SQL API and Extended Support for BI Tools

cube

Igor Lukanin

Community Code of Conduct

- We want to foster an open and welcoming environment where everyone feels they belong in the Cube.js community.
- The full text of our Code of Conduct is available at https://github.com/cube-js/cube.js/blob/master/CODE_OF_CONDUCT.md
- Any instances of inappropriate/unacceptable behavior can be reported to conduct@cube.dev.

Some quick notes 🗒️

- If you have any questions during the workshop, please feel free to type them in the “Q&A” (on Zoom & YouTube).
- We will be using Cube Cloud for “hands-on” demos.
  - See the workshop preparation guide.
- Recording of the workshop will be posted on the Cube’s YouTube channel.
- All attendees will receive a post-event survey and we’d appreciate your feedback to help us with future events.
  - You will have a chance to select a charity that Cube will make a donation to:
    - Lesbians who tech or oSTEM
- Look for an announcement on the next workshop soon 😀
  - Targeting late July

What we will discuss today

- Recap of the recent updates to the product (headless BI, SQL API)
- **Connecting and using SQL API** + Q & A session
- **Digging deep into SQL API features** + Q & A session
- **Working with BI tools and data notebooks** + Q & A session

Headless BI

cube.dev/blog/headless-bi

cube.dev

cube blog

Data Engineering | Application Development | All Things Cube | Join us on Slack | Get Cube

What is Headless BI?

The four essential components of a headless BI tool

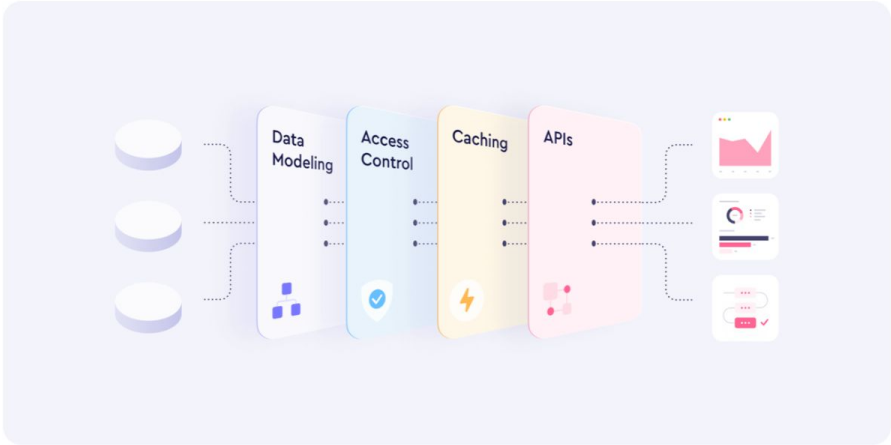

March 24, 2022

Deep Dives

Artyom Keydunov

Co-founder and CEO at Cube

Headless Business Intelligence

- Universal connectivity to data sources
  - Even multiple data sources at once
- **Data modeling layer** provides consistent metrics definitions
- **Access control layer** provides consistent security policy
- **Caching layer** provides consistent freshness and fast access to data
- **API layer** provides universal connectivity to data consumers
  - REST and GraphQL APIs — mostly for programmatic access
  - SQL API — mostly for BI tools and data notebooks

