VOLTRON DATA

carahsoft.

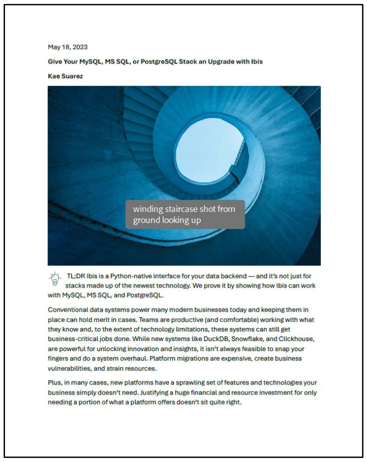
May 18, 2023

Give Your MySQL, MS SQL, or PostgreSQL Stack an Upgrade with Ibis

Kae Suarez

winding staircase shot from ground looking up

TL;DR Ibis is a Python-native interface for your data backend — and it's not just for stacks made up of the newest technology. We prove it by showing how Ibis can work with MySQL, MS SQL, and PostgreSQL.

Conventional data systems power many modern businesses today and keeping them in place can hold merit in cases. Teams are productive (and comfortable) working with what they know and, to the extent of technology limitations, these systems can still get business-critical jobs done. While new systems like DuckDB, Snowflake, and Clickhouse, are powerful for unlocking innovation and insights, it isn't always feasible to snap your fingers and do a system overhaul. Platform migrations are expensive, create business vulnerabilities, and strain resources.

Plus, in many cases, new platforms have a sprawling set of features and technologies your business simply doesn't need. Justifying a huge financial and resource investment for only needing a portion of what a platform offers doesn't sit quite right.

DᶁǾŜ ¸Ƶdzǋ aề{v[Ʃ a{ {v[Ʃ Ƶǋ tƵǎǘƎǋŜ{v[{ǘl·Ŏʅ l·ỿ ¦ǇƎǋl·ŘŜ ǿᶁǘK Lōᶁǎ

Thank you for downloading this Voltron Data resource. Carahsoft is the distributor for Voltron Data's AI and ML solutions available via NASA SEWP V, ITES-SW2, NASPO ValuePoint, and many more contract vehicles.

To learn how to take the next step toward acquiring Voltron Data's solutions, please check out the following resources and information:

For Voltron Data overview:
carah.io/Voltron-Data

For upcoming events:
carah.io/Voltron-Events

For additional resources:
carah.io/Voltron-Resources

For additional Artificial Intelligence:
carah.io/ai-solutions

To set up a meeting:
VoltronData@carahsoft.com

To purchase, check out the contract vehicles available for procurement:
carah.io/procurement

For more information, contact Carahsoft or our reseller partners:
VoltronData@carahsoft.com | 888-606-2770

