

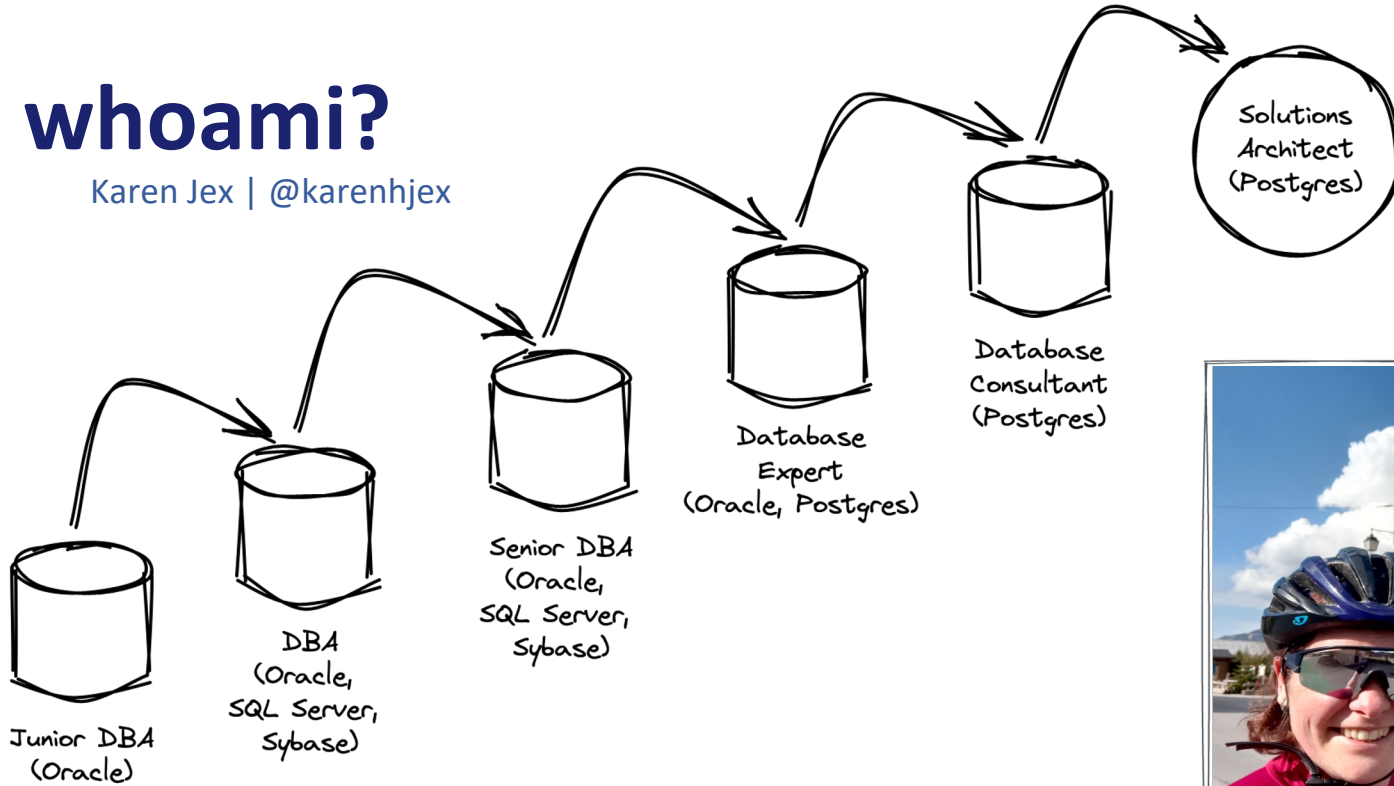


Everything You Wanted to Know about Databases as a Developer but Were Too Afraid to Ask Your DBA

Karen Jex | Senior Solutions Architect @ Crunchy Data
PGConf Europe, Berlin | October 2022

whoami?

Karen Jex | @karenhjex



Introduction

- Databases are essential to most applications
- Most developers aren't trained in database administration
- DBAs are ~~grumpy~~ busy people

Agenda

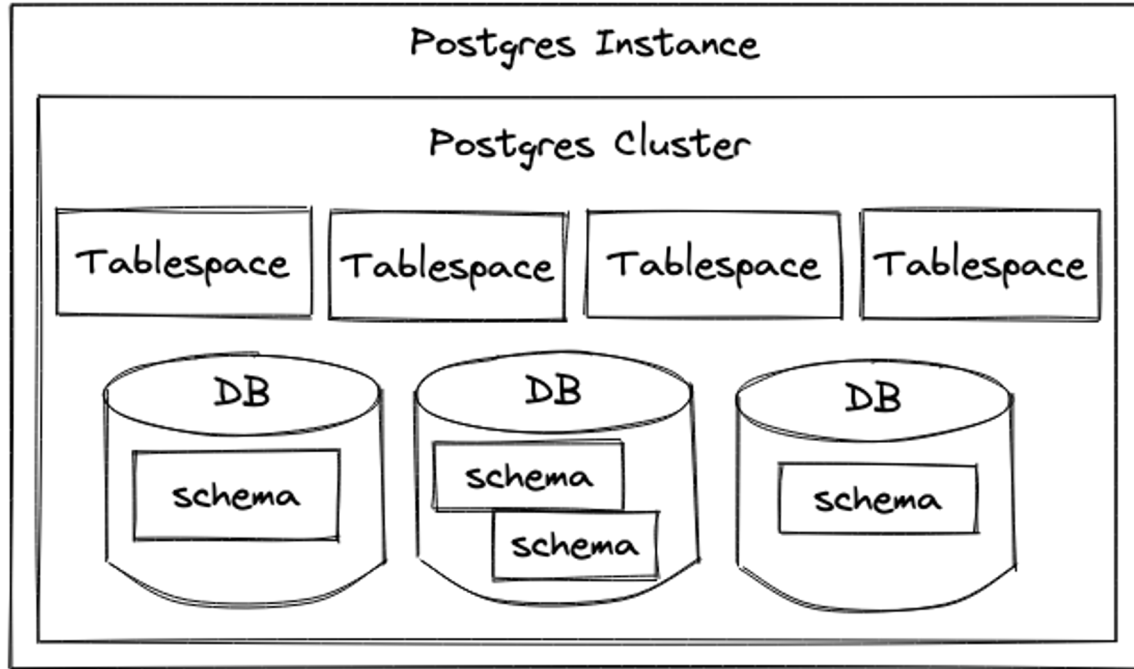
- Database Architecture
- Users and Roles
- Database Objects
- Database Connections
- Database Operations and Transactions
- WAL
- Documentation

Agenda

- **Database Architecture**
- Users and Roles
- Database Objects
- Database Connections
- Database Operations and Transactions
- WAL
- Documentation

Database Architecture

<https://www.postgresql.org/docs/current/tutorial-arch.html>



How do I Install PostgreSQL?

How do I Install Postgres?

- Install from source code

<https://www.postgresql.org/docs/current/install-procedure.html>

- Install package for given platform

<https://www.postgresql.org/download/>

- Choose a managed service

<https://crunchybridge.com/register>

- Try out the Postgres Playground (Postgres in your browser using WASM):

<https://www.crunchydata.com/developers/tutorials>

How do I Install Postgres?