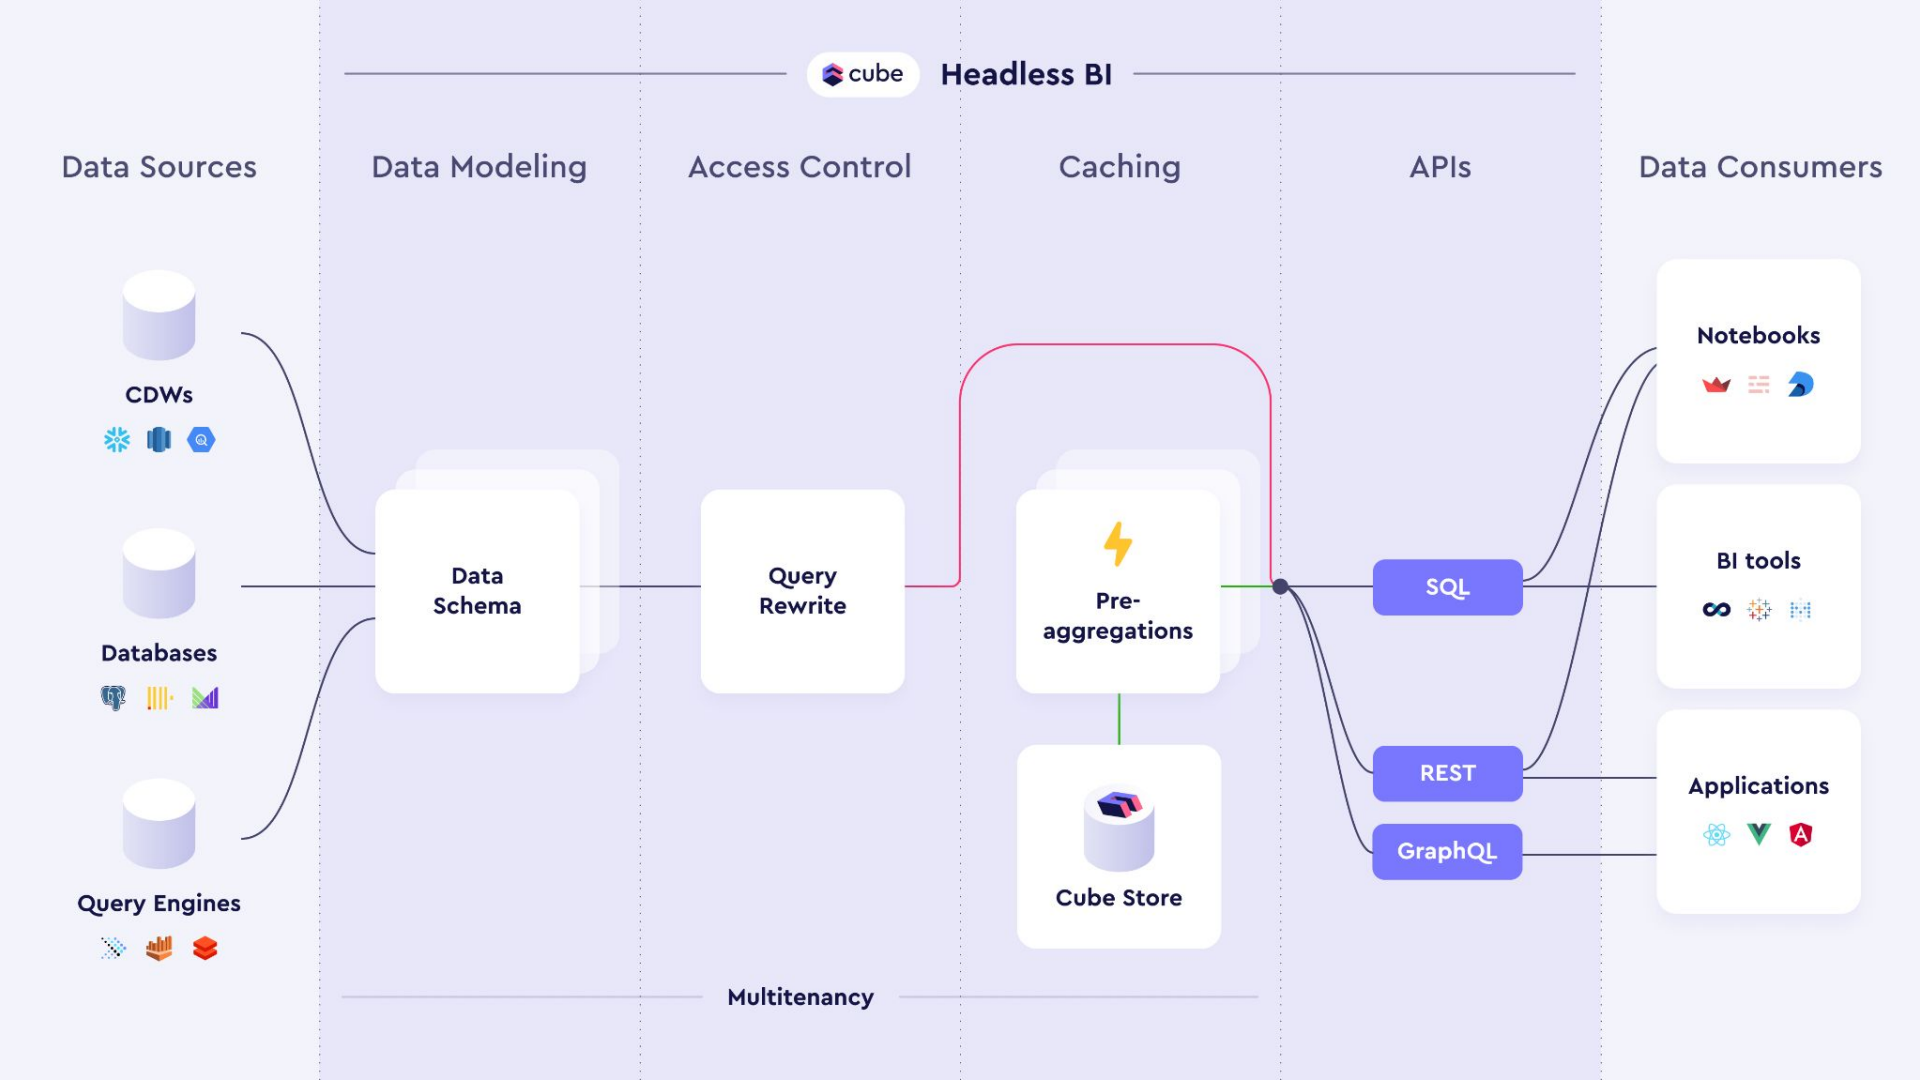
cube
Headless BI
Data Sources
Data Modeling
Access Control
Caching
APIs
Data Consumers
CDWs
Databases
Query Engines
Data Schema
Query Rewrite
Pre-aggregations
Cube Store
SQL
REST
GraphQL
Notebooks
BI tools
Applications
Multitenancy

