.2 sql-migrate

```
-- +migrate Up
-- SQL in section 'Up' is executed when this migration is applied
CREATE TABLE people (id int);

-- +migrate Down
-- SQL section 'Down' is executed when this migration is rolled back
DROP TABLE people;
```

```
package db

type People struct {
    ID      int32
}
```

12.1.3 tern

```
CREATE TABLE comment (id int NOT NULL, text text NOT NULL);
---- create above / drop below ----
DROP TABLE comment;
```

```
package db

type Comment struct {
    ID      int32
    Text   string
}
```


12.1.4 golang-migrate

In 20060102.up.sql:

```
CREATE TABLE post (  
  id    int NOT NULL,  
  title text,  
  body  text,  
  PRIMARY KEY(id)  
);
```

In 20060102.down.sql:

```
DROP TABLE post;
```

```
package db  
  
type Post struct {  
  ID    int  
  Title sql.NullString  
  Body  sql.NullString  
}
```

12.1.5 dbmate

```
-- migrate:up  
CREATE TABLE foo (bar INT NOT NULL);  
  
-- migrate:down  
DROP TABLE foo;
```

```
package db  
  
type Foo struct {  
  Bar int32  
}
```


CONFIGURING GENERATED STRUCTS

13.1 Naming scheme

Structs generated from tables will attempt to use the singular form of a table name if the table name is pluralized.

```
CREATE TABLE authors (  
  id SERIAL PRIMARY KEY,  
  name text NOT NULL  
);
```

```
package db  
  
// Struct names use the singular form of table names  
type Author struct {  
  ID int  
  Name string  
}
```

13.2 JSON tags

```
CREATE TABLE authors (  
  id SERIAL PRIMARY KEY,  
  created_at timestamp NOT NULL  
);
```

sqlc can generate structs with JSON tags. The JSON name for a field matches the column name in the database.

```
package db  
  
import (  
  "time"  
)  
  
type Author struct {  
  ID int `json:"id"`  
  CreatedAt time.Time `json:"created_at"`  
}
```



```
Usage:
  sqlc [command]

Available Commands:
  compile      Statically check SQL for syntax and type errors
  generate     Generate Go code from SQL
  help        Help about any command
  init        Create an empty sqlc.yaml settings file
  version     Print the sqlc version number

Flags:
  -h, --help  help for sqlc

Use "sqlc [command] --help" for more information about a command.
```


CONFIGURATION FILE (VERSION 1)

The `sqlc` tool is configured via a `sqlc.yaml` or `sqlc.json` file. This file must be in the directory where the `sqlc` command is run.

```
version: "1"
packages:
- name: "db"
  path: "internal/db"
  queries: "./sql/query/"
  schema: "./sql/schema/"
  engine: "postgresql"
  emit_prepared_queries: true
  emit_interface: false
  emit_exact_table_names: false
  emit_empty_slices: false
  emit_exported_queries: false
  emit_json_tags: true
  json_tags_case_style: "camel"
  output_db_file_name: "db.go"
  output_models_file_name: "models.go"
  output_querier_file_name: "querier.go"
```

Each package document has the following keys:

- `name`:
 - The package name to use for the generated code. Defaults to `path` basename
- `path`:
 - Output directory for generated code
- `queries`:
 - Directory of SQL queries or path to single SQL file; or a list of paths
- `schema`:
 - Directory of SQL migrations or path to single SQL file; or a list of paths
- `engine`:
 - Either `postgresql` or `mysql`. Defaults to `postgresql`.
- `sql_package`:
 - Either `pgx/v4` or `database/sql`. Defaults to `database/sql`.
- `emit_db_tags`:
 - If true, add DB tags to generated structs. Defaults to `false`.

- `emit_prepared_queries`:
 - If true, include support for prepared queries. Defaults to `false`.
- `emit_interface`:
 - If true, output a `Querier` interface in the generated package. Defaults to `false`.
- `emit_exact_table_names`:
 - If true, struct names will mirror table names. Otherwise, sqlc attempts to singularize plural table names. Defaults to `false`.
- `emit_empty_slices`:
 - If true, slices returned by `:many` queries will be empty instead of `nil`. Defaults to `false`.
- `emit_exported_queries`:
 - If true, autogenerated SQL statement can be exported to be accessed by another package.
- `emit_json_tags`:
 - If true, add JSON tags to generated structs. Defaults to `false`.
- `json_tags_case_style`:
 - `camel` for `camelCase`, `pascal` for `PascalCase`, `snake` for `snake_case` or `none` to use the column name in the DB. Defaults to `none`.
- `output_db_file_name`:
 - Customize the name of the db file. Defaults to `db.go`.
- `output_models_file_name`:
 - Customize the name of the models file. Defaults to `models.go`.
- `output_querier_file_name`:
 - Customize the name of the querier file. Defaults to `querier.go`.
- `output_files_suffix`:
 - If specified the suffix will be added to the name of the generated files.

