PostgreSQL SQL Reference

Data Definition Language

- Create a database named `refcard` :

```
CREATE DATABASE refcard;
```

then connect to this new database.

- Create a schema named `myevents` and use it :

```
CREATE SCHEMA myevents;
SET search_path = myevents, public;
ALTER DATABASE refcard
  SET search_path FROM CURRENT;
```

- Create two tables named `users` and `schedule` :

```
CREATE TABLE users (
  id bigint GENERATED BY DEFAULT AS IDENTITY,
  owner text,
  e_mail text,
  valid boolean NOT NULL DEFAULT true );

CREATE TABLE schedule (
  id bigint GENERATED BY DEFAULT AS IDENTITY,
  owner bigint NOT NULL,
  date_ev tstzrange NOT NULL,
  description text,
  metadata jsonb );
```

- Alter tables to add primary keys :

```
ALTER TABLE users ADD PRIMARY KEY ( id );
ALTER TABLE schedule ADD PRIMARY KEY ( id );
```

each constraint implies a `btree` index.

- Add a foreign key between both tables :

```
ALTER TABLE schedule ADD FOREIGN KEY ( owner )
  REFERENCES users ( id );
```

- Create some GIN index :

```
CREATE INDEX schedule_metadata_idx
  ON schedule USING GIN (metadata);

CREATE INDEX schedule_desc_tsvector_idx
  ON schedule
  USING GIN (to_tsvector('english', description ));
```

- Add an extension and use it for create an index :

```
CREATE EXTENSION pg_trgm ;
CREATE INDEX users_owner_idx
  ON users USING GIN (owner gin_trgm_ops);
```

- Create an exclusion constraint :

```
CREATE EXTENSION btree_gist ;
ALTER TABLE schedule
  ADD CONSTRAINT unique_event_by_owner
  EXCLUDE USING gist
  ( owner WITH =, date_ev WITH && );
```

- Create a view :

```
CREATE VIEW schedule_working_hours AS
SELECT *
FROM schedule
WHERE lower(date_ev)::time
  BETWEEN time '9:00' and time '18:00';
```

Role management

- Create roles :

```
CREATE ROLE events_grp;
CREATE ROLE events_owner;
CREATE ROLE user_tests;
CREATE ROLE app_events LOGIN PASSWORD 'azerty'
  IN ROLE events_grp;
```

- Set Owner :

```
ALTER DATABASE refcard OWNER TO events_grp;
ALTER SCHEMA myevents OWNER TO events_grp;
ALTER TABLE users OWNER TO events_grp;
```

- Reassign roles :

```
REASSIGN OWNED BY events_grp TO events_owner;
```

- Add role into another role :

```
GRANT events_owner TO user_tests
  WITH INHERIT FALSE;
```

- Give some authorisation to roles :

```
GRANT USAGE ON SCHEMA myevents TO events_grp;
GRANT SELECT, INSERT, UPDATE, DELETE
  ON TABLE users, schedule TO events_grp;
```

- Drop Roles :

```
DROP ROLE IF EXISTS user_tests;
```

Data Manipulation Language

- Insert an `users` entry :

```
INSERT INTO users (owner, e_mail, valid)
VALUES ('Elaine Robinson'
       , 'e.rob@myhoster.dom', default );
```

- Insert a `schedule` entry :

```
INSERT INTO schedule (owner, date_ev, description, metadata)
VALUES ( 2, '[2026-01-07 18:30+01:00, 2026-01-07 23:30+01)'
       , 'meeting with Mrs.', '{"location": "her home"
          , "attendees": ["Glenda", "Me"]}'::jsonb);
```

- Copy data from a file :

```
COPY users FROM
  PROGRAM 'zcat /tmp/data/refcard.users.csv.gz'
  WITH FORMAT CSV;
```

- Copy data to a file :

```
COPY schedule TO '/tmp/data/refcard.schedule.csv'
  WITH FORMAT CSV;
```

- Update `users` table :

```
UPDATE users SET valid = true
WHERE e_mail = 'cs@smith.uny' ;
```

- Insert on conflict :

```
INSERT INTO users as u (id, owner, e_mail, valid )
VALUES (5, 'Carl Smith', 'cs@smith.uny', true )
      ON CONFLICT (id)
      DO UPDATE SET valid = EXCLUDED.valid
      WHERE u.valid <> EXCLUDED.valid ;
```

- Merge :

```
MERGE INTO users AS u
  USING (VALUES (5, 'Carl Smith', 'cs@smith.uny', true ) )
        AS d (id, owner, e_mail, valid )
  ON u.id = d.id
  WHEN MATCHED AND u.valid <> d.valid
    THEN UPDATE SET valid = d.valid
  WHEN NOT MATCHED
    THEN INSERT (id, owner, e_mail, valid )
      VALUES ( d.id, d.owner, d.e_mail, d.valid );
```