Ubuntu 20.04 (Focal)

1. Configure the PostgreSQL repository

```
user@my_vm$ sudo sh -c 'echo "deb http://apt.postgresql.org/pub/repos/apt $(lsb_release -cs)-pgdg main"
> /etc/apt/sources.list.d/pgdg.list'
```

2. Import the repository signing key

```
user@my_vm$ wget --quiet -O - https://www.postgresql.org/media/keys/ACCC4CF8.asc | sudo apt-key add -
```

3. Update the package lists

```
user@my_vm$ sudo apt-get update
```

3. Install PostgreSQL

```
user@my_vm$ sudo apt-get -y install postgresql
```

How do I Install Postgres?

Centos7

1. Enable the pgdg repository

```
user@my_vm$ sudo yum -y install https://download.postgresql.org/pub/repos/yum/reporpms/EL-7-x86_64/pgdg-redhat-repo-latest.noarch.rpm
```

2. Install extra packages for Enterprise Linux

```
user@my_vm$ sudo yum -y install https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```

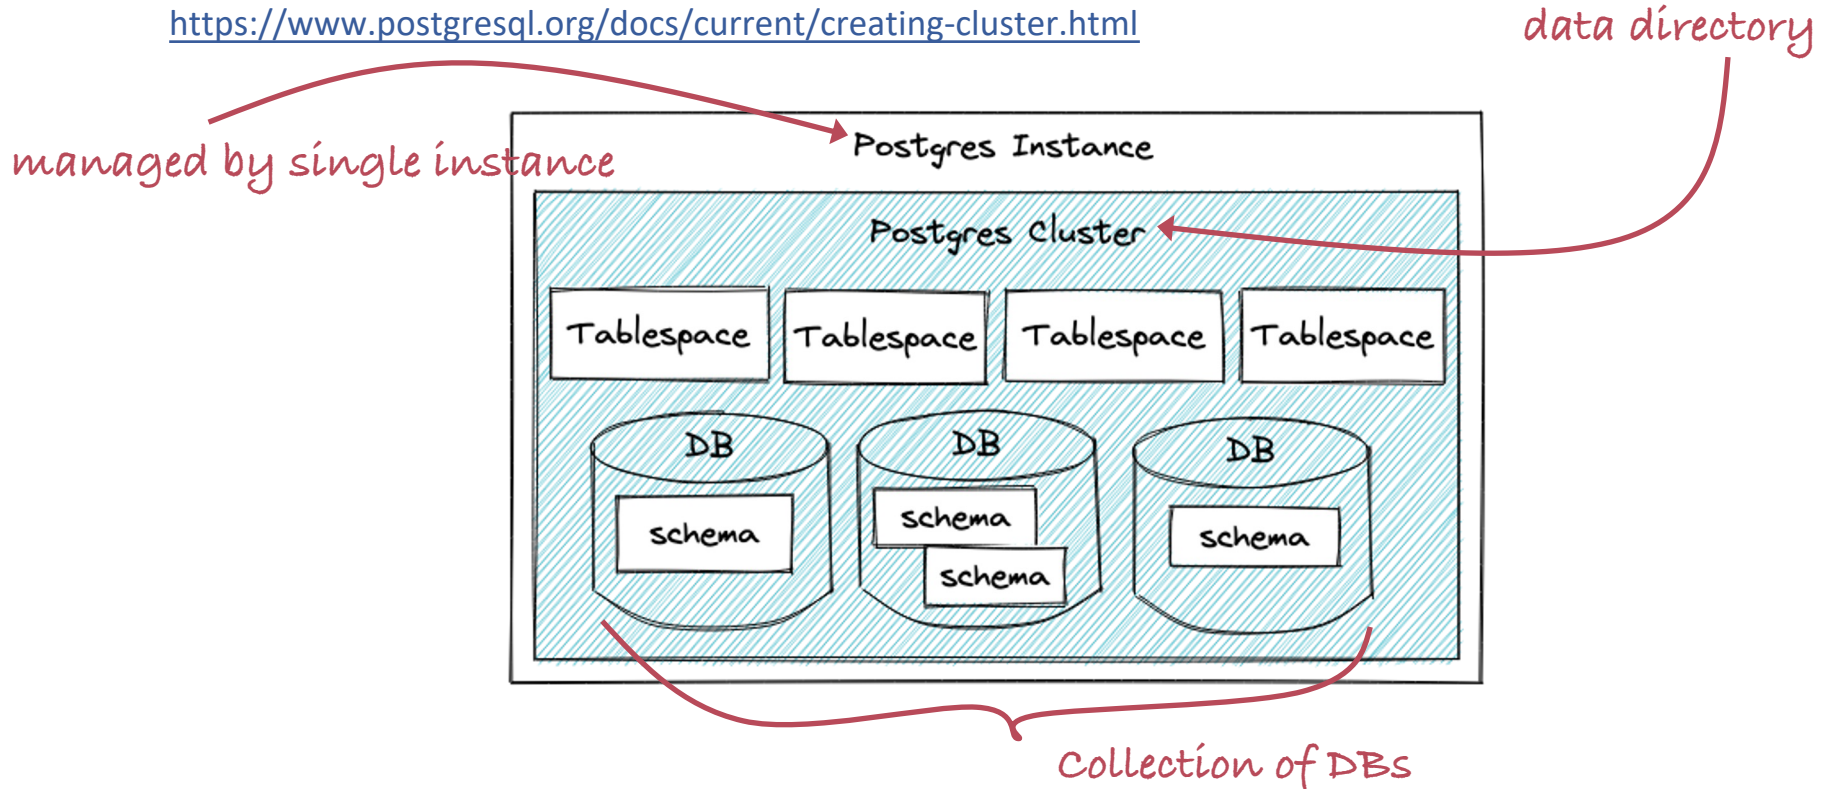
3. Install PostgreSQL

```
user@my_vm$ yum -y install postgresql14 postgresql14-server
```

What is a PostgreSQL Cluster?

What is a PostgreSQL Cluster?

<https://www.postgresql.org/docs/current/creating-cluster.html>



How do I Create a PostgreSQL Cluster?

How do I create a PostgreSQL Cluster?

<https://www.postgresql.org/docs/current/app-initdb.html>

```
postgres@my_vm$ export PATH=$PATH:/usr/pgsql-14/bin
```

```
postgres@my_vm$ export PGDATA=/var/lib/pgsql/14/data
```

```
postgres@my_vm$ initdb
```

```
...
```

Success. You can now start the database server using:

```
/usr/pgsql-14/bin/pg_ctl -D /var/lib/pgsql/14/data -l logfile start
```

How do I Start (or Stop) PostgreSQL?

How do I Start Postgres?

<https://www.postgresql.org/docs/current/app-pg-ctl.html>

```
postgres@my_vm$ pg_ctl -l logfile start
waiting for server to start.... done
server started
```


Postgres Processes

```
postgres@my_vm$ ps -ef|grep postgres
postgres  4236      1  0 11:59 ?        00:00:00 /usr/pgsql-14/bin/postgres
postgres  4237    4236  0 11:59 ?        00:00:00 postgres: logger
postgres  4239    4236  0 11:59 ?        00:00:00 postgres: checkpointer
postgres  4240    4236  0 11:59 ?        00:00:00 postgres: background writer
postgres  4241    4236  0 11:59 ?        00:00:00 postgres: walwriter
postgres  4242    4236  0 11:59 ?        00:00:00 postgres: autovacuum launcher
postgres  4243    4236  0 11:59 ?        00:00:00 postgres: stats collector
postgres  4244    4236  0 11:59 ?        00:00:00 postgres: logical replication launcher
```

How do I Stop Postgres?

<https://www.postgresql.org/docs/current/app-pg-ctl.html>

```
postgres@my_vm$ pg_ctl stop  
waiting for server to shut down.... done  
server stopped
```

How do I Control Postgres using systemd?

