# Django SQL Explorer

**May 25, 2021**

# Contents:

# SQL Explorer

SQL Explorer aims to make the flow of data between people fast, simple, and confusion-free. It is a Django-based application that you can add to an existing Django site, or use as a standalone business intelligence tool.

Quickly write and share SQL queries in a simple, usable SQL editor, preview the results in the browser, share links, download CSV, JSON, or Excel files (and even expose queries as API endpoints, if desired), and keep the information flowing!

Comes with support for multiple connections, to many different SQL database types, a schema explorer, query history (e.g. lightweight version control), a basic security model, in-browser pivot tables, and more.

SQL Explorer values simplicity, intuitive use, unobtrusiveness, stability, and the principle of least surprise.

SQL Explorer is inspired by any number of great query and reporting tools out there.

The original idea came from Stack Exchange's Data Explorer, but also owes credit to similar projects like Redash and Blazer.

Sql Explorer is MIT licensed, and pull requests are welcome.

**A view of a query**

## Viewing all queries



## Quick access to DB schema info

**Snapshot query results to S3 & download as csv**

## 1.1 Features

### 1.1.1 Security

- Let's not kid ourselves - this tool is all about giving people access to running SQL in production. So if that makes you nervous (**and it should**) - you've been warned. Explorer makes an effort to not allow terrible things to happen, but be careful! It's recommended you setup read-only roles for each of your database connections and only use these particular connections for your queries through the `EXPLORER_CONNECTIONS` setting.

- Explorer supports two different permission checks for users of the tool. Users passing the `EXPLORER_PERMISSION_CHANGE` test can create, edit, delete, and execute queries. Users who do not pass this test but pass the `EXPLORER_PERMISSION_VIEW` test can only execute queries. Other users cannot access any part of Explorer. Both permission groups are set to is_staff by default and can be overridden in your settings file.

- Enforces a SQL blacklist so destructive queries don't get executed (delete, drop, alter, update etc). This is not bulletproof and it's recommended that you instead configure a read-only database role, but when not possible the blacklist provides reasonable protection.

### 1.1.2 Easy to get started

- Built on Django's ORM, so works with Postgresql, Mysql, and Sqlite. And, between you and me, it works fine on RedShift as well.

- Small number of dependencies.

- Just want to get in and write some ad-hoc queries? Go nuts with the Playground area.

### 1.1.3 Snapshots

- Tick the 'snapshot' box on a query, and Explorer will upload a .csv snapshot of the query results to S3. Configure the snapshot frequency via a celery cron task, e.g. for daily at 1am:

```
'explorer.tasks.snapshot_queries': {
    'task': 'explorer.tasks.snapshot_queries',
    'schedule': crontab(hour=1, minute=0)
}
```

- Requires celery, obviously. Also uses djcelery and tinys3. All of these deps are optional and can be installed with `pip install -r optional-requirements.txt`

- The checkbox for opting a query into a snapshot is ALL THE WAY on the bottom of the query view (underneath the results table).

- You must also have the setting `EXPLORER_TASKS_ENABLED` enabled.

### 1.1.4 Email query results

- Click the email icon in the query listing view, enter an email address, and the query results (zipped .csv) will be sent to you asynchronously. Very handy for long-running queries.

### 1.1.5 Parameterized Queries

- Use $$foo$$ in your queries and Explorer will build a UI to fill out parameters. When viewing a query like `SELECT * FROM table WHERE id=$$id$$`, Explorer will generate UI for the `id` parameter.

- Parameters are stashed in the URL, so you can share links to parameterized queries with colleagues

- Use `$$paramName:defaultValue$$` to provide default values for the parameters.

### 1.1.6 Schema Helper

- `/explorer/schema/<connection-alias>` renders a list of your table and column names + types that you can refer to while writing queries. Apps can be excluded from this list so users aren't bogged down with tons of irrelevant tables. See settings documentation below for details.

- This is available quickly as a sidebar helper while composing queries (see screenshot)

- Quick search for the tables you are looking for. Just start typing!

- Explorer uses Django DB introspection to generate the schema. This can sometimes be slow, as it issues a separate query for each table it introspects. Therefore, once generated, Explorer caches the schema information. There is also the option to generate the schema information asyncronously, via Celery. To enable this, make sure Celery is installed and configured, and set `EXPLORER_ENABLE_TASKS` and `EXPLORER_ASYNC_SCHEMA` to `True`.

