



Monthly Community Call

October 12, 2022



Community Code of Conduct

- We want to foster an open and welcoming environment where everyone feels they belong in the Cube 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.

What we will discuss today

- Community news
- Product update (plus some sneak peeks!)
- Q&A

Community news

Headless BI with streaming data



September 22, 2022

All Things Cube



Artyom Keydunov  

Co-founder and CEO at Cube

Check out the blog post at
<https://cube.dev/blog/headless-bi-with-streaming-data>

Events recap (<https://cube.dev/events>)



Past Events

Oct 6, 2022 **Building a fast & open-source data stack with ClickHouse and Cube** 

Check out all the videos on our Youtube channel [View on YouTube](#)


You can find both videos and slides from past events!


Oct 4, 2022  **Replacing Looker with Cube and Apache Superset**

 **Shreesham Mukherjee**
Developer Relations Engineer at Preset




Oct 4, 2022 **Cube at Current 2022**

Sep 28, 2022   **Building lightning-fast data apps with Firebolt and Cube**


 **Luka Lovosevic**
Solution Architect at Firebolt



Cube at Current'2022 (<https://2022.currentevent.io>)




Next workshop:



- Date: October 13, 2022 at 9am Pacific Time
- Topic: Headless BI with streaming data
- Register for the event at <https://cube.dev/events/building-real-time-streaming-data-applications-with-cube> 



Online

Headless BI with streaming data

 cube

Oct 13, 2022 @ 9:00 AM 🔥 2 days left

 **Artyom Keydunov** 
Co-founder and CEO at Cube

 **Pavel Tiunov** 
PhD, Co-founder and CTO at Cube

In this 60-minute workshop, we will be discussing our latest product enhancements to support real-time/streaming data in Cube.

In particular, we'll demonstrate:

- Uniting with your source data
- Ingesting real-time data into cube from [ksqldb](#), [Materialize](#), etc.
- [Streaming pre-aggregations](#)
- and much more 🚀

02 : 15 : 05 : 05
days hours min sec

We are sponsors at the Trino Summit

<https://www.starburst.io/info/trinosummit/>

Hosted by  Starburst

 **trino** SUMMIT

NOVEMBER 10TH, 2022



THANK YOU TO OUR SPONSORS!

 privacera

 MC MONTE CARLO

 MMUTA™

 cube

 hightouch

Product updates + sneak peeks

Q3'2022 roadmap update

Roadmap on GitHub: github.com/cube-js/cube.js/projects/13

github.com

cube-js / cube.js > Projects > Roadmap 2022 Q3

Filter cards

+ Add cards Exit fullscreen Menu

1 Epics

- Epic: Replace Redis with Cube Store** #5042 opened by rpaik
 - backend:cube-store Epic
 - Roadmap: 2022 Q3

4 In progress

- Epic: Build metrics in Cube data model** #5044 opened by rpaik
 - Epic Roadmap: 2022 Q3
 - 2 linked pull requests
- feat: Cube Views implementation** #5278 opened by paveltiunov
 - 3 of 4 tasks
 - Epic Roadmap: 2022 Q3
- feat: Includes and Excludes directives for Cube Views** #5437 opened by paveltiunov
 - 3 of 4 tasks
 - Epic Roadmap: 2022 Q3
- Epic: Add YAML support** #5045 opened by rpaik
 - Epic Roadmap: 2022 Q3

6 Done

- Epic: Streaming pre-aggregations** #5043 opened by rpaik
 - backend:pre-aggregations Epic
 - Roadmap: 2022 Q3
- feat: ksql support** #3507 opened by paveltiunov
 - 3 of 4 tasks
 - Roadmap: 2022 Q3
- docs(schema): document unionWithSourceData and rollupLambda** #5297 opened by hassankhan
 - 4 tasks done
 - docs
- docs: add lambda pre-aggs and ksqldb** #5321 opened by keydunov
 - docs Roadmap: 2022 Q3
- examples: enhance ksql example**

Q3'2022 roadmap update

Roadmap on GitHub: github.com/cube-js/cube.js/projects/13

Real-time update:

- [ksqlDB integration](#)
- [Streaming pre-aggregations](#)
- [Lambda pre-aggregations](#)

Data modeling update:

- All-new [Views](#) in Cube's data modeling layer

ksqlDB integration

ksqlDB integration

- Cube is built to serve as a single, headless data access layer, regardless of whether the data is a *stream*, a *batch*, or a *union of both*
- [Apache Kafka](#) is the most popular event streaming platform in the world
- [ksqlDB](#) is the purpose-built streaming database for Kafka
 - Also works with [Redpanda](#)!
- Now you can access data in Kafka with Cube

Docs: cube.dev/docs/config/databases/ksqldb

Demo: [ksql-demo](#)



ksql-demo.netlify.app



All Examples ▾

Discuss on Slack

Get started with Cube

Real-Time Dashboard with Kafka, ksqlDB and Cube

This live demo shows a real-time dashboard built with Kafka, ksqlDB and Cube.
Read about [ksqlDB and Cube integration](#) or explore the [source code](#) to learn more.