```
root@my_vm$ systemctl start|stop|restart postgresql-14
```

```
root@my_vm$ systemctl status postgresql-14
```

- postgresql-14.service - PostgreSQL 14 database server

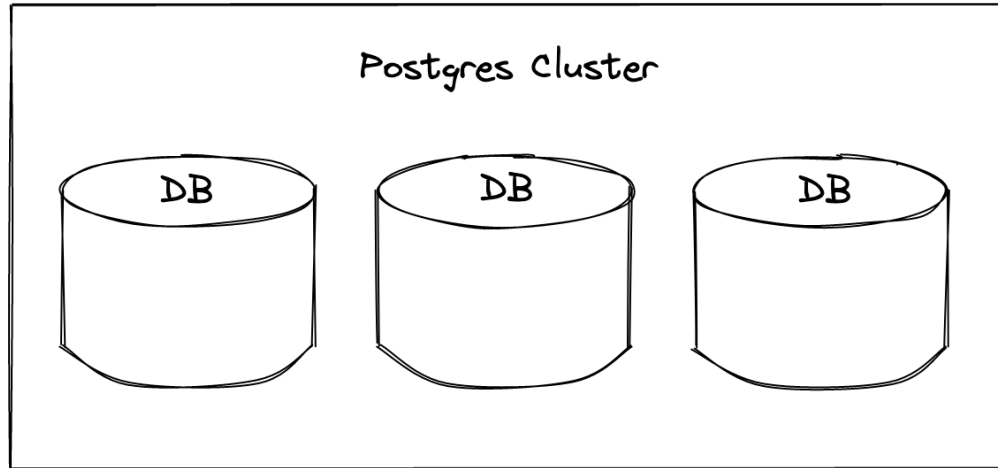
```
Loaded: loaded (/usr/lib/systemd/system/postgresql-14.service; disabled; vendor preset: disabled)
```

```
Active: active (running) since ...
```

What is a Database?

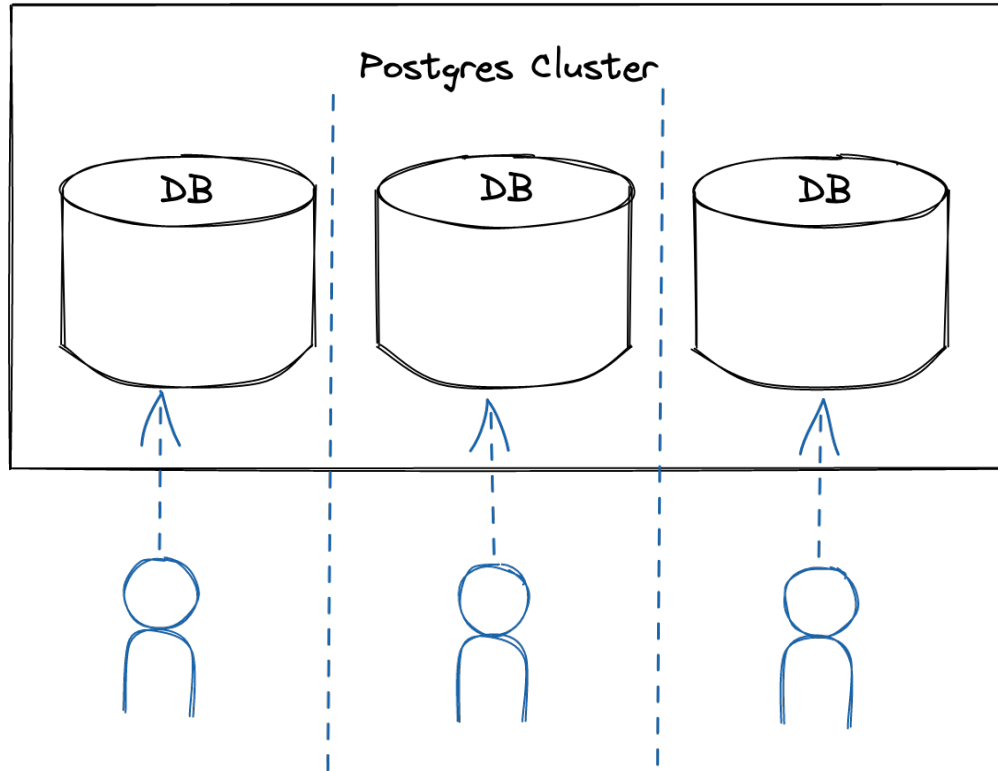
What is a Database?

<https://www.postgresql.org/docs/current/managing-databases.html>



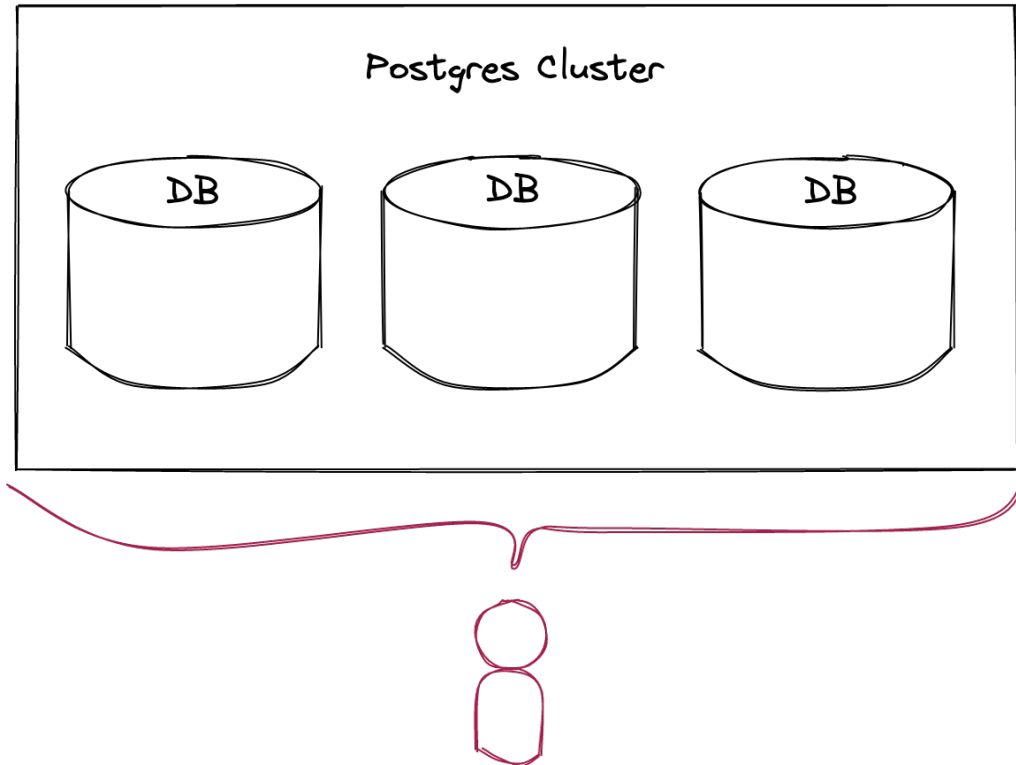
What is a Database?

<https://www.postgresql.org/docs/current/managing-databases.html>



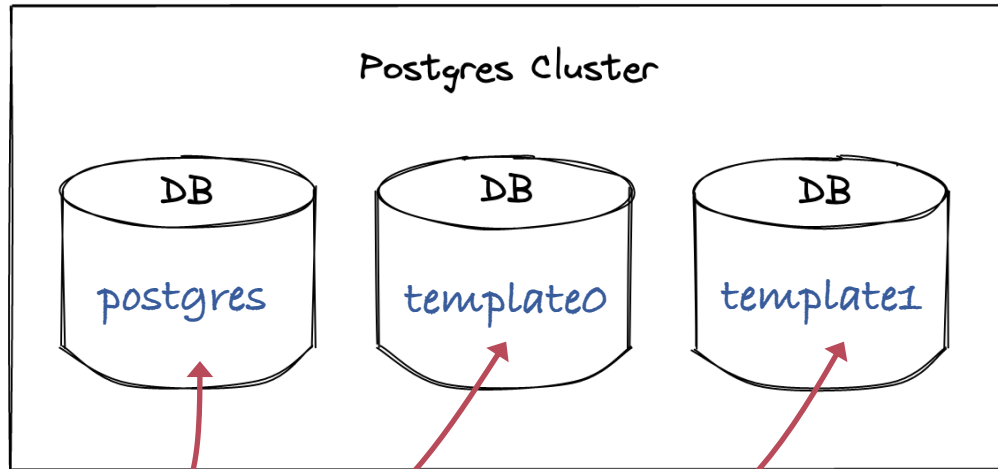
What is a Database?

<https://www.postgresql.org/docs/current/managing-databases.html>




What is a Database?

<https://www.postgresql.org/docs/current/managing-databases.html>



Pre-defined databases



Aside: What is psql?

Aside: What is a psql?

<https://www.postgresql.org/docs/current/app-psql.html>

- Command line tool
- Execute commands and/or scripts against the database
 - SQL
 - psql commands: e.g. \d to view table details

How do I Create a Database?

How do I Create a Database?

<https://www.postgresql.org/docs/current/sql-createdatabase.html>