### 1.1.7 Template Columns

- Let's say you have a query like 'select id, email from user' and you'd like to quickly drill through to the profile page for each user in the result. You can create a "template" column to do just that.

- Just set up a template column in your settings file:

```
EXPLORER_TRANSFORMS = [('user', '<a href="https://yoursite.com/profile/{0}/
">{0}</a>')]
```

- And change your query to `SELECT id AS "user", email FROM user`. Explorer will match the "user" column alias to the transform and merge each cell in that column into the template string. Cool!

- Note you *must* set `EXPLORER_UNSAFE_RENDERING` to `True` if you want to see rendered HTML (vs string literals) in the output. And be aware of the implications of enabling that setting.

### 1.1.8 Pivot Table

- Go to the Pivot tab on query results to use the in-browser pivot functionality (provided by Pivottable JS).

- Hit the link icon on the top right to get a URL to recreate the exact pivot setup to share with colleagues.

### 1.1.9 Query Logs

- Explorer will save a snapshot of every query you execute so you can recover lost ad-hoc queries, and see what you've been querying.

- This also serves as cheap-and-dirty versioning of Queries, and provides the 'run count' property and average duration in milliseconds, by aggregating the logs.

- You can also quickly share playground queries by copying the link to the playground's query log record – look on the top right of the sql editor for the link icon.

- If Explorer gets a lot of use, the logs can get beefy. explorer.tasks contains the 'truncate_querylogs' task that will remove log entries older than <days> (30 days and older in the example below).

```
'explorer.tasks.truncate_querylogs': {
    'task': 'explorer.tasks.truncate_querylogs',
    'schedule': crontab(hour=1, minute=0),
    'kwargs': {'days': 30}
}
```

### 1.1.10 Multiple Connections

- Have data in more than one database? No problemo. Just set up multiple Django database connections, register them with Explorer, and you can write, save, and view queries against all of your different data sources. Compatible with any database support by Django. Note that the target database does *not* have to contain any Django schema, or be related to Django in any way. See connections.py for more documentation on multi-connection setup.

### 1.1.11 Power tips

- On the query listing page, focus gets set to a search box so you can just navigate to /explorer and start typing the name of your query to find it.

- Quick search also works after hitting "Show Schema" on a query view.

- Command+Enter and Ctrl+Enter will execute a query when typing in the SQL editor area.

- Hit the "Format" button to format and clean up your SQL (this is non-validating – just formatting).

- Use the Query Logs feature to share one-time queries that aren't worth creating a persistent query for. Just run your SQL in the playground, then navigate to `/logs` and share the link (e.g. `/explorer/play/?querylog_id=2428`)

- Click the 'history' link towards the top-right of a saved query to filter the logs down to changes to just that query.

- If you need to download a query as something other than csv but don't want to globally change delimiters via `settings.EXPLORER_CSV_DELIMETER`, you can use `/query/download?delim=|` to get a pipe (or whatever) delimited file. For a tab-delimited file, use `delim=tab`. Note that the file extension will remain .csv

- If a query is taking a long time to run (perhaps timing out) and you want to get in there to optimize it, go to `/query/123/?show=0`. You'll see the normal query detail page, but the query won't execute.

- Set env vars for `EXPLORER_TOKEN_AUTH_ENABLED=TRUE` and `EXPLORER_TOKEN=<SOME TOKEN>` and you have an instant data API. Just:

  ```
  curl --header "X-API-TOKEN: <TOKEN>" https://www.your-site.com/explorer/
  <QUERY_ID>/stream?format=csv
  ```

  You can also pass the token with a query parameter like this:

  ```
  curl https://www.your-site.com/explorer/<QUERY_ID>/stream?
  format=csv&token=<TOKEN>
  ```

## 1.2 Change Log

This document records all notable changes to django-sql-explorer. This project adheres to Semantic Versioning.

