

## Engine & Connection URLs

One `create_engine` call builds a pooled gateway to any database from a URL.

PURPOSE	COMMAND	RESULT
Build an engine (SQLite file).	<code>engine = create_engine("sqlite:///app.db")</code>	
Connect to PostgreSQL.	<code>create_engine("postgresql+psycopg://user:pw@localhost:5432/mydb")</code>	
In-memory DB for tests.	<code>create_engine("sqlite://")</code>	
See the SQL it emits.	<code>create_engine(url, echo=True)</code>	
Build a URL safely.	<code>URL.create("postgresql+psycopg", username="u", ...)</code>	
Open a connection block.	<code>with engine.connect() as conn: ...</code>	

`create_engine(url)` builds one lazy, pooled Engine; open short with `engine.connect()` blocks that return the connection on exit

## Raw SQL with text()

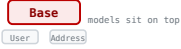





Run any SQL string with bound parameters, never string formatting.

PURPOSE	COMMAND	RESULT
Execute a query.	<code>conn.execute(text("SELECT * FROM users"))</code>	
Bind parameters (no f-strings).	<code>conn.execute(text("SELECT * FROM users WHERE id =:id"), {"id": 7})</code>	
One scalar value.	<code>conn.execute(text("SELECT count(*) FROM users")).scalar_one()</code>	
Rows as dicts.	<code>conn.execute(stmt).mappings().all()</code>	
Insert many rows at once.	<code>conn.execute(text("INSERT INTO t VALUES (:x)"), [{"x": 1}, {"x": 2}])</code>	
Commit the transaction.	<code>with engine.begin() as conn: ...</code>	

Wrap raw SQL in `text()` and pass `:name` placeholders with a `params` dict; `engine.begin()` commits on clean exit, rolls back on error

## ORM Models




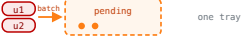
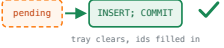

Map a Python class to a table with Mapped and mapped\_column.

PURPOSE	COMMAND	RESULT
Declare a base class.	<pre>class Base(DeclarativeBase): pass</pre>	
Map a class to a table.	<pre>class User(Base):     __tablename__ = "user_account"</pre>	
Typed primary-key column.	<pre>id: Mapped[int] = mapped_column(     primary_key=True)</pre>	
Typed column with length.	<pre>name: Mapped[str] = mapped_column(String(30))</pre>	
Nullable column (Optional).	<pre>fullname: Mapped[str   None]</pre>	
Create all tables.	<pre>Base.metadata.create_all(engine)</pre>	

2.0 style: subclass DeclarativeBase, set \_\_tablename\_\_, declare each column as Mapped[...] + mapped\_column(); the Python type drives the SQL type

## Session & Unit of Work

A Session collects changes; commit flushes them in one transaction.

PURPOSE	COMMAND	RESULT
Open a session.	<pre>with Session(engine) as session:</pre>	
A reusable session factory.	<pre>Session = sessionmaker(engine)</pre>	
Stage one new object.	<pre>session.add(user)</pre>	
Stage several at once.	<pre>session.add_all([u1, u2])</pre>	
Commit the unit of work.	<pre>session.commit()</pre>	
Undo on error.	<pre>session.rollback()</pre>	

add / add\_all stage objects as pending; nothing hits the database until commit() flushes the whole batch in one transaction

## Querying with select()

Build a SELECT as an object, then run it through the Session.

PURPOSE	COMMAND	RESULT
Build a SELECT.	<code>stmt = select(User).where(User.name == "ada")</code>	
Get model objects back.	<code>session.scalars(stmt).all()</code>	
Exactly one (or error).	<code>session.scalars(stmt).one()</code>	
Order and limit.	<code>select(User).order_by(User.name).limit(10)</code>	
Look up by primary key.	<code>session.get(User, 7)</code>	
Aggregate / count.	<code>session.scalar(select(func.count()).select_from(User))</code>	

`select()` builds the statement; `scalars()` returns mapped objects, `get()` hits the identity-map cache, `one()` raises if not exactly one

## Insert, Update, Delete

Mutate through the Session (ORM) or with Core DML statements.

PURPOSE	COMMAND	RESULT
Insert (ORM).	<code>session.add(User(name="ada")); .commit()</code>	
Update (ORM, just assign).	<code>user.fullname = "Ada L."; session.commit()</code>	
Delete (ORM).	<code>session.delete(user); session.commit()</code>	
Bulk update (Core).	<code>update(User).where(User.id == 7).values(name="x")</code>	
Bulk delete (Core).	<code>delete(User).where(User.name == "spam")</code>	
Run the Core statement.	<code>session.execute(stmt); session.commit()</code>	

ORM: add / assign / delete then commit; Core update()/delete() compile to one bulk statement, run via session.execute then commit

## Relationships & Joins

Link tables with relationship(); traverse or join across them.

PURPOSE	COMMAND	RESULT
Foreign-key column.	<pre>user_id: Mapped[int] = mapped_column(ForeignKey("user_account.id"))</pre>	
One-to-many relationship.	<pre>addresses: Mapped[list["Address"]] = relationship(back_populates="user")</pre>	
Many-to-one back-reference.	<pre>user: Mapped["User"] = relationship(back_populates="addresses")</pre>	
Traverse in Python.	<pre>user.addresses.append(Address(email="a@x.com"))</pre>	
JOIN two tables.	<pre>select(User.name, Address.email).join(User.addresses)</pre>	
Avoid the N+1 problem.	<pre>select(User).options(selectinload(User.addresses))</pre>	

ForeignKey ties rows in the DB; relationship() exposes the link in Python; join() and selectinload() pull related data without the N+1 storm

## Read straight into pandas

Hand pandas a SQLAlchemy connection and a statement; get a DataFrame.

PURPOSE	COMMAND	RESULT
Read a SQL string.	<pre>pd.read_sql(text("SELECT * FROM ..."), conn)</pre>	
Read a select() statement.	<pre>pd.read_sql_query(select(User.id, User.name), conn)</pre>	
Read a whole table.	<pre>pd.read_sql_table("user_account", engine)</pre>	
Pass params safely.	<pre>pd.read_sql(text("... WHERE id = :id"), conn, params={"id": 7})</pre>	
Read in chunks (big tables).	<pre>pd.read_sql(stmt, conn, chunksize=10_000)</pre>	
Write a DataFrame back.	<pre>df.to_sql("scratch", engine, if_exists="replace")</pre>	

SQLAlchemy owns the connection and SQL, pandas owns the table: read\_sql/read\_sql\_query/read\_sql\_table in, to\_sql out