```
postgres@my_vm$ psql
psql (14.0)
Type "help" for help.

postgres=# CREATE DATABASE my_database;
CREATE DATABASE
```

How do I List my Databases?

How do I List my Databases?

```
postgres=# \l
```

```
List of databases
```

Name	Owner	Encoding	Collate	Ctype	Access privileges
my_database	postgres	UTF8	en_US.UTF-8	en_US.UTF-8	
Postgres	postgres	UTF8	en_US.UTF-8	en_US.UTF-8	
template0	postgres	UTF8	en_US.UTF-8	en_US.UTF-8	=c/postgres + postgres=CtC/postgres
template1	postgres	UTF8	en_US.UTF-8	en_US.UTF-8	=c/postgres + postgres=CtC/postgres

```
(4 rows)
```

How do I List my Databases?

<https://www.postgresql.org/docs/current/catalogs.html>

```
postgres=# SELECT datname FROM pg_database;
```

```
datname
```

```
-----
```

```
my_database
```

```
Postgres
```

```
template0
```

```
template1
```

```
(4 rows)
```

What is a Tablespace?

What is a Tablespace?

And why should I create one?

- Physical location of the database objects
- Default tablespace `pg_default`
- Allows control of the disk layout of a PostgreSQL installation
- Accessible to the **entire cluster**

How do I Create a Tablespace?

How do I Create a Tablespace?

<https://www.postgresql.org/docs/current/sql-createtablespace.html>

```
postgres@my_vm$ mkdir -p /my_tablespace/tbsp_1
```

```
postgres@my_vm$ psql -c "CREATE TABLESPACE tablespace_1 location '/my_tablespace/tbsp_1'"
```

```
CREATE TABLESPACE
```

How do I List my Tablespaces?

How do I List my Tablespaces?

```
postgres=# \db
```

```
List of tablespaces
```

Name	Owner	Location
pg_default	postgres	
pg_global	postgres	
tablespace_1	postgres	/my_tablespaces/tbsp_1

(3 rows)

How do I List my Tablespaces?

```
postgres=# \db+
```

```
List of tablespaces
```

Name	Owner	Location	Access privileges	Options	Size	Description
pg_default	postgres				33 MB	
pg_global	postgres				560 kB	
tablespace_1	postgres	/my_tablespaces/tbsp_1			0 bytes	

```
(3 rows)
```

How do I List my Tablespaces?

<https://www.postgresql.org/docs/current/catalogs.html>

```
postgres@my_vm$ psql -c "select spcname from pg_tablespace"
```

```
      spcname
```

```
-----
```

```
pg_default
```

```
pg_global
```

```
tablespace_1
```

```
(3 rows)
```

What is a Schema?

What is a Schema?

Collection of objects within the database

- Logical grouping - no impact on physical location of objects
- Schema and owner of objects need not be the same
- Specific to the **database** in which it is created
- Namespace

How do I Create a Schema?

How do I Create a Schema?

<https://www.postgresql.org/docs/current/sql-createschema.html>

```
postgres@my_vm$ psql my_database -c "CREATE SCHEMA my_schema"
```

```
CREATE SCHEMA
```

How do I List my Schemas?