- Delete tuples :

```
DELETE FROM schedule
  WHERE date_ev && '[2024-01-07 17:05+00:00
                  , 2024-01-07 17:10:00+00)';
```

- Truncate tables and reset sequences :

```
TRUNCATE TABLE users, schedule RESTART IDENTITY;
```

© 2026 LOXODATA, licensed under CC BY-NC-SA 4.0

Data Query Language

- Join tables :

```
SELECT u.id, u.owner
      , lower( date_ev ) AS start_ev , description
  FROM users AS u
      INNER JOIN schedule AS p ON u.id=p.owner
 WHERE u.id = 2
 ORDER BY start_ev DESC ;
```

- Use Window Fonction :

```
SELECT u.id, u.owner
      , lower( date_ev ) AS start_ev
      , count( p.id ) OVER( PARTITION BY
         extract ('month' from lower( date_ev )))
  FROM users AS u
      INNER JOIN schedule AS p ON u.id=p.owner;
```

- Use CTE :

```
WITH cte_schedule AS (
SELECT u.id, u.owner , lower( p.date_ev ) AS start_ev
  , upper(lag( p.date_ev )
    OVER( PARTITION BY u.id ORDER BY p.date_ev )
       ) AS previous_end_ev
  FROM users AS u
      INNER JOIN schedule AS p ON u.id=p.owner )
SELECT id, owner, start_ev
      , start_ev - previous_end_ev AS itv_wh_prev_ev
  FROM cte_schedule;
```

- Use JSONPATH :

```
SELECT id
  , jsonb_path_query( metadata, '$.attendees.size()')
    as nb_attendees
  FROM myevents.schedule;
```

- Use grouping and aggregation :

```
SELECT count(*), metadata->'location' as location
  , date_trunc('day', lower(date_ev)) as day
  FROM schedule
 GROUP BY ROLLUP( metadata->'location'
              , date_trunc('day', lower(date_ev)))
 ORDER BY location, day;
```

Transaction Control Language

- Try Some `INSERT` :

```
BEGIN;
INSERT INTO schedule ( owner, date_ev
  , description, metadata ) VALUES
  ( 1, '[2026-01-07 19:30+01:00, 2026-01-07 21:30+01)'
    , 'meeting with Elain', '{"location": ["at the zoo"]
     , "attendees": {"Her": 1, "Me": 1}}'::jsonb);
INSERT INTO schedule ( owner, date_ev
  , description, metadata ) VALUES
  ( 2, '[2026-01-07 19:30+01:00, 2026-01-07 21:30+01)'
    , 'meeting with Ben', '{"location": ["at the zoo"]
     , "attendees": {"Him": 1, "Me": 1}}'::jsonb);
COMMIT;
```

Due to exclusion constraint, the second fails, then the whole transaction is rollback.

Maintenance commands

- Clean the table and collect statistics :

```
VACUUM ANALYZE users, schedule ;
```

- Rebuild the table and their index in new file :

```
VACUUM FULL users ;
```

- Remove transaction IDs for the oldest tuples :

```
VACUUM FREEZE schedule ;
```

- Rebuild an index :

```
REINDEX INDEX CONCURRENTLY schedule_desc_tsvector_idx ;
```

- Examine the execution plan for a query :

```
EXPLAIN (ANALYZE ON, COSTS OFF )
  SELECT id
    , jsonb_path_query( metadata, '$.attendees.size()')
     as nb_attendees FROM myevents.schedule;

ProjectSet (actual time=0.017..0.022 rows=6.00 loops=1)
  Buffers: shared hit=1
  -> Seq Scan on schedule (actual time=0.009..0.009
                               rows=6.00 loops=1)
       Buffers: shared hit=1
Planning Time: 0.051 ms
Execution Time: 0.038 ms
```

© 2026 LOXODATA, licensed under CC BY-NC-SA 4.0

get this card : https://www.loxodata.com/refcard

LOXODATA

Your PostgreSQL Partner

PostgreSQL SQL Reference Card

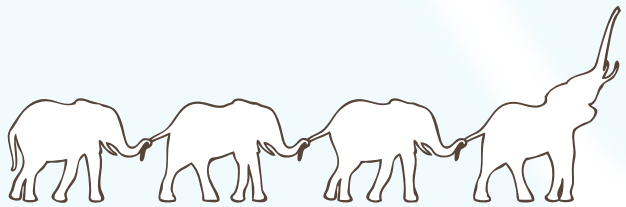

contact@loxodata.com +33 1 797 2 5775
www.loxodata.com