### 1.2.1 2.1.0 (2021-01-13)

- **BREAKING CHANGE**: `request` object now passed to `EXPLORER_PERMISSION_CHANGE` and `EXPLORER_PERMISSION_VIEW` (#417 to fix #396)

Major Changes

- #413: Static assets now served directly from the application, not CDN. (#418 also)
- #414: Better blacklist checking - Fix #371 and #412
- #415: Fix for MySQL following change for Oracle in #337

Minor Changes

- #370: Get the CSRF cookie name from django instead of a hardcoded value
- #410 and #416: Sphinx docs
- #420: Formatting change in templates
- #424: Collapsable SQL panel
- #425: Ensure a *Query* object contains SQL

### 1.2.2 2.0.0 (2020-10-09)

- **BREAKING CHANGE**: #403: Dropping support for EOL Python 2.7 and 3.5

Major Changes

- #404: Add support for Django 3.1 and drop support for (EOL) <2.2
- #408: Refactored the application, updating the URLs to use path and the views into a module

Minor Changes

- #334: Django 2.1 support
- #337: Fix Oracle query failure caused by *TextField* in a group by clause
- #345: Added (some) Chinese translation
- #366: Changes to Travis django versions
- #372: Run queries as atomic requests
- #382: Django 2.2 support
- #383: Typo in the README
- #385: Removed deprecated *render_to_response* usage
- #386: Bump minimum django version to 2.2
- #387: Django 3 support
- #390: README formatting changes
- #393: Added option to install *XlsxWriter* as an extra package

- #397: Bump patch version of django 2.2
- #406: Show some love to the README
- Fix #341: PYC files excluded from build

### 1.2.3 1.1.3 (2019-09-23)

- #347: URL-friendly parameter encoding
- #354: Updating dependency reference for Python 3 compatibility
- #357: Include database views in list of tables
- #359: Fix unicode issue when generating migration with py2 or py3
- #363: Do not use "message" attribute on exception
- #368: Update EXPLORER_SCHEMA_EXCLUDE_TABLE_PREFIXES

Minor Changes

- release checklist included in repo
- readme updated with new screenshots
- python dependencies/optional-dependencies updated to latest (six, xlsxwriter, factory-boy, sqlparse)

### 1.2.4 1.1.2 (2018-08-14)

- Fix #269
- Fix bug when deleting query
- Fix bug when invalid characters present in Excel worksheet name

Major Changes

- Django 2.0 compatibility
- Improved interface to database connection management

Minor Changes

- Documentation updates
- Load images over same protocol as originating page

### 1.2.5 1.1.1 (2017-03-21)

- Fix #288 (incorrect import)

### 1.2.6 1.1.0 (2017-03-19)

- **BREAKING CHANGE**: EXPLORER_DATA_EXPORTERS setting is now a list of tuples instead of a dictionary. This only affects you if you have customized this setting. This was to preserve ordering of the export buttons in the UI.
- **BREAKING CHANGE**: Values from the database are now escaped by default. Disable this behavior (enabling potential XSS attacks) with the EXPLORER_UNSAFE_RENDERING setting.

Major Changes

- Django 1.10 and 2.0 compatibility

- Theming & visual updates

- PDF export

- Query-param based authentication ([#254](#254))

- Schema built via SQL querying rather than Django app/model introspection. Paves the way for the tool to be pointed at any DB, not just Django DBs

Minor Changes

- Switched from TinyS3 to Boto (will switch to Boto3 in next release)

- Optionally show row numbers in results preview pane

- Full-screen view (icon on top-right of preview pane)

- Moved 'open in playground' to icon on top-right on SQL editor

- Save-only option (does not execute query)

- Show the time that the query was rendered (useful if you've had a tab open a while)

## 1.2.7 1.0.0 (2016-06-16)

- **BREAKING CHANGE**: Dropped support for Python 2.6. See `.travis.yml` for test matrix.

- **BREAKING CHANGE**: The 'export' methods have all changed. Those these weren't originally designed to be external APIs, folks have written consuming code that directly called export code.

  If you had code that looked like:

  ```
  explorer.utils.csv_report(query)
  ```

  You will now need to do something like:

  ```
  explorer.exporters.get_exporter_class('csv')(query).
  get_file_output()
  ```

- There is a new export system! v1 is shipping with support for CSV, JSON, and Excel (xlsx). The availablility of these can be configured via the EXPLORER_DATA_EXPORTERS setting. * *Note* that for Excel export to work, you will need to install `xlsxwriter` from `optional-requirements.txt`.