```
my_database=# \dn+
```

```
List of schemas
```

Name	Owner	Access privileges	Description
my_schema	postgres		
public	postgres	postgres=UC/postgres+ =UC/postgres	standard public schema

```
(2 rows)
```

Agenda

- Database Architecture
- **Users and Roles**
- Database Objects
- Database Connections
- Database Operations and Transactions
- WAL
- Documentation

What's the Difference Between a User and a Role?

What's the Difference Between a User and a Role?

- **CREATE USER** and **CREATE ROLE** are synonyms except:
 - CREATE USER: LOGIN by default
 - CREATE ROLE: NOLOGIN by default
- A **role** can be considered a **user** or a **group** (or both)
- A role is available to the **entire cluster**
- A role can own database objects

How do I Create a Role?

How do I Create a Role?

<https://www.postgresql.org/docs/current/sql-createrole.html>

```
postgres@my_vm$ psql -c "CREATE ROLE my_role"
```

```
CREATE ROLE
```

```
postgres@my_vm$ psql -d my_database -U my_role
```

```
psql: error: connection to server on socket "/var/run/postgresql/.s.PGSQL.5432" failed:
```

```
FATAL:  role "my_role" is not permitted to log in
```

How do I Create a User?

How do I Create a User?

<https://www.postgresql.org/docs/current/sql-createuser.html>

```
postgres@my_vm$ psql -c "CREATE USER my_user"
```

How do I Create a ~~User~~ Role?

<https://www.postgresql.org/docs/current/sql-createuser.html>

```
postgres@my_vm$ psql -c "CREATE USER my_user"
```

```
CREATE ROLE
```

```
postgres@my_vm$ psql -d my_database -U my_user
```

```
psql (14.0)
```

```
Type "help" for help.
```

What is a Privilege?

What is a Privilege?

Permission to perform certain action(s) on given object(s)

- Granted by the **owner** of the object or by a **superuser**
 - Objects: DATABASE, FUNCTION, SCHEMA, TABLE ...
 - Privileges: SELECT, INSERT, UPDATE, DELETE, TRUNCATE, CREATE ...
- ALTER DEFAULT PRIVILEGES

Granting Privileges

<https://www.postgresql.org/docs/current/ddl-priv.html>

```
postgres@my_vm$ psql -d my_database
```

```
my_database=# GRANT CREATE ON DATABASE my_database TO my_user;  
GRANT
```

```
my_database=# GRANT CREATE ON SCHEMA my_schema TO my_user;  
GRANT
```

```
my_database=# GRANT CREATE ON TABLESPACE tablespace_1 TO my_user;  
GRANT
```

Agenda

- Database Architecture
- Users and Roles
- **Database Objects**
- Database Connections
- Database Operations and Transactions
- WAL
- Documentation

Database Objects

- Table
- Index
- Constraint
- View
- Materialized View
- Sequence

What is a Table?

What is a Table?

- A “relation”
- Data arranged in columns and rows
- Rows are not ordered

dept_id	dept_name

dept

emp_id	emp_name	dept_id

emp

How do I Create a Table?

How do I Create a Table?

<https://www.postgresql.org/docs/current/sql-createtable.html>

```
postgres@my_vm$ psql -d my_database -U my_user
```

```
my_database=> CREATE TABLE my_schema.dept (  
my_database(>   dept_id integer,  
my_database(>   dept_name varchar);  
my_database-> TABLESPACE tablespace_1;  
CREATE TABLE
```

What is a Sequence?

What is a Sequence?

Object that generates a sequence of integers

- Specific to a **schema**
- Used to generate unique numeric identifiers
- Implemented as a single row table

How do I Create a Sequence?

How do I Create a Sequence?

<https://www.postgresql.org/docs/current/sql-createsequence.html>

```
my_database=> CREATE SEQUENCE my_schema.s_dept
my_database-> INCREMENT BY 1 MINVALUE 1 NO MAXVALUE;
CREATE SEQUENCE

my_database=> ALTER TABLE my_schema.dept
my_database-> ALTER COLUMN dept_id SET DEFAULT nextval('my_schema.s_dept');
ALTER TABLE
```

How do I Auto Generate an ID Column?

How do I Auto Generate an ID Column?

<https://www.postgresql.org/docs/current/sql-createtable.html>

```
my_database=> CREATE TABLE my_schema.emp (  
my_database(>   emp_id integer,  
my_database(>   emp_name varchar,  
my_database(>   dept_id integer)  
my_database-> TABLESPACE tablespace_1;
```

How do I Auto Generate an ID Column?

<https://www.postgresql.org/docs/current/sql-createtable.html>

```
my_database=> CREATE TABLE my_schema.emp (  
my_database(>   emp_id integer generated always as identity,  
my_database(>   emp_name varchar,  
my_database(>   dept_id integer)  
my_database-> TABLESPACE tablespace_1;  
CREATE TABLE
```

How do I View the Table Definitions?

How do I View the Table Definitions?

```
my_database=> \d my_schema.dept
```

```
Table "my_schema.dept"
  Column |          Type          | Collation | Nullable |          Default
-----+-----+-----+-----+-----
 dept_id | integer                |           |          | nextval('s_dept'::regclass)
 dept_name | character varying      |           |          |
Tablespace: "tablespace_1"
```

How do I View the Table Definitions?

```
my_database=> \d my_schema.emp
```

```
Table "my_schema.emp"
  Column |          Type          | Collation | Nullable |          Default
-----+-----+-----+-----+-----
 emp_id  | integer                |           | not null | generated always as identity
 emp_name | character varying     |           |          |
 dept_id | integer                |           |          |
Tablespace: "tablespace_1"
```

What is an Index?

What is an Index?

Ordered list of entries containing a value and a pointer to the table row

- Can create on one or more columns or expressions
- Can speed up searches based on values of column(s) in the index
- Default type is **btree**
- Various types available in Postgres

<https://www.postgresql.org/docs/current/indexes-types.html>

How do I Create an Index?

How do I Create an Index?

<https://www.postgresql.org/docs/current/sql-createindex.html>

```
my_database=> CREATE INDEX emp_dept_id  
my_database-> ON my_schema.emp (dept_id)  
my_database-> TABLESPACE tablespace_1;  
CREATE INDEX
```

What is a Constraint?

What is a Constraint?

<https://www.postgresql.org/docs/current/ddl-constraints.html>

- NOT NULL constraint
- CHECK constraint
- PRIMARY KEY (PK) constraint
- FOREIGN KEY (FK) constraint

What is a NOT NULL Constraint?

- Column must contain a value

What is a CHECK Constraint?

- Value must conform to certain rules
- for example: dept_name must contain only letters and spaces

What is a Primary Key?

One or more columns that allow a row in a table to be identified uniquely

- Enforced through PK constraint + unique index
- A table may have **only one** PK constraint
- The columns in the PK must be **NOT NULL**

How do I Create a Primary Key Constraint?

How do I Create a PK Constraint?

<https://www.postgresql.org/docs/current/sql-altertable.html>

```
my_database=> ALTER TABLE my_schema.emp
my_database-> ADD CONSTRAINT emp_pk PRIMARY KEY (emp_id)
my_database-> USING INDEX TABLESPACE tablespace_1;
ALTER TABLE
my_database=> \d my_schema.emp
```

```
Table "my_schema.emp"
  Column |          Type          | Collation | Nullable |          Default
-----+-----+-----+-----+-----
 emp_id  | integer                |           | not null | generated always as identity
 emp_name | character varying     |           |          |
 dept_id | integer                |           |          |
```

Indexes:

```
    "emp_pk" PRIMARY KEY, btree (emp_id), tablespace "tablespace_1"
```

```
    "emp_dept_id" btree (dept_id), tablespace "tablespace_1"
```

Tablespace: "tablespace_1"

How do I Create a PK Constraint?

<https://www.postgresql.org/docs/current/sql-altertable.html>

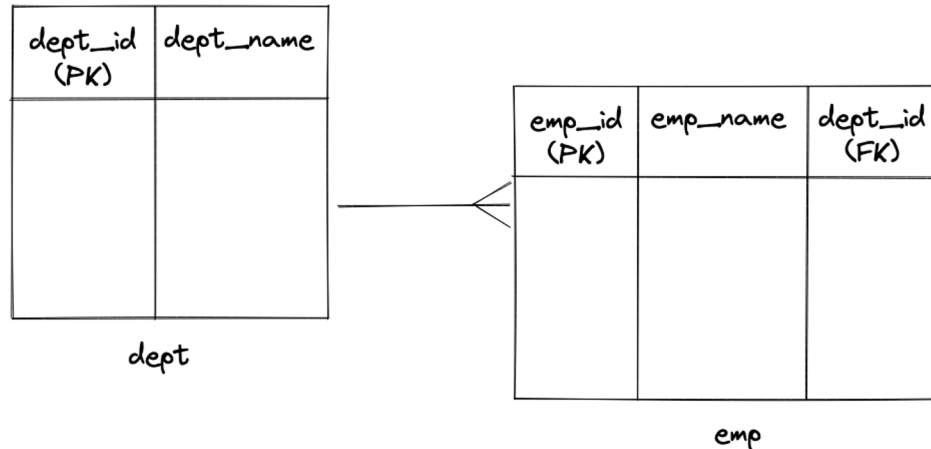
```
my_database=> ALTER TABLE my_schema.dept
my_database-> ADD CONSTRAINT dept_pk PRIMARY KEY (dept_id)
my_database-> USING INDEX TABLESPACE tablespace_1;
ALTER TABLE
my_database=> \d my_schema.dept
```

```
Table "my_schema.dept"
  Column |          Type          | Collation | Nullable |          Default
-----+-----+-----+-----+-----
 dept_id | integer                |           | not null | generated always as identity
 dept_name | character varying     |           |         |
Indexes:
    "dept_pk" PRIMARY KEY, btree (dept_id), tablespace "tablespace_1"
Tablespace: "tablespace_1"
```

What is a Foreign Key?

A relationship between a “parent” and a “child” table

- A way to enforce “referential integrity”
- References the primary key or another unique key in the parent table



How do I Create a Foreign Key Constraint?

How do I Create a FK Constraint?

<https://www.postgresql.org/docs/current/sql-altertable.html>

```
my_database=> ALTER TABLE my_schema.emp
my_database-> ADD CONSTRAINT emp_dept_fk FOREIGN KEY (dept_id)
my_database-> REFERENCES my_schema.dept(dept_id);
ALTER TABLE
my_database=> \d my_schema.emp
...
Foreign-key constraints:
    "emp_dept_fk" FOREIGN KEY (dept_id) REFERENCES my_schema.dept(dept_id)
...

my_database=> \d my_schema.dept
...
Referenced by:
    TABLE "emp" CONSTRAINT "emp_dept_fk" FOREIGN KEY (dept_id) REFERENCES dept(dept_id)
...
```

Aside: What is a Search Path?

Aside: What is a Search Path?

Schema(s) that will be searched if I don't refer to an object using a fully-qualified object name

```
my_database=> SHOW search_path;
search_path
-----
"$user", public

my_database=> select * from emp;
ERROR:  relation "emp" does not exist
LINE 1: select * from emp
                    ^
```

Aside: What is a Search Path?

Schema(s) that will be searched if I don't refer to an object using a fully-qualified object name

```
my_database=> SET search_path = my_schema;
SET
my_database=> SHOW search_path;
 search_path
-----
 my_schema
(1 row)
my_database=> select * from emp;
 emp_id | emp_name | dept_id
-----+-----+-----
(0 rows)
```


Aside: What is a Search Path?

I can set the search path (and other options) automatically

```
postgres@my_vm$ echo 'set search_path to my_schema' >> ~/.psqlrc
```

```
postgres@my_vm$ psql -d my_database -U my_user
```

```
SET
```

```
psql (14.0)
```

```
Type "help" for help.
```

```
my_database=> SHOW search_path;
```

```
search_path
```

```
-----
```

```
my_schema
```

```
(1 row)
```

How do I Populate my Tables?

How do I Populate my Tables?

<https://www.postgresql.org/docs/current/sql-insert.html>

```
my_database=> INSERT INTO dept (dept_name)
my_database-> VALUES ('Sales'),('Consulting'),('Product'),('HR');
INSERT 0 4
my_database=> SELECT * FROM dept ORDER BY dept_id;
 dept_id | dept_name
-----+-----
      1 | Sales
      2 | Consulting
      3 | Product
      4 | HR
(4 rows)
```

How do I Populate my Tables?

<https://www.postgresql.org/docs/current/sql-insert.html>

```
my_database=> INSERT INTO emp (emp_name, dept_id)
my_database->VALUES ('Ay Bee',2),('Cee Dee',2),('E.F. Gee',1),('Aitch Eye',null),('Jay Kay',4);
INSERT 0 4
my_database=> SELECT * FROM emp ORDER BY emp_id;
 emp_id | emp_name  | dept_id
-----+-----+-----
      1 | Ay Bee   |      2
      2 | Cee Dee  |      2
      3 | E.F. Gee |      1
      4 | Aitch Eye|
      5 | Jay Kay  |      4
(5 rows)
```

What if I try to Insert an Invalid dept_id?

<https://www.postgresql.org/docs/current/sql-insert.html>

```
my_database=> INSERT INTO emp (emp_name, dept_id)
my_database-> VALUES ('No Good',7);
ERROR:  insert or update on table "emp" violates foreign key constraint "emp_dept_fk"
DETAIL:  Key (dept_id)=(7) is not present in table "dept".
```

What is a View?

What is a View?

A virtual table, based on a query

- Does not take up any space
- Executed in real-time
- Shorthand for a long query
- Present just certain data to certain users

How do I Create a View?

How do I Create a View?

<https://www.postgresql.org/docs/current/sql-createview.html>

```
my_database=> CREATE VIEW emp_name_view AS
my_database-> SELECT emp_name AS "employee_name" FROM emp ORDER BY emp_name;
CREATE VIEW
my_database=> SELECT * FROM emp_name_view;
 employee_name
-----
Aitch Eye
Ay Bee
Cee Dee
E.F. Gee
Jay Kay
(5 rows)
```

What is a Materialized View?

What is a Materialized View?

A table that contains the results of a query

- Query is not executed in real-time
- Can be “refreshed” (query re-executed to gather latest results)
- Useful for aggregating data
- Avoids re-executing long-running/frequently executed queries

How do I Create a Materialized View?

How do I Create a Materialized View?

<https://www.postgresql.org/docs/current/sql-creatematerializedview.html>

```
my_database=> CREATE MATERIALIZED VIEW emp_mview
my_database-> TABLESPACE tablespace_1 AS
my_database-> SELECT emp_id, emp_name AS "employee_name" FROM emp WHERE dept_id in (1,2,3);
SELECT 3
```

```
my_database=> SELECT * FROM emp_mview;
```

```
emp_id | employee_name
-----+-----
      1 | Ay Bee
      2 | Cee Dee
      3 | E.F. Gee
```

```
(3 rows)
```

How do I Refresh a Materialized View?

First, update some data...