May 18, 2023

Give Your MySQL, MS SQL, or PostgreSQL Stack an Upgrade with Ibis

Kae Suarez

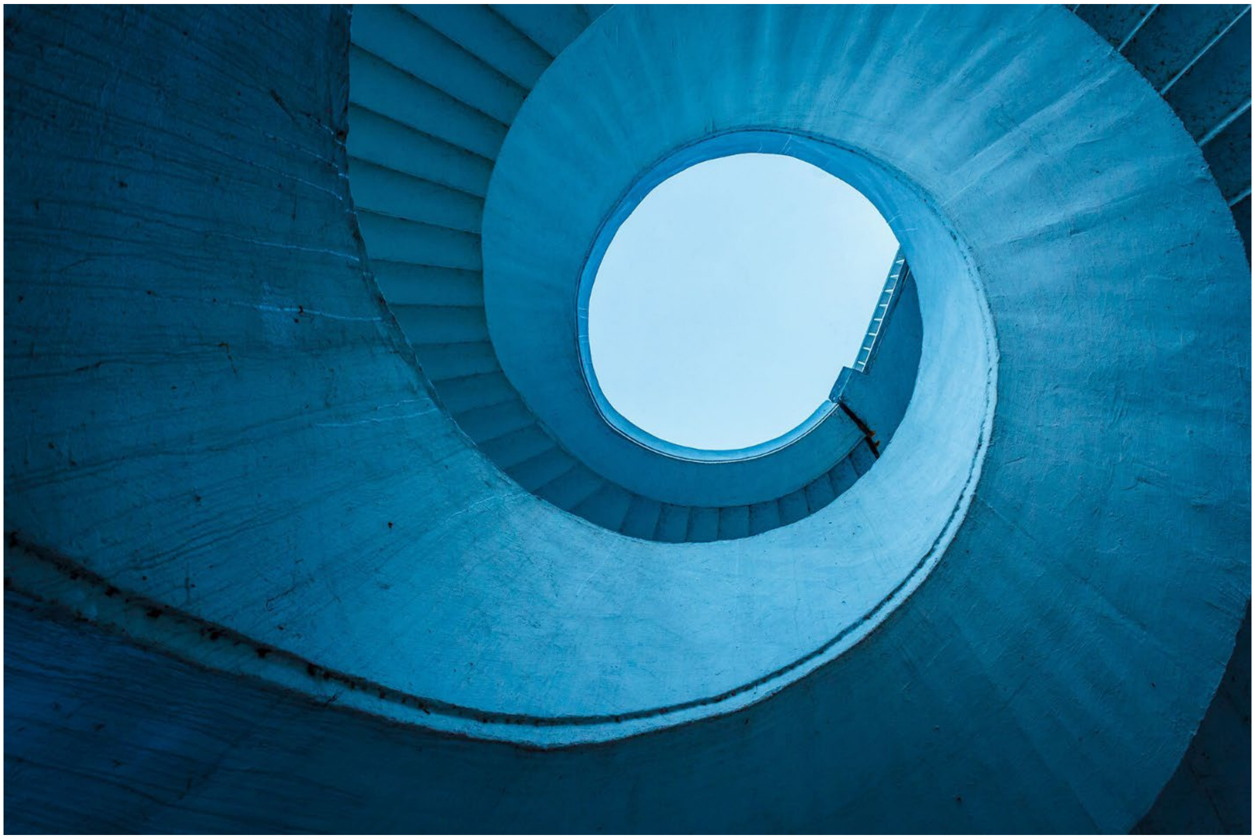

💡 TL;DR Ibis is a Python-native interface for your data backend — and it's not just for stacks made up of the newest technology. We prove it by showing how Ibis can work with MySQL, MS SQL, and PostgreSQL.

Conventional data systems power many modern businesses today and keeping them in place can hold merit in cases. Teams are productive (and comfortable) working with what they know and, to the extent of technology limitations, these systems can still get business-critical jobs done. While new systems like DuckDB, Snowflake, and Clickhouse, are powerful for unlocking innovation and insights, it isn't always feasible to snap your fingers and do a system overhaul. Platform migrations are expensive, create business vulnerabilities, and strain resources.

Plus, in many cases, new platforms have a sprawling set of features and technologies your business simply doesn't need. Justifying a huge financial and resource investment for only needing a portion of what a platform offers doesn't sit quite right.

This is why we push for a strategy of augmenting data systems — carefully considering what you have and then optimizing with open standards.

The Ibis project exemplifies the practice of augmentation because it seamlessly plugs into your system and optimizes it. Ibis is a Python-native interface for data, which brings an easy-to-use and powerful analytics API to your data, along with the rich Python data ecosystem. It isn't just for the new database systems — Ibis is dedicated to supporting the all-time hits, such as PostgreSQL, MySQL, and MS SQL, so you get all it offers without replacing your hard-won stack.

Perhaps best of all, it takes only a couple of minutes to do it.

Setting up with Ibis

Ibis is packaged on PyPi and conda-forge to enable quick and easy connection to the database system of your choice, and we'll showcase that here with PostgreSQL, MySQL, and MS SQL so you can get started using Ibis with your systems right away.