- Introduced Query History link. Find it towards the top right of a saved query.

- Front end performance improvements and library upgrades.

- Allow non-admins with permission to log into explorer.

- Added a proper test_project for an easier entry-point for contributors, or folks who want to kick the tires.

- Loads of little bugfixes.

## 1.2.8 0.9.2 (2016-02-02)

- Fixed readme issue (.1) and `setup.py` issue (.2)

### 1.2.9 0.9.1 (2016-02-01)

Major changes

- Dropped support for Django 1.6, added support for Django 1.9. See .travis.yml for test matrix.

- Dropped charted.js & visualization because it didn't work well.

- Client-side pivot tables with pivot.js. This is ridiculously cool!

Minor (but awesome!) changes

- Cmd-/ to comment/uncomment a block of SQL

- Quick 'shortcut' links to the corresponding querylog to more quickly share results. Look at the top-right of the editor. Also works for playground!

- Prompt for unsaved changes before navigating away

- Support for default parameter values via $$paramName:defaultValue$$

- Optional Celery task for truncating query logs as entries build up

- Display historical average query runtime

- Increased default number of rows from 100 to 1000

- Increased SQL editor size (5 additional visible lines)

- CSS cleanup and streamlining (making better use of foundation)

- Various bugfixes (blacklist not enforced on playground being the big one)

- Upgraded front-end libraries

- Hide Celery-based features if tasks not enabled.

### 1.2.10 0.8.0 (2015-10-21)

- Snapshots! Dump the csv results of a query to S3 on a regular schedule. More details in readme.rst under 'features'.

- Async queries + email! If you have a query that takes a long time to run, execute it in the background and Explorer will send you an email with the results when they are ready. More details in readme.rst

- Run counts! Explorer inspects the query log to see how many times a query has been executed.

- Column Statistics! Click the . . . on top of numeric columns in the results pane to see min, max, avg, sum, count, and missing values.

- Python 3! * Django 1.9!

- Delimiters! Export with delimiters other than commas.

- Listings respect permissions! If you've given permission to queries to non-admins, they will see only those queries on the listing page.

### 1.2.11 0.7.0 (2015-02-18)

- Added search functionality to schema view and explorer view (using list.js).

- Python 2.6 compatibility.

- Basic charts via charted (from Medium via charted.co).

- SQL formatting function.

- Token authentication to retrieve csv version of queries.

- Fixed south_migrations packaging issue.

- Refactored front-end and pulled CSS and JS into dedicated files.

### 1.2.12  0.6.0 (2014-11-05)

- Introduced Django 1.7 migrations. See readme.rst for info on how to run South migrations if you are not on Django 1.7 yet.

- Upgraded front-end libraries to latest versions.

- Added ability to grant selected users view permissions on selected queries via the `EXPLORER_USER_QUERY_VIEWS` parameter

- Example usage: `EXPLORER_USER_QUERY_VIEWS = {1:  [3,4], 2:[3]}`

- This would grant user with PK 1 read-only access to query with PK=3 and PK=4 and user 2 access to query 3.

- Bugfixes

- Navigating to an explorer URL without the trailing slash now redirects to the intended page (e.g. `/logs ->
/logs/`)

- Downloading a .csv and subsequently re-executing a query via a keyboard shortcut (cmd+enter) would re-submit the form and re-download the .csv. It now correctly just refreshes the query.

- Django 1.7 compatibility fix

### 1.2.13  0.5.1 (2014-09-02)

Bugfixes

- Created_by_user not getting saved correctly

- Content-disposition .csv issue

- Issue with queries ending in `...like  '%...` clauses

- Change the way customer user model is referenced

- Pseudo-folders for queries. Use "Foo * Ba1", "Foo * Bar2" for query names and the UI will build a little "Foo" pseudofolder for you in the query list.

### 1.2.14  0.5.0 (2014-06-06)

- Query logs! Accessible via `explorer/logs/`. You can look at previously executed queries (so you don't, for instance, lose that playground query you were working, or have to worry about mucking up a recorded query). It's quite usable now, and could be used for versioning and reverts in the future. It can be accessed at `explorer/logs/`

- Actually captures the creator of the query via a ForeignKey relation, instead of just using a Char field.