```
my_database=> UPDATE emp SET emp_name = 'CEE DEE' WHERE emp_id = 2;
UPDATE 1
```

```
my_database=> SELECT emp_name FROM emp WHERE emp_id = 2;
 emp_name
-----
 CEE DEE
```

```
my_database=> SELECT employee_name FROM emp_name_view WHERE employee_name like 'C%';
 employee_name
-----
 CEE DEE
(1 row)
```

How do I Refresh a Materialized View?

<https://www.postgresql.org/docs/current/sql-refreshmaterializedview.html>

```
my_database=>SELECT employee_name FROM emp_mview where emp_id = 2;
employee_name
-----
Cee Dee
(1 row)
```

```
my_database=> REFRESH MATERIALIZED VIEW emp_mview;
REFRESH MATERIALIZED VIEW
```

```
my_database=> SELECT employee_name FROM emp_mview where emp_id = 2;
employee_name
-----
CEE DEE
```

Agenda

- Database Architecture
- Users and Roles
- Database Objects
- **Database Connections**
- Database Operations and Transactions
- WAL
- Documentation

What Information do I Need?

- host
- port
- database
- username
- password/certificate

What Information do I Need?

Format depends on client tool/driver

name of database to connect to

database host

psql example: `psql -h 127.0.0.1 -p 5432 -U myuser mydatabase`

jdbc example: `jdbc:postgresql://myuser@127.0.0.1:5432/mydatabase`

database username

database port (5432 by default)

What Client Tools are Available?

Many client tools allow connection to Postgres, including:

- psql
- DBeaver (multi-platform)

<https://dbeaver.io/>

- pgAdmin4

<https://www.pgadmin.org/>

Agenda

- Database Architecture
- Users and Roles
- Database Objects
- Database Connections
- **Database Operations and Transactions**
- WAL
- Documentation

Database Operations and Transactions

- Transactions
- Commit
- Rollback
- SQL, DML, DDL
- Join Types
- Execution Plans

What is a Transaction?

What is a Transaction?

A single unit of work that consists of one or more operations

- Ends with a COMMIT or a ROLLBACK:
 - COMMIT: operations permanently applied to database
 - ROLLBACK: operations cancelled
- Changes only visible to other transactions after COMMIT
- Certain locks retained until end of transaction
- ACID (Atomic, Consistent, Isolated, Durable)

How do I Begin/End a Transaction?

How do I Begin/End a Transaction?

Example for psql: default behaviour is autocommit

- No need to explicitly BEGIN or COMMIT any transactions
- Each command is a distinct transaction with implicit COMMIT
- To manage a transaction manually:
 - Issue BEGIN to start a transaction
 - Execute the operations that comprise the transaction
 - Issue COMMIT to make the changes permanent or/
 - Issue ROLLBACK to undo the changes

What is SQL/DML/DDL?

What is SQL/DML/DDL?

<https://www.postgresql.org/docs/current/sql.html>

SQL Structured Query Language

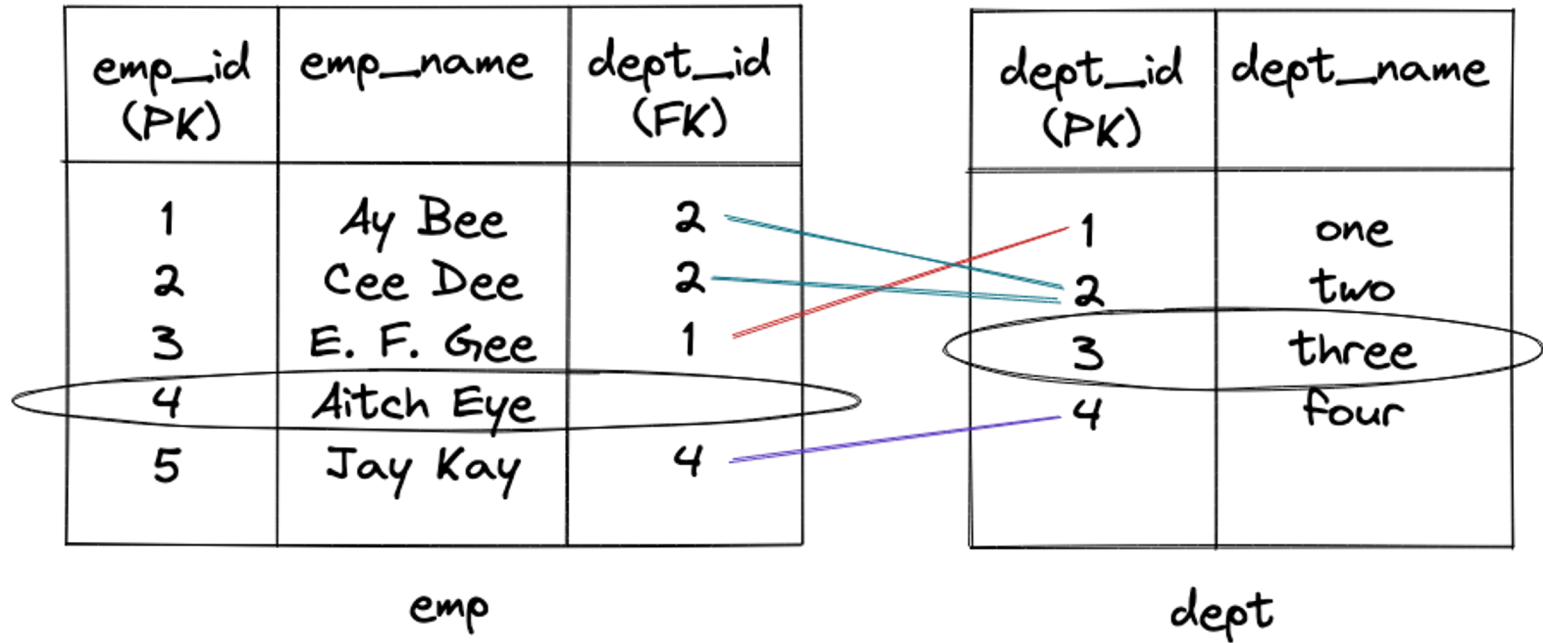
DML Data Manipulation Language SELECT, INSERT, UPDATE, DELETE...

DDL Data Definition Language CREATE, ALTER, DROP, RENAME ...

What is a Join?

What is a Join?

A way to select from multiple tables in one statement



What are the Different Types of Join?

What are the Different Types of Join?


INNER JOIN

```
my_database=> SELECT emp.emp_id, emp.emp_name, dept.dept_name  
my_database-> FROM emp INNER JOIN dept  
my_database-> ON emp.dept_id = dept.dept_id;
```

What are the Different Types of Join?

INNER JOIN or JOIN default join type

```
my_database=> SELECT emp.emp_id, emp.emp_name, dept.dept_name
my_database-> FROM   emp JOIN dept
my_database-> ON     emp.dept_id = dept.dept_id;
 emp_id | emp_name | dept_name
-----+-----+-----
      1 | Ay Bee   | Consulting
      2 | Cee Dee  | Consulting
      3 | E.F. Gee | Sales
      5 | Jay Kay  | HR
(4 rows)
```



What are the Different Types of Join?


LEFT OUTER JOIN