Update to SQL API

cube.dev/blog/expanded-bi-support

cube.dev

cube blog

Data Engineering | Application Development | All Things Cube | Join us on Slack | Get Cube

Announcing expanded BI support with our SQL API

Connect your favorite visualization tools to Cube.

May 24, 2022

News & Product Updates

Artyom Keydunov

Co-founder and CEO at Cube

Expanded BI support with the SQL API

- Preview of the SQL API announced in November 2021
- Initial support for data notebooks and BI tools, e.g., Apache Superset
- A lot of thoughtful and positive feedback from Cube's users
- A lot of engineering and product work
  - PostgreSQL dialect as the SQL dialect
  - Apache Arrow Datafusion as the SQL execution engine
- Update to the SQL API announced in May 2022
- Support for many BI tools and data notebooks—and counting

Expanded BI support with the SQL API

- Preview of the SQL API announced in November 2021
- Initial support for data notebooks and BI tools, e.g., Apache Superset
- A lot of thoughtful and positive feedback from Cube's users
- A lot of engineering and product work
  - PostgreSQL dialect as the SQL dialect
  - Apache Arrow Datafusion as the SQL execution engine
- Update to the SQL API announced in May 2022
- Support for many BI tools and data notebooks—and counting

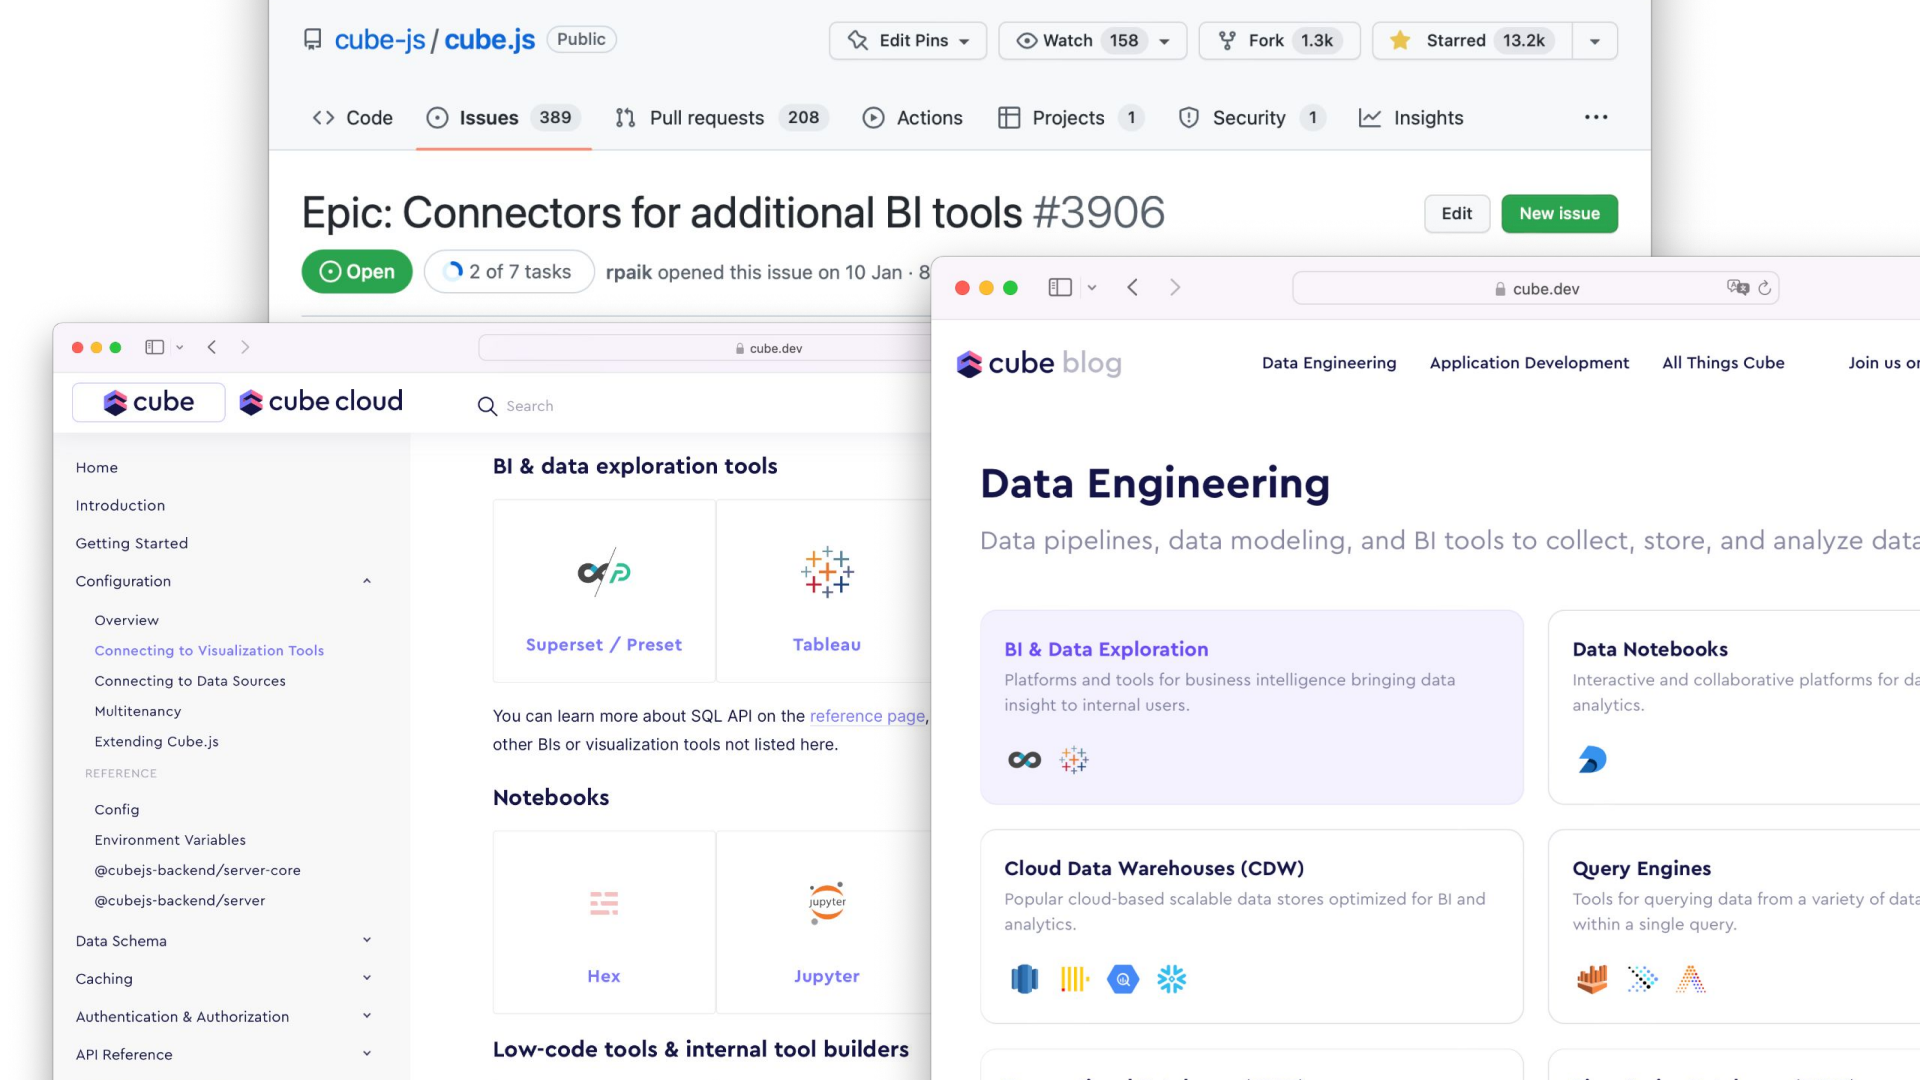
cube-js / cube.js