PostgreSQL

In the case that you have PostgreSQL served somewhere, and don't have Ibis, then the first step is installing Ibis, packaged up with dependencies for PostgreSQL. If you're using pip:

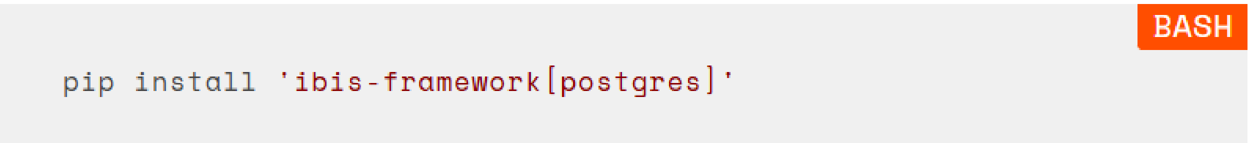

```bash
pip install 'ibis-framework[postgres]'
```

Or mamba/conda:

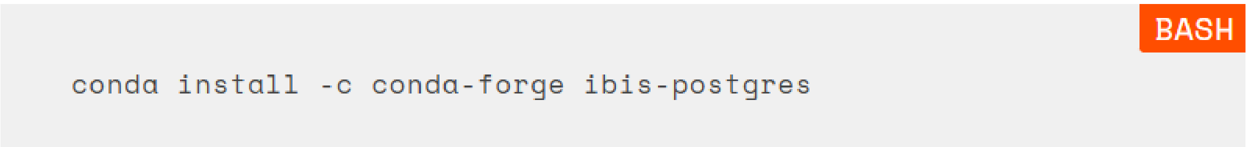

```bash
conda install -c conda-forge ibis-postgres
```

Then, just get into Python, and you can connect to your database:

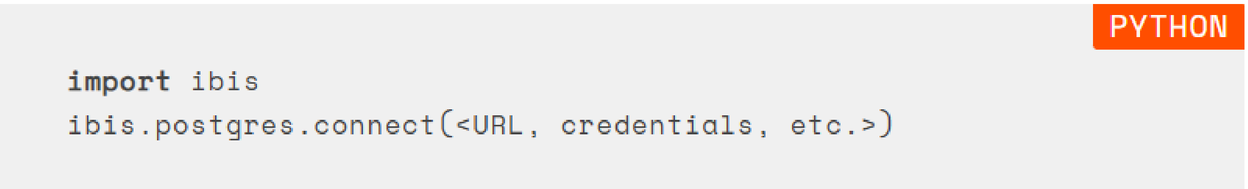

```python
import ibis
ibis.postgres.connect(<URL, credentials, etc.>)
```

Once that’s done, congratulations! Ibis is connected to your PostgreSQL endpoint, and you can start experimenting.

MySQL

But what if you’re using MySQL, instead? Well, then you’ll just have to use the packaging that gets you set up for that. With pip:

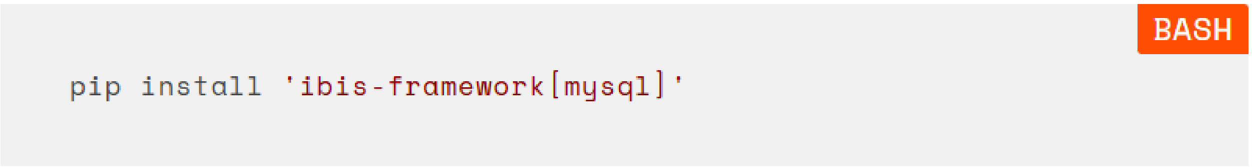

```bash
pip install 'ibis-framework[mysql]'
```

Or mamba/conda:

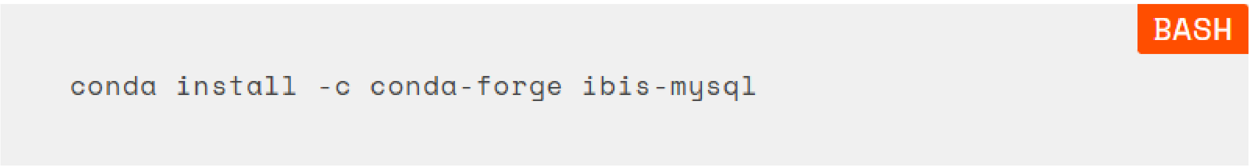

```bash
conda install -c conda-forge ibis-mysql
```