```
my_database=> SELECT emp.emp_id, emp.emp_name, dept.dept_name
my_database-> FROM emp LEFT OUTER JOIN dept
my_database-> ON emp.dept_id = dept.dept_id;
```

What are the Different Types of Join?

LEFT OUTER JOIN or LEFT JOIN

```
my_database=> SELECT emp.emp_id, emp.emp_name, dept.dept_name
my_database-> FROM emp LEFT JOIN dept
my_database-> ON emp.dept_id = dept.dept_id;
 emp_id | emp_name | dept_name
-----+-----+-----
      1 | Ay Bee   | Consulting
      2 | Cee Dee  | Consulting
      3 | E. F. Gee | Sales
      4 | Aitch Eye |
      5 | Jay Kay  | HR
(5 rows)
```



What are the Different Types of Join?

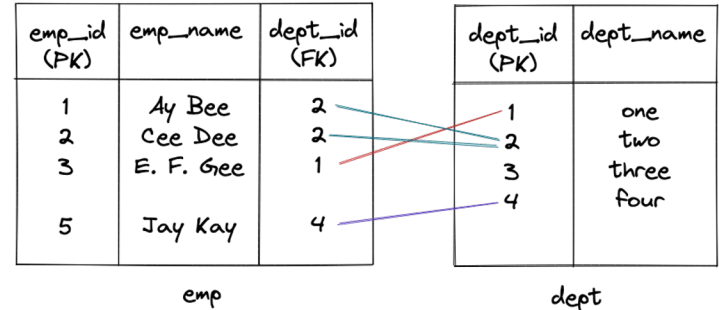
RIGHT OUTER JOIN

```
my_database=> SELECT emp.emp_id, emp.emp_name, dept.dept_name
my_database-> FROM emp RIGHT OUTER JOIN dept
my_database-> ON emp.dept_id = dept.dept_id;
```

What are the Different Types of Join?

RIGHT OUTER JOIN or RIGHT JOIN

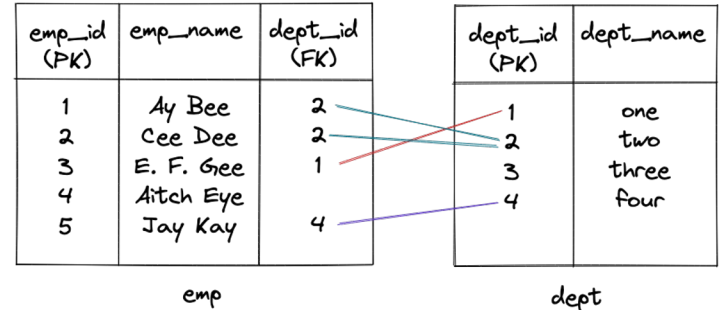
```
my_database=> SELECT emp.emp_id, emp.emp_name, dept.dept_name
my_database-> FROM emp RIGHT JOIN dept
my_database-> ON emp.dept_id = dept.dept_id;
 emp_id | emp_name | dept_name
-----+-----+-----
      1 | Ay Bee   | Consulting
      2 | Cee Dee  | Consulting
      3 | E. F. Gee | Sales
      5 | Jay Kay  | HR
          |          | Product
(5 rows)
```



What are the Different Types of Join?

FULL OUTER JOIN

```
my_database=> SELECT emp.emp_id, emp.emp_name, dept.dept_name
my_database-> FROM emp FULL OUTER JOIN dept
my_database-> ON emp.dept_id = dept.dept_id;
 emp_id | emp_name | dept_name
-----+-----+-----
      1 | Ay Bee   | Consulting
      2 | Cee Dee  | Consulting
      3 | E. F. Gee | Sales
      4 | Aitch Eye |
      5 | Jay Kay  | HR
          |          | Product
(6 rows)
```



What is an Execution Plan?

What is an Execution Plan?

The steps that Postgres takes to execute a SQL statement

- Access path: full table scan, index lookup
- Join algorithm: nested loop, hash, merge
- Estimated # rows
- Estimated cost

How do I Generate an Execution Plan?

How do I Generate an Execution Plan?

<https://www.postgresql.org/docs/current/using-explain.html>

```
my_database=> EXPLAIN SELECT * FROM emp_name_view;
```

How do I Generate an Execution Plan?

<https://www.postgresql.org/docs/current/using-explain.html>

```
my_database=> EXPLAIN SELECT * FROM emp_name_view;
               QUERY PLAN
-----
Sort  (cost=1.11..1.12 rows=5 width=32)
  Sort Key: emp.emp_name
    -> Seq Scan on emp  (cost=0.00..1.05 rows=5 width=32)
(3 rows)
```

How do I Generate an Execution Plan?

<https://www.postgresql.org/docs/current/using-explain.html>

```
my_database=> EXPLAIN ANALYZE SELECT * FROM emp_name_view;
```

How do I Generate an Execution Plan?

<https://www.postgresql.org/docs/current/using-explain.html>

```
my_database=> EXPLAIN ANALYZE SELECT * FROM emp_name_view;
                QUERY PLAN
-----
Sort  (cost=83.37..86.37 rows=1200 width=32) (actual time=0.225..0.226 rows=5 loops=1)
  Sort Key: emp.emp_name
  Sort Method: quicksort  Memory: 25kB
  -> Seq Scan on emp  (cost=0.00..22.00 rows=1200 width=32) (actual time=0.009..0.010 rows=5 loops=1)
Planning Time: 0.310 ms
Execution Time: 0.434 ms
(6 rows)
```

Agenda

- Database Architecture
- Users and Roles
- Database Objects
- Database Connections
- Database Operations and Transactions
- **WAL**
- Documentation

What is WAL?

What is WAL?

<https://www.postgresql.org/docs/current/wal-intro.html>

- Write-ahead log
- Details of changes executed against the Postgres cluster
- Allows crash-recovery, online backups and restores
- WAL files can be archived to allow point in time recovery (PITR)
- **DON'T DELETE THEM!**

Where are my WAL Files?

```
postgres@my_vm$ ls -ltr $PGDATA/pg_wal
```


Where are my WAL Files?

```
postgres@my_vm$ ls -ltr $PGDATA/pg_wal
total 32768
drwx----- . 2 postgres postgres      6 Nov  9 11:00 archive_status
-rw----- . 1 postgres postgres 16777216 Nov 14 08:41 000000010000000000000001
-rw----- . 1 postgres postgres 16777216 Nov 14 08:41 000000010000000000000002
-rw----- . 1 postgres postgres 16777216 Nov 14 08:46 000000010000000000000003
```

Agenda

- Database Architecture
- Users and Roles
- Database Objects
- Database Connections
- Database Operations and Transactions
- **Documentation**

Where is the PostgreSQL Documentation?

PostgreSQL Documentation

<https://www.postgresql.org/docs>

Current online version: <https://www.postgresql.org/docs/current/index.html>

Where else can I get information?

Slack Channel

<https://postgres-slack.herokuapp.com/>

- Over 100 channels
- 20k members

Mailing Lists

<https://lists.postgresql.org/>

- Linked to PostgreSQL Community account
- Many different lists

Congratulations!



Thank You!

Karen Jex | @karenhjex | karen.jex@crunchydata.com