15.1 Type Overrides

The default mapping of PostgreSQL types to Go types only uses packages outside the standard library when it must.

For example, the `uuid` PostgreSQL type is mapped to `github.com/google/uuid`. If a different Go package for UUIDs is required, specify the package in the `overrides` array. In this case, I'm going to use the `github.com/gofrs/uuid` instead.

```
version: "1"
packages: [...]
overrides:
- go_type: "github.com/gofrs/uuid.UUID"
  db_type: "uuid"
```

Each override document has the following keys:

- `db_type`:
 - The PostgreSQL type to override. Find the full list of supported types in [postgresql_type.go](https://github.com/sqlc-dev/sqlc/blob/main/postgresql_type.go).

- `go_type`:
 - A fully qualified name to a Go type to use in the generated code.
- `nullable`:
 - If true, use this type when a column is nullable. Defaults to `false`.

15.2 Per-Column Type Overrides

Sometimes you would like to override the Go type used in model or query generation for a specific field of a table and not on a type basis as described in the previous section.

This may be configured by specifying the `column` property in the override definition. `column` should be of the form `table.column` but you may be even more specific by specifying `schema.table.column` or `catalog.schema.table.column`.

```
version: "1"
packages: [...]
overrides:
  - column: "authors.id"
    go_type: "github.com/segmentio/ksuid.KSUID"
```

15.3 Package Level Overrides

Overrides can be configured globally, as demonstrated in the previous sections, or they can be configured on a per-package which scopes the override behavior to just a single package:

```
version: "1"
packages:
  - overrides: [...]
```

15.4 Renaming Struct Fields

Struct field names are generated from column names using a simple algorithm: split the column name on underscores and capitalize the first letter of each part.

```
account      -> Account
spotify_url  -> SpotifyUrl
app_id       -> AppID
```

If you're not happy with a field's generated name, use the `rename` dictionary to pick a new name. The keys are column names and the values are the struct field name to use.

```
version: "1"
packages: [...]
rename:
  spotify_url: "SpotifyURL"
```


DATATYPES

16.1 Arrays

PostgreSQL arrays are materialized as Go slices. Currently, only one-dimensional arrays are supported.

```
CREATE TABLE places (  
    name text not null,  
    tags text[]  
);
```

```
package db  
  
type Place struct {  
    Name string  
    Tags []string  
}
```

16.2 Dates and Time

All PostgreSQL time and date types are returned as `time.Time` structs. For null time or date values, the `NullTime` type from `database/sql` is used.

```
CREATE TABLE authors (  
    id SERIAL PRIMARY KEY,  
    created_at timestamp NOT NULL DEFAULT NOW(),  
    updated_at timestamp  
);
```

```
package db  
  
import (  
    "time"  
    "database/sql"  
)  
  
type Author struct {  
    ID int  
    CreatedAt time.Time  
    UpdatedAt sql.NullTime  
}
```

16.3 Enums

PostgreSQL `enums` are mapped to an aliased string type.

```
CREATE TYPE status AS ENUM (  
    'open',  
    'closed'  
);  
  
CREATE TABLE stores (  
    name text PRIMARY KEY,  
    status status NOT NULL  
);
```

```
package db  
  
type Status string  
  
const (  
    StatusOpen Status = "open"  
    StatusClosed Status = "closed"  
)  
  
type Store struct {  
    Name string  
    Status Status  
}
```

16.4 Null

For structs, null values are represented using the appropriate type from the `database/sql` package.

```
CREATE TABLE authors (  
    id SERIAL PRIMARY KEY,  
    name text NOT NULL,  
    bio text  
);
```

```
package db  
  
import (  
    "database/sql"  
)  
  
type Author struct {  
    ID int  
    Name string  
    Bio sql.NullString  
}
```

16.5 UUIDs

The Go standard library does not come with a `uuid` package. For UUID support, sqlc uses the excellent `github.com/google/uuid` package.

```
CREATE TABLE records (  
  id    uuid PRIMARY KEY  
);
```

```
package db  
  
import (  
    "github.com/google/uuid"  
)  
  
type Author struct {  
    ID    uuid.UUID  
}
```