Public
Edit Pins
Watch 158
Fork 1.3k
Starred 13.2k
Code
Issues 389
Pull requests 208
Actions
Projects 1
Security 1
Insights
Epic: Connectors for additional BI tools #3906
Edit
New issue
Open
2 of 7 tasks
rpaik opened this issue on 10 Jan · 8
cube.dev
cube
cube cloud
Search
Home
Introduction
Getting Started
Configuration
Overview
Connecting to Visualization Tools
Connecting to Data Sources
Multitenancy
Extending Cube.js
REFERENCE
Config
Environment Variables
@cubejs-backend/server-core
@cubejs-backend/server
Data Schema
Caching
Authentication & Authorization
API Reference
BI & data exploration tools
Superset / Preset
Tableau
You can learn more about SQL API on the reference page,
other BIs or visualization tools not listed here.
Notebooks
Hex
Jupyter
Low-code tools & internal tool builders
cube.dev
cube blog
Data Engineering
Application Development
All Things Cube
Join us on
Data Engineering
Data pipelines, data modeling, and BI tools to collect, store, and analyze data
BI & Data Exploration
Platforms and tools for business intelligence bringing data insight to internal users.
Data Notebooks
Interactive and collaborative platforms for da
analytics.
Cloud Data Warehouses (CDW)
Popular cloud-based scalable data stores optimized for BI and analytics.
Query Engines
Tools for querying data from a variety of data
within a single query.