Then, it’s just a matter of connection:

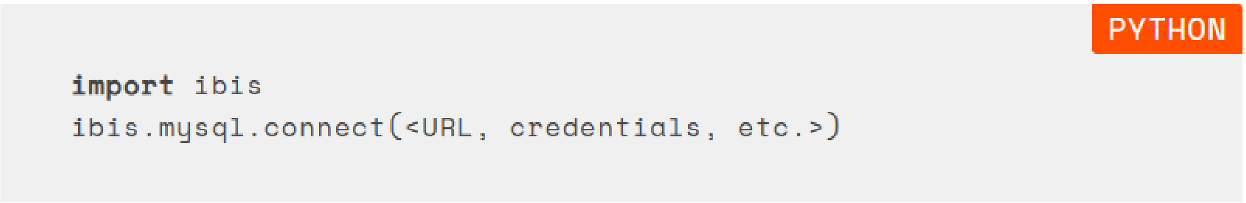

```python
import ibis
ibis.mysql.connect(<URL, credentials, etc.>)
```

That will have you set up with everything you need to start using Ibis and the Python ecosystem for your MySQL database.

Microsoft SQL Server (MS SQL)

If you’re suspicious that this will be repetitive, then you’re catching onto what makes Ibis portable and easy to use across any backend of choice. To get Ibis set up with MS SQL, first, you install. With pip:

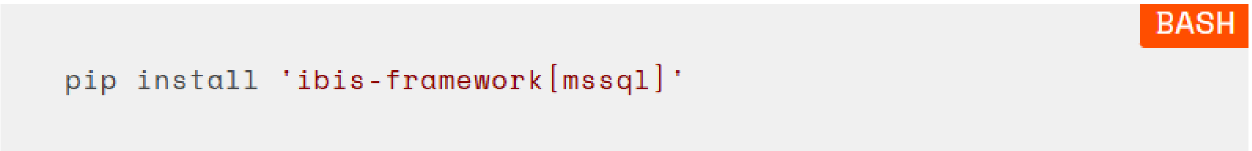

```bash
pip install 'ibis-framework[mssql]'
```

Then, connect:

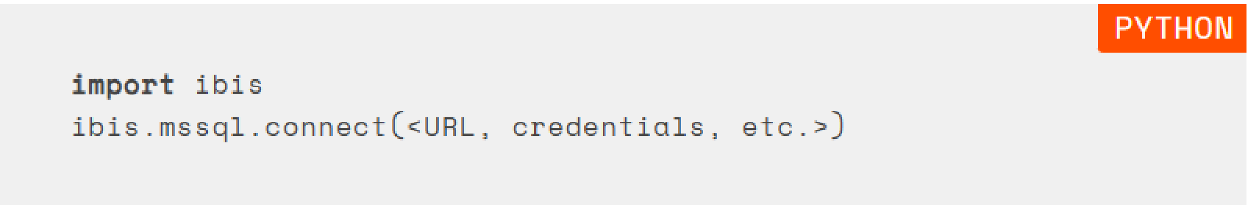

```python
import ibis
ibis.mssql.connect(<URL, credentials, etc.>)
```

With that, you're good to go with MS SQL and Ibis.

Finding New Insights with Ibis

Once you're set up with your backend of choice, it's time to get going. So, what did you actually get out of this? There are several possibilities, but to enumerate a few:

- Ibis can be a brand-new analytics tool, with easy-to-use selectors and aggregations that are portable across backends and generate the SQL for you.
- Ibis can offer a path into the Python ecosystem — and Python is the language for machine learning, so this is an easy entry point to take the data from your legacy systems and use them for modern data analytics.
- Ibis can become your new user-facing query language. Python is a popular language, and being able to use it opens up the field of who already knows, or can learn, to interact with the stack you already have in place.

All of these benefits are made possible by a simple addition to your stack. No replacement, no overhauls, just augmentation.

To get started, refer to Ibis's Getting Started guide, which will help you learn what you can do with your brand new Ibis connection. If you're interested in optimizing your existing data system and want to learn more about Voltron Data can support, contact us.

Photo by Agnieszka Cymbalak