QUERY ANNOTATIONS

sqlc requires each query to have a small comment indicating the name and command. The format of this comment is as follows:

```
-- name: <name> <command>
```

17.1 :exec

The generated method will return the error from `ExecContext`.

```
-- name: DeleteAuthor :exec  
DELETE FROM authors  
WHERE id = $1;
```

```
func (q *Queries) DeleteAuthor(ctx context.Context, id int64) error {  
    _, err := q.db.ExecContext(ctx, deleteAuthor, id)  
    return err  
}
```

17.2 :execresult

The generated method will return the `sql.Result` returned by `ExecContext`.

```
-- name: DeleteAllAuthors :execresult  
DELETE FROM authors;
```

```
func (q *Queries) DeleteAllAuthors(ctx context.Context) (sql.Result, error) {  
    return q.db.ExecContext(ctx, deleteAllAuthors)  
}
```

17.3 :execrows

The generated method will return the number of affected rows from the `result` returned by `ExecContext`.

```
-- name: DeleteAllAuthors :execrows
DELETE FROM authors;
```

```
func (q *Queries) DeleteAllAuthors(ctx context.Context) (int64, error) {
    _, err := q.db.ExecContext(ctx, deleteAllAuthors)
    // ...
}
```

17.4 :many

The generated method will return a slice of records via `QueryContext`.

```
-- name: ListAuthors :many
SELECT * FROM authors
ORDER BY name;
```

```
func (q *Queries) ListAuthors(ctx context.Context) ([]Author, error) {
    rows, err := q.db.QueryContext(ctx, listAuthors)
    // ...
}
```

17.5 :one

The generated method will return a single record via `QueryRowContext`.

```
-- name: GetAuthor :one
SELECT * FROM authors
WHERE id = $1 LIMIT 1;
```

```
func (q *Queries) GetAuthor(ctx context.Context, id int64) (Author, error) {
    row := q.db.QueryRowContext(ctx, getAuthor, id)
    // ...
}
```


DATABASE AND LANGUAGE SUPPORT

Language	MySQL	PostgreSQL
Go	Stable	Stable
Kotlin	Beta	Beta
Python	Experimental	Experimental

18.1 Future Language Support

- C#
- TypeScript

18.2 Future Database Support

- SQLite

DEVELOPING SQLC

19.1 Building

For local development, install `sqlc` under an alias. We suggest `sqlc-dev`.

```
go build -o ~/go/bin/sqlc-dev ./cmd/sqlc
```

19.2 Running Tests

```
go test ./...
```

To run the tests in the examples folder, use the `examples` tag.

```
go test --tags=examples ./...
```

These tests require locally-running database instances. Run these databases using [Docker Compose](#).

```
docker-compose up -d
```

The tests use the following environment variables to connect to the database

19.2.1 For PostgreSQL

Variable	Default Value
PG_HOST	127.0.0.1
PG_PORT	5432
PG_USER	postgres
PG_PASSWORD	mysecretpassword
PG_DATABASE	dinotest

19.2.2 For MySQL

Variable	Default Value
MYSQL_HOST	127.0.0.1
MYSQL_PORT	3306
MYSQL_USER	root
MYSQL_ROOT_PASSWORD	mysecretpassword
MYSQL_DATABASE	dinotest

19.3 Regenerate expected test output

If you need to update a large number of expected test output in the `internal/endtoend/testdata` directory, run the `regenerate.sh` script.

```
make regen
```

Note that this uses the `sqlc-dev` binary, not `sqlc` so make sure you have an up to date `sqlc-dev` binary.

PRIVACY AND DATA COLLECTION

These days, it feels like every piece of software is tracking you. From your browser, to your phone, to your terminal, programs collect as much data about you as possible and send it off to the cloud for analysis.

We believe the best way to keep data safe is to never collect it in the first place.

20.1 Our Privacy Pledge

The `sqlc` program does not collect any information. It does not send crash reports to a third-party. It does not gather anonymous aggregate user behaviour analytics.

No analytics. No finger-printing. No tracking.

Not now and not in the future.

20.2 Distribution Channels

We distribute `sqlc` using popular package managers such as [Homebrew](#) and [Snapcraft](#). These package managers and their associated command-line tools do collect usage metrics.

We use these services to make it easy to for users to install `sqlc`. There will always be an option to download `sqlc` from a stable URL.