Supported BI & data exploration tools

See the list of supported BI tools in the docs. Track progress on GitHub.

- Apache Superset
- Tableau
- Microsoft Power BI
- Metabase — coming soon
- Google Data Studio — coming soon
- Google Sheets — planned
- Microsoft Excel — planned

Supported data notebooks

See the list of supported data notebooks in the docs.

- Jupyter
- Deepnote
- Hex
- Streamlit
- Observable
- and many more

Feel free to suggest more tools on GitHub.

Connecting to SQL API

Connecting and using SQL API (self-hosted)

See the docs.

- Set `CUBEJS_PG_SQL_PORT` to enable the SQL API
- Provide `CUBEJS_SQL_USER` to set the user name
- Provide `CUBEJS_SQL_PASSWORD` to set the password
- Test with `psql -h 127.0.0.1 --port 5432 -U user --password`

If you're running Cube in Docker, don't forget to expose (map) the port.

Feel free check this Docker Compose example config.

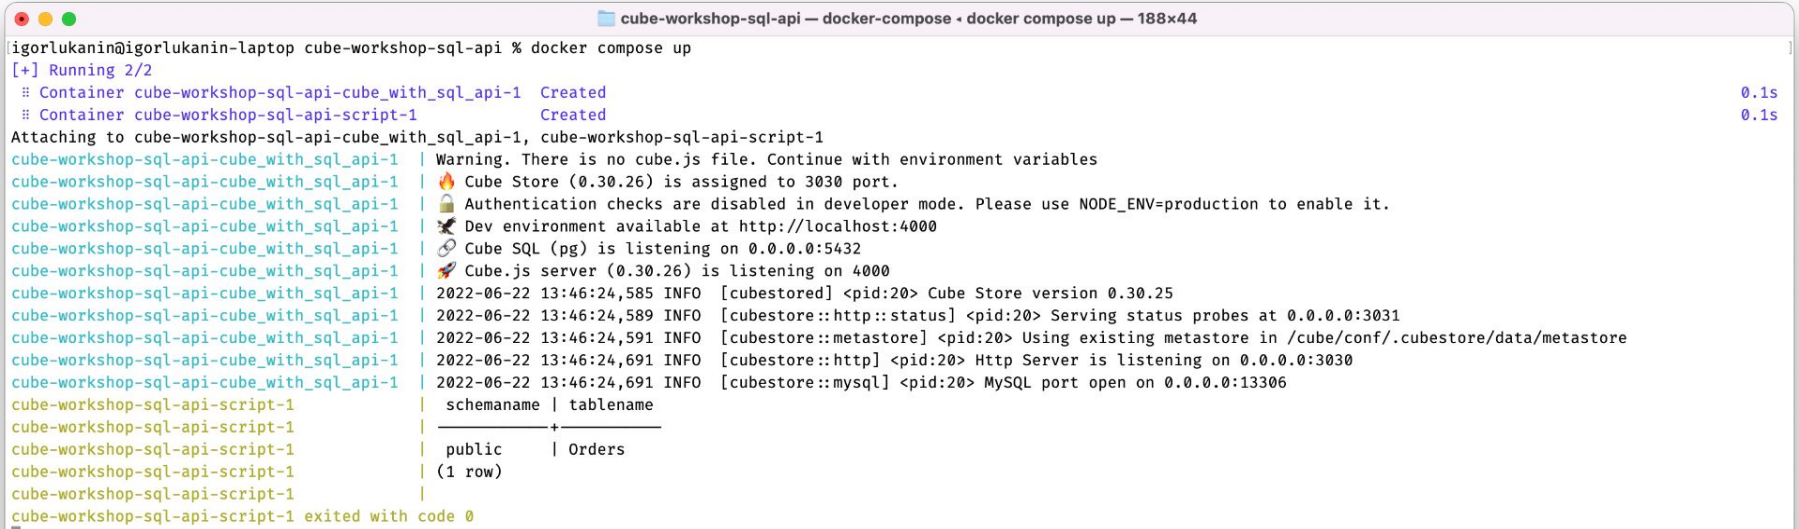
cube-workshop-sql-api — docker-compose ◂ docker compose up — 188×44
igorlukanin@igorlukanin-laptop cube-workshop-sql-api % docker compose up
[+] Running 2/2
⠿ Container cube-workshop-sql-api-cube_with_sql_api-1 Created 0.1s
⠿ Container cube-workshop-sql-api-script-1 Created 0.1s
Attaching to cube-workshop-sql-api-cube_with_sql_api-1, cube-workshop-sql-api-script-1
cube-workshop-sql-api-cube_with_sql_api-1 | Warning. There is no cube.js file. Continue with environment variables
cube-workshop-sql-api-cube_with_sql_api-1 | 🔥 Cube Store (0.30.26) is assigned to 3030 port.
cube-workshop-sql-api-cube_with_sql_api-1 | 🔓 Authentication checks are disabled in developer mode. Please use NODE_ENV=production to enable it.
cube-workshop-sql-api-cube_with_sql_api-1 | 🦅 Dev environment available at http://localhost:4000
cube-workshop-sql-api-cube_with_sql_api-1 | 🔗 Cube SQL (pg) is listening on 0.0.0.0:5432
cube-workshop-sql-api-cube_with_sql_api-1 | 🚀 Cube.js server (0.30.26) is listening on 4000
cube-workshop-sql-api-cube_with_sql_api-1 | 2022-06-22 13:46:24,585 INFO [cubestored] <pid:20> Cube Store version 0.30.25
cube-workshop-sql-api-cube_with_sql_api-1 | 2022-06-22 13:46:24,589 INFO [cubestore::http::status] <pid:20> Serving status probes at 0.0.0.0:3031
cube-workshop-sql-api-cube_with_sql_api-1 | 2022-06-22 13:46:24,591 INFO [cubestore::metastore] <pid:20> Using existing metastore in /cube/conf/.cubestore/data/metastore
cube-workshop-sql-api-cube_with_sql_api-1 | 2022-06-22 13:46:24,691 INFO [cubestore::http] <pid:20> Http Server is listening on 0.0.0.0:3030
cube-workshop-sql-api-cube_with_sql_api-1 | 2022-06-22 13:46:24,691 INFO [cubestore::mysql] <pid:20> MySQL port open on 0.0.0.0:13306
cube-workshop-sql-api-script-1 | schemaname | tablename
cube-workshop-sql-api-script-1 | ------------+-----------
cube-workshop-sql-api-script-1 | public | Orders
cube-workshop-sql-api-script-1 | (1 row)
cube-workshop-sql-api-script-1 |
cube-workshop-sql-api-script-1 exited with code 0

Connecting and using SQL API (Cube Cloud)

See the workshop preparation guide.