- Re-introduced type information in the schema helpers.

- Proper relative URL handling after downloading a query as CSV.

- Users with view permissions can use query parameters. There is potential for SQL injection here. I think about the permissions as being about preventing users from borking up queries, not preventing them from viewing data. You've been warned.

- Refactored params handling for extra safety in multi-threaded environments.

### 1.2.15 0.4.1 (2014-02-24)

- Renaming template blocks to prevent conflicts

### 1.2.16 0.4 (2014-02-14 *Happy Valentine's Day!*)

- Templatized columns for easy linking
- Additional security config options for splitting create vs. view permissions
- Show many-to-many relation tables in schema helper

### 1.2.17 0.3 (2014-01-25)

- Query execution time shown in query preview
- Schema helper available as a sidebar in the query views
- Better defaults for sql blacklist
- Minor UI bug fixes

### 1.2.18 0.2 (2014-01-05)

- Support for parameters
- UI Tweaks
- Test coverage

### 1.2.19 0.1.1 (2013-12-31)

Bug Fixes

- Proper SQL blacklist checks
- Downloading CSV from playground

### 1.2.20 0.1 (2013-12-29)

Initial Release

## 1.3 Install

Requires Python 3.6 or higher. Requires Django 2.2 or higher.

Set up a Django project with the following:

```
$ pip install django
$ django-admin startproject project
```

More information in the [django tutorial](#).

Install with pip from pypi:

```
pip install django-sql-explorer
```

If you would also like to support downloading Excel files install with the dependency using:

```
pip install django-sql-explorer[xls]
```

Add to your `INSTALLED_APPS`, located in the `settings.py` file in your project folder:

```
INSTALLED_APPS = (
    ...,
    'explorer',
    ...
)
```

Add the following to your urls.py (all Explorer URLs are restricted via the `EXPLORER_PERMISSION_VIEW` and `EXPLORER_PERMISSION_CHANGE` settings. See Settings section below for further documentation.):

```
from django.urls import path

urlpatterns = [
    ...
    path('explorer/', include('explorer.urls')),
    ...
]
```

Configure your settings to something like:

```
EXPLORER_CONNECTIONS = { 'Default': 'readonly' }
EXPLORER_DEFAULT_CONNECTION = 'readonly'
```

The first setting lists the connections you want to allow Explorer to use. The keys of the connections dictionary are friendly names to show Explorer users, and the values are the actual database aliases used in `settings.` `DATABASES`. It is highly recommended to setup read-only roles in your database, add them in your project's `DATABASES` setting and use these read-only cconnections in the `EXPLORER_CONNECTIONS`.

If you want to quickly use django-sql-explorer with the existing default connection **and know what you are doing** (or you are on development), you can use the following settings:

```
EXPLORER_CONNECTIONS = { 'Default': 'default' }
EXPLORER_DEFAULT_CONNECTION = 'default'
```

Finally, run migrate to create the tables:

```
python manage.py migrate
```

You can now browse to [https://yoursite/explorer/](https://yoursite/explorer/) and get exploring!

There are a handful of features (snapshots, emailing queries) that rely on Celery and the dependencies in optional-requirements.txt. If you have Celery installed, set `EXPLORER_TASKS_ENABLED=True` in your settings.py to enable these features.

## 1.4 Dependencies

An effort has been made to keep the number of dependencies to a minimum.

### 1.4.1 Python

| Name | Version | License |
|------|---------|---------|
| sqlparse | 0.4.0 | BSD |

- sqlparse is used for SQL formatting

**Python - Optional Dependencies**

| Name | Version | License |
|------|---------|---------|
| celery | >=3.1,<4 | BSD |
| django-celery | >=3.3.1 | BSD |
| Factory Boy | >=3.1.0 | MIT |
| xlsxwriter | >=1.3.6 | BSD |
| boto | >=2.49 | MIT |

- Factory Boy is required for tests

- celery is required for the 'email' feature, and for snapshots

- boto is required for snapshots

- xlsxwriter is required for Excel export (csv still works fine without it)

JavaScript

| Name | Version | License |
|------|---------|---------|
| Twitter Boostrap | 3.3.6 | MIT |
| jQuery | 2.1.4 | MIT |
| jQuery Cookie | 1.4.1 | MIT |
| jQuery UI | 1.11.4 | MIT |
| Underscore | 1.7.0 | MIT |
| Codemirror | 5.15.2 | MIT |
| floatThead | 1.4.0 | MIT |
| list.js | 1.2.0 | MIT |
| pivottable.js | 2.0.2 | MIT |

- All are served locally, with jQuery UI being a custom build.

pivottable.js relies on jQuery UI but only for the `Sortable` method.