I will click myself in 5 seconds or [Click Me Now](#)

Users Online

1

Real Time Events



Button Clicks (Last Hour)

Streaming pre-aggregations

Streaming pre-aggregations

- Streaming pre-aggregations follow the “push” model
- Cube subscribes to the updates and keeps pre-aggregation up-to-date
 - You don't need to define **refreshKey**
- Whether pre-aggregation is streaming or not is defined by the data source
 - Currently, Cube supports only ksqldb
- Next data sources to come: Materialize, Flink SQL, Spark Streaming
- Let us know if these or other streaming engines: cube.dev/contact

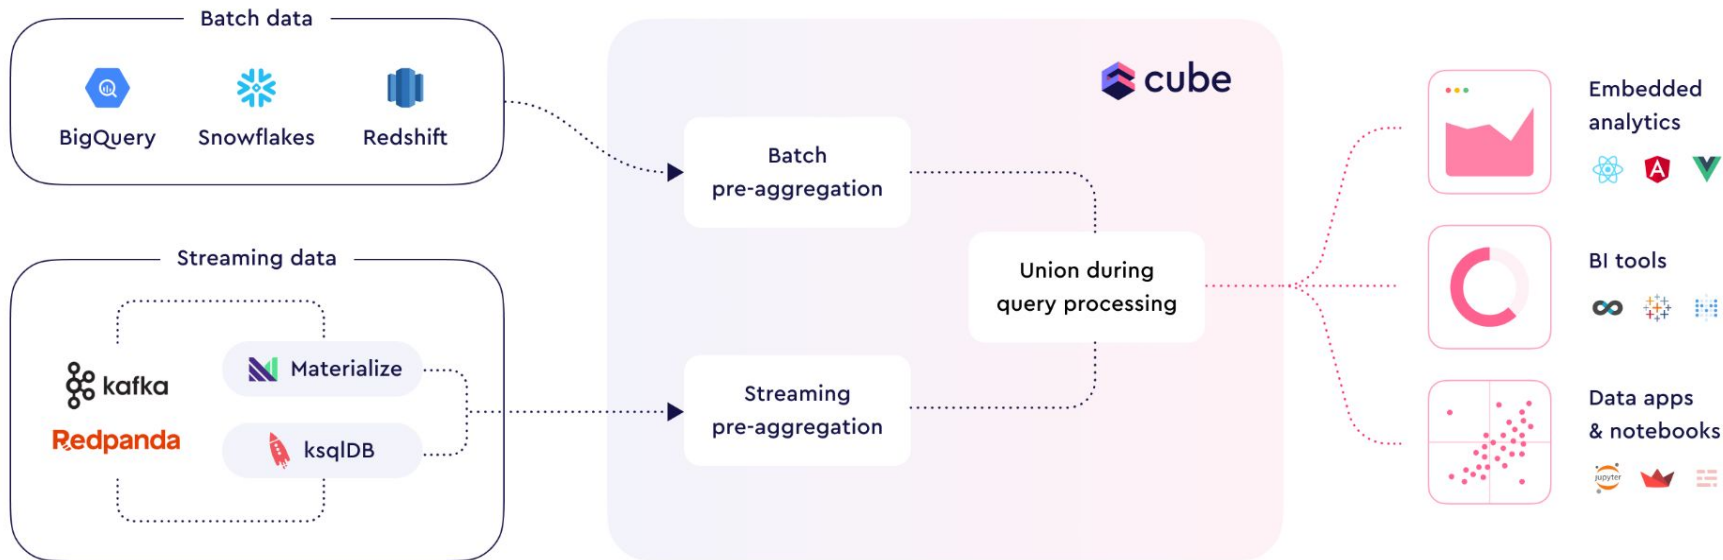
cube.dev/docs/caching/using-pre-aggregations#streaming-pre-aggregations