- We need a Cube instance — let's take one in Cube Cloud

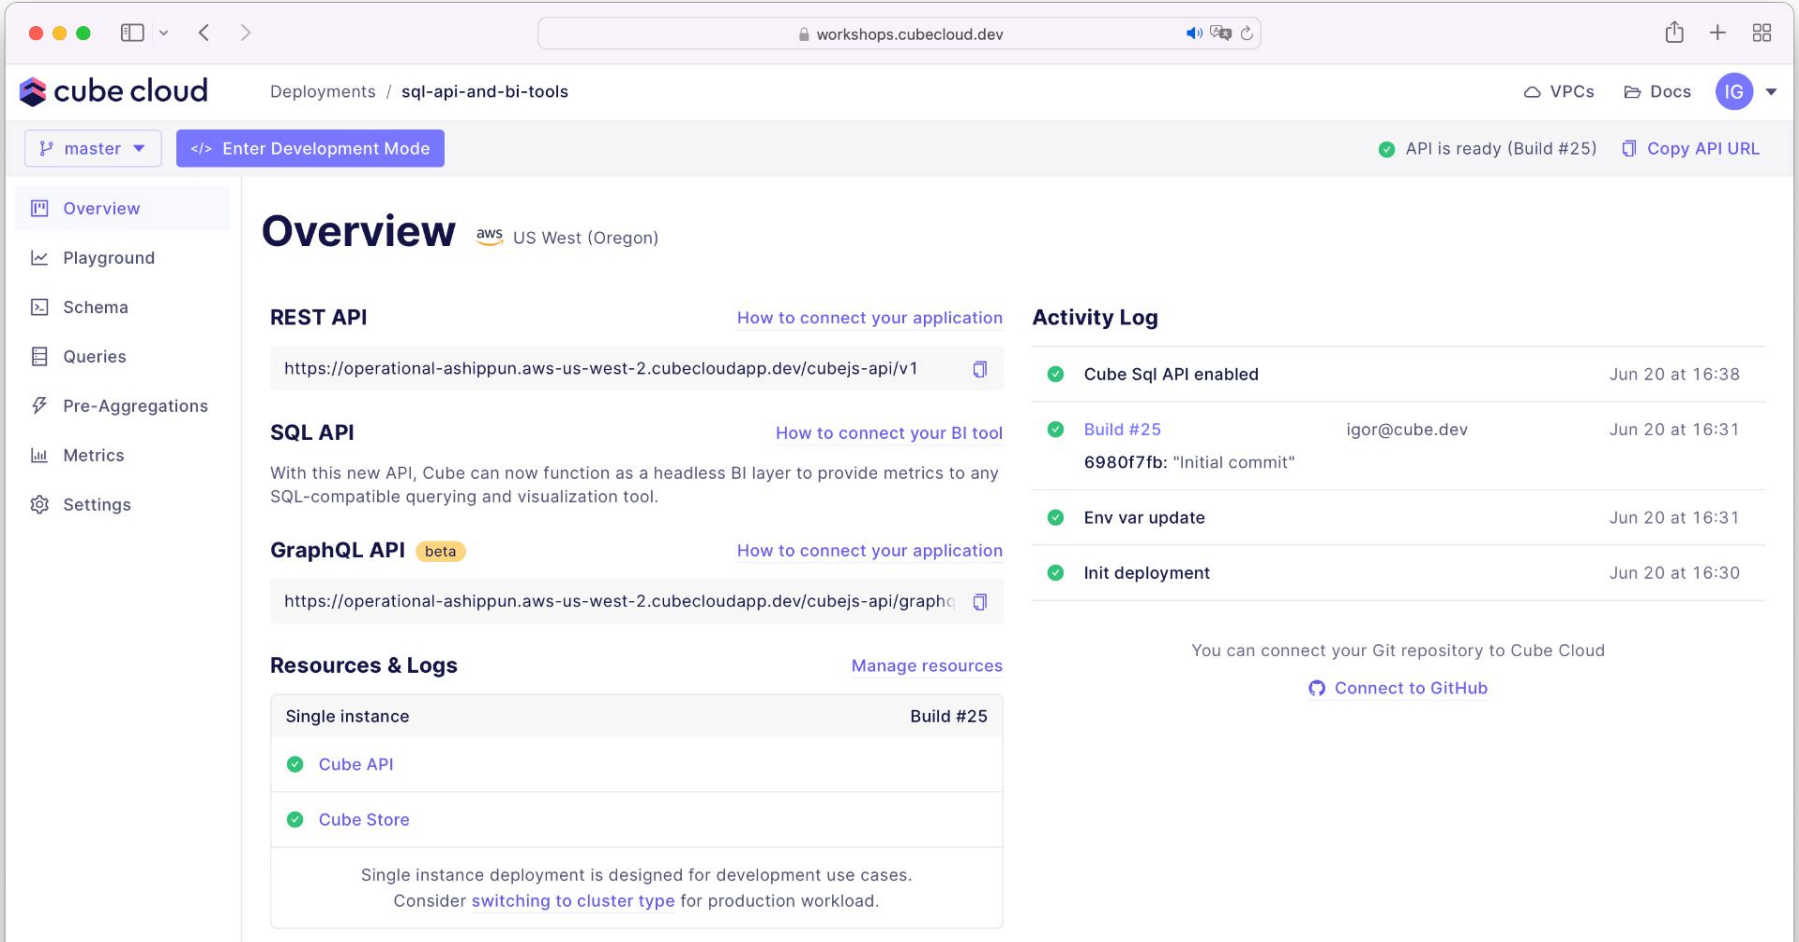
workshops.cubecloud.dev
cube cloud
Deployments / sql-api-and-bi-tools
VPCs
Docs
IG
master
</> Enter Development Mode
API is ready (Build #25)
Copy API URL
Overview
Playground
Schema
Queries
Pre-Aggregations
Metrics
Settings
Overview
aws US West (Oregon)
REST API
How to connect your application
https://operational-ashippun.aws-us-west-2.cubecloudapp.dev/cubejs-api/v1
SQL API
How to connect your BI tool
With this new API, Cube can now function as a headless BI layer to provide metrics to any SQL-compatible querying and visualization tool.
GraphQL API
beta
How to connect your application
https://operational-ashippun.aws-us-west-2.cubecloudapp.dev/cubejs-api/graphq
Resources & Logs
Manage resources
Single instance
Build #25
Cube API
Cube Store
Single instance deployment is designed for development use cases.
Consider switching to cluster type for production workload.
Activity Log
Cube Sql API enabled
Jun 20 at 16:38
Build #25
igor@cube.dev
Jun 20 at 16:31
6980f7fb: "Initial commit"
Env var update
Jun 20 at 16:31
Init deployment
Jun 20 at 16:30
You can connect your Git repository to Cube Cloud
Connect to GitHub