### 1.4.2 Tests

Factory Boy is needed if you'd like to run the tests, which can you do easily:

```
python manage.py test
```

and with coverage:

```
coverage run --source='.' manage.py test
```

then:

```
coverage report
```

. . . 97%! Huzzah!

### 1.4.3 Running Locally

There is also a test_project that you can use to kick the tires. Just create a new virtualenv, cd into `test_project` and run `start.sh` (or walk through the steps yourself) to get a test instance of the app up and running.

# 1.5 Settings

| Setting | Description | Default |
|---|---|---|
| EX-PLORER_SQL_BLACKLIST | Disallowed words in SQL queries to prevent destructive actions. | ('ALTER', 'RENAME ', 'DROP', 'TRUNCATE', 'IN-SERT INTO', 'UPDATE', 'REPLACE', 'DELETE', 'ALTER', 'CREATE TABLE', 'SCHEMA', 'GRANT', 'OWNER TO') |
| EX-PLORER_SQL_WHITELIST | These phrases are allowed, even though a part of the phrase appears in the blacklist. | ('CREATED', 'UPDATED', 'DELETED','REGEXP_REPLACE') |
| EX-PLORER_DEFAULT_ROWS | The number of rows to show by default in the preview pane. | 1000 |
| EX-PLORER_SCHEMA_INCLUDE_TABLE_PREFIXES | If not None, show schema only for tables starting with these prefixes. "Wins" if in conflict with EXCLUDE | None # shows all tables |
| EX-PLORER_SCHEMA_EXCLUDE_TABLE_PREFIXES | Don't show schema for tables starting with these prefixes, in the schema helper. | ('django.contrib.auth', 'django.contrib.contenttypes', 'django.contrib.sessions', 'django.contrib.admin') |
| EX-PLORER_SCHEMA_INCLUDE_VIEWS | Include database views | False |
| EX-PLORER_ASYNC_SCHEMA | Generate DB schema asynchronously. Requires Celery and EXPLORER_TASKS_ENABLED | False |
| EX-PLORER_DEFAULT_CONNECTION | The name of the Django database connection to use. Ideally set this to a connection with read only permissions | None # Must be set for the app to work, as this is required |
| EX-PLORER_CONNECTIONS | A dictionary of { 'Friendly Name': django_db_alias'}. | {} # At a minimum, should be set to something like { 'Default': 'readonly' } or similar. See connections.py for more documentation. |
| EX-PLORER_PERMISSION_VIEW | Callback to check if the user is allowed to view and execute stored queries | lambda u: u.is_staff |
| EX-PLORER_PERMISSION_CHANGE | Callback to check if the user is allowed to change/delete queries | lambda u: u.is_staff |
| EX-PLORER_TRANSFORMS | List of tuples like [('alias', 'Template for {0}')]. See features section of this doc for more info. | [] |
| EX-PLORER_RECENT_QUERY_COUNT | The number of recent queries to show at the top of the query listing. | 10 |
| EX-PLORER_GET_USER_QUERY_VIEWS | A dict granting view permissions on specific queries of the form {userId:[queryId, ...], ...} | {} |
| EX-PLORER_TOKEN_AUTH_ENABLED | Bool indicating whether token-authenticated requests should be enabled. See "Power Tips", above. | False |
| EX-PLORER_TOKEN | Access token for query results. | "CHANGEME" |
| EX-PLORER_TASKS_ENABLED | Turn on if you want to use the snapshot queries celery task, or email report functionality in tasks.py | False |
| EX-PLORER_S3_ACCESS_KEY | S3 Access Key for snapshot upload | None |
| EX-PLORER_S3_SECRET_KEY | S3 Secret Key for snapshot upload | None |
| EX-PLORER_S3_BUCKET | S3 Bucket for snapshot upload | None |

## 1.6 Indices and tables

- genindex
- modindex
- search