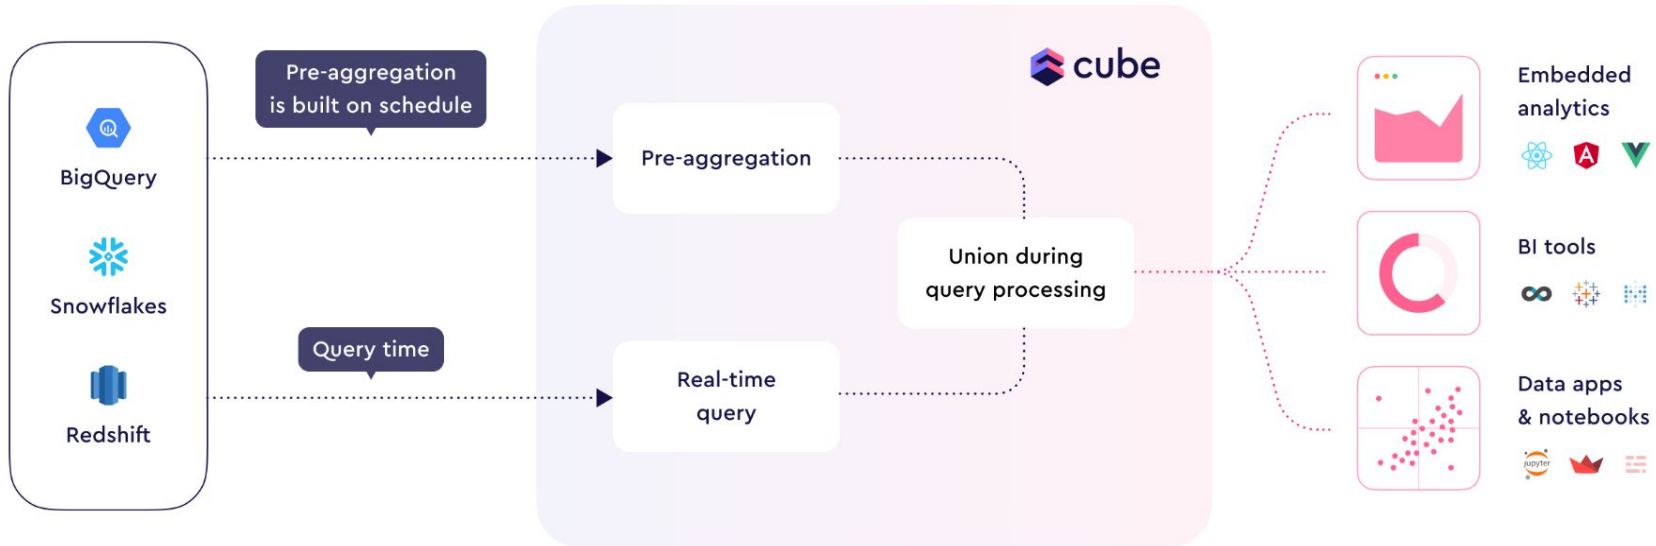
Lambda pre-aggregations

Lambda pre-aggregations

- Lambda pre-aggregations follow the [Lambda architecture](#) design to union real-time and batch data
- Batch layer: regular *pre-aggregations*
- Speed layer: *streaming pre-aggregations* or *source data*

Docs: cube.dev/docs/caching/pre-aggregations/lambda-pre-aggregations





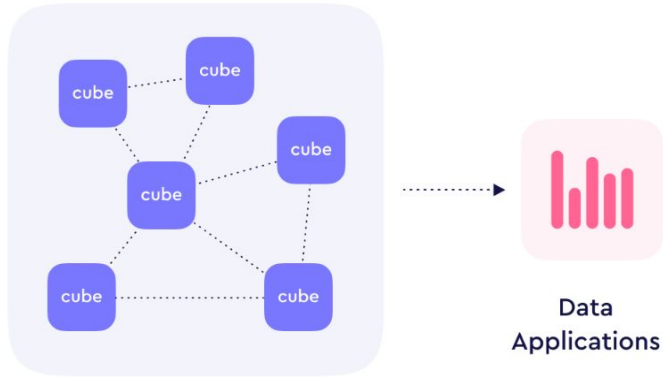
Introducing Views

Views in Cube's data modeling layer

- Views sit on top of cubes and create a facade of your whole data model
- Views are similar to cubes but they can't have their own members
- You can include select measures and dimensions into a view from cubes

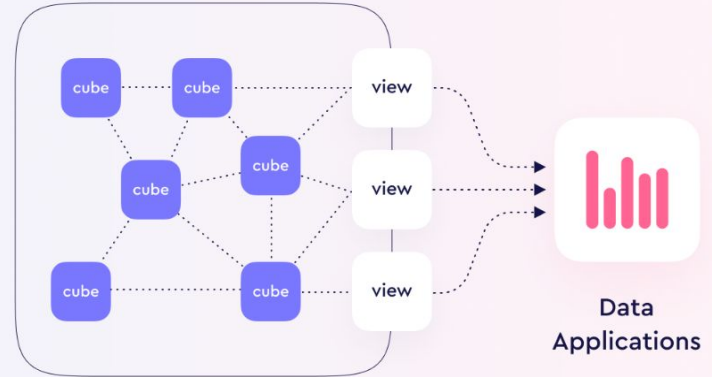
Docs: cube.dev/docs/schema/reference/view

Before



How to control join path?
How to manage visibility of cubes to different teams/users?
How to define metrics?

Now



Views are the place to define metrics, control ambiguous join paths, and manage governance and accessibility for data consumers.

Use cases for views

- Defining metrics
 - Views with a single measure can be used to define metrics
- Managing visibility and governance
 - Views have the `shown` property to control their visibility, e.g.,
`shown: COMPILER_CONTEXT.permissions['finance']`
- Controlling join paths
 - Views allow to define an exact join path to a measure, e.g.,
`Orders.Users.Countries.name`