Connecting and using SQL API (Cube Cloud)

See the workshop preparation guide.

- We need a Cube instance — let's take one in Cube Cloud
- We need a BI tool — let's take Tableau (also, check this guide)

Connecting and using SQL API (Cube Cloud)

See the workshop preparation guide.

- We need a Cube instance — let's take one in Cube Cloud
- We need a BI tool — let's take Tableau (also, check this guide)
- We need a BI tool — let's take Superset (also, check this guide)

Connecting and using SQL API (Cube Cloud)

See the workshop preparation guide.

- We need a Cube instance — let's take one in Cube Cloud
- We need a BI tool — let's take Tableau (also, check this guide)
- We need a BI tool — let's take Superset (also, check this guide)
- We need a data notebook — let's take Deepnote (also, check this guide)
  - SQL API supports development mode in Cube Cloud
  - Let's get into development mode!

Q&A

SQL API features

SQL API features

- Exposed entities, e.g., tables and columns
- Query types: aggregated and non-aggregated
- Selection: filtering, ordering, limiting
- Projection: star, aggregate functions, aliases

Feel free to check the SQL API documentation.

Known limitations

Not supported yet:

- Joins. Please use proxy dimensions to model joins, they work great
- Time zones: https://github.com/cube-js/cube.js/issues/4142
- Data blending: https://github.com/cube-js/cube.js/issues/4524

Joins via proxy dimensions

- In the data model, you can cross-reference cube members
  - `sql: ` ${Orders.count}` `
- You can add a join by adding a *proxy dimension* to any cube
  - ...a reference to a member of a joined cube
- We plan to support regular join syntax in future releases

Q&A

Pre-aggregations

Pre-aggregations

- Pre-aggregations are a part of the caching layer
- They accelerate queries and make all queries sub-second

Multitenancy

Multitenancy

- Multitenancy spans across all Cube's layers
- Supported in SQL API with the custom authentication

Extra: More tools

To learn more visit...

- Documentation:
  - SQL API: https://cube.dev/docs/backend/sql
- Blog:
  - https://cube.dev/blog/expanded-bi-support
  - https://cube.dev/blog/headless-bi
  - https://cube.dev/blog/category/data
- Community support:
  - Slack: https://slack.cube.dev


cube

cube.dev/blog/category/data