

Rill Data

Flying Fast at Scale with DuckDB

Brian Holmes

The Rill logo is rendered in a bold, white, sans-serif font. The letter 'R' is significantly larger and more prominent than the 'ill' which follows it. The background of the slide features a dark blue gradient with abstract, wavy shapes in a slightly lighter shade of blue, creating a sense of motion and depth.

Customer	# ↓	Δ%
Zoom	\$125.9k	-13%
Salesforce	\$82.9k	-15%
Shopify	\$78.8k	-17%
Dropbox	\$52.8k	-9%
Twilio	\$46.6k	-38%
Zapier	\$45.9k	-10%
Intercom	\$28.9k	-24%

(Expand Table)



demo / i18-cost-monitoring / Customer Unit Eco

Last 4 Weeks Apr 3 - May 1, 2023 11:00AM-12:00AM

Cost by App Name Ingestion-1

All Measures

2023 Apr 06

Total Cost \$60.0k

Search

- Adobe
- Airtable
- Algolia
- Algosec
- Asana
- Atlassian
- Auth0
- Basecamp
- BigCommerce
- Calendly
- Canva
- Cloudflare

Select all Exclude

Metric trends by day

Cost by App Name # ↓

storage-1	\$71.6k
ingestion-1	\$60.0k
null	no data

Plan Name # ↓

Standard Plan	\$60.0k
---------------	---------

Cost by SKU # ↓

Data Delivery	\$39.5k
Products	\$14.5k

Cost by Data Pipeline # ↓

shopify_stream	\$43.4k
shopify_ads	\$7.9k
dropbox	\$5.1k
shopify_orders	\$3.3k
shopify_v2	\$411.14
dropbox_test	\$4.50

Cost by Environment # ↓

production	\$442.9k	-13%
stage	\$48.5k	15%
dev	\$44.7k	-34%
null	no data	no data



Create Alert ✕

1 Data 2 Criteria 3 Delivery

Filters
These are inherited from the underlying dashboard view.

Year to Date

Alert data
Select the measures you want to monitor.

Measure
Number of commits

Split by dimension (optional)
Username

Data preview
Here's a look at the data you've selected above.

Username	Number of commits
Laurens Kuiper	503
Carlo Piovesan	572
Pedro Holanda	885

Cancel Next

Schedule report

Email recurring exports to recipients.

Report title
My report

Frequency **Day** **Time** **Time zone**

Weekly Thursday 7:45am PDT (Local)

Format
CSV

Row limit (optional)
1000

Recipients

brian.holmes@rilldata.com

Enter an email address

+ Add email

Cancel Create

[Copy URL](#) [Create public URL](#)

Create a public URL that you can send to anyone.

- Measures and dimensions will be limited to current visible set.
- Filters will be locked and hidden.

Set expiration

Access expires 10/14/2024

Create and copy URL

Why is Rill's BI tool different?

10X Faster Dashboards

Data and compute are co-located, queries return instantly

BI-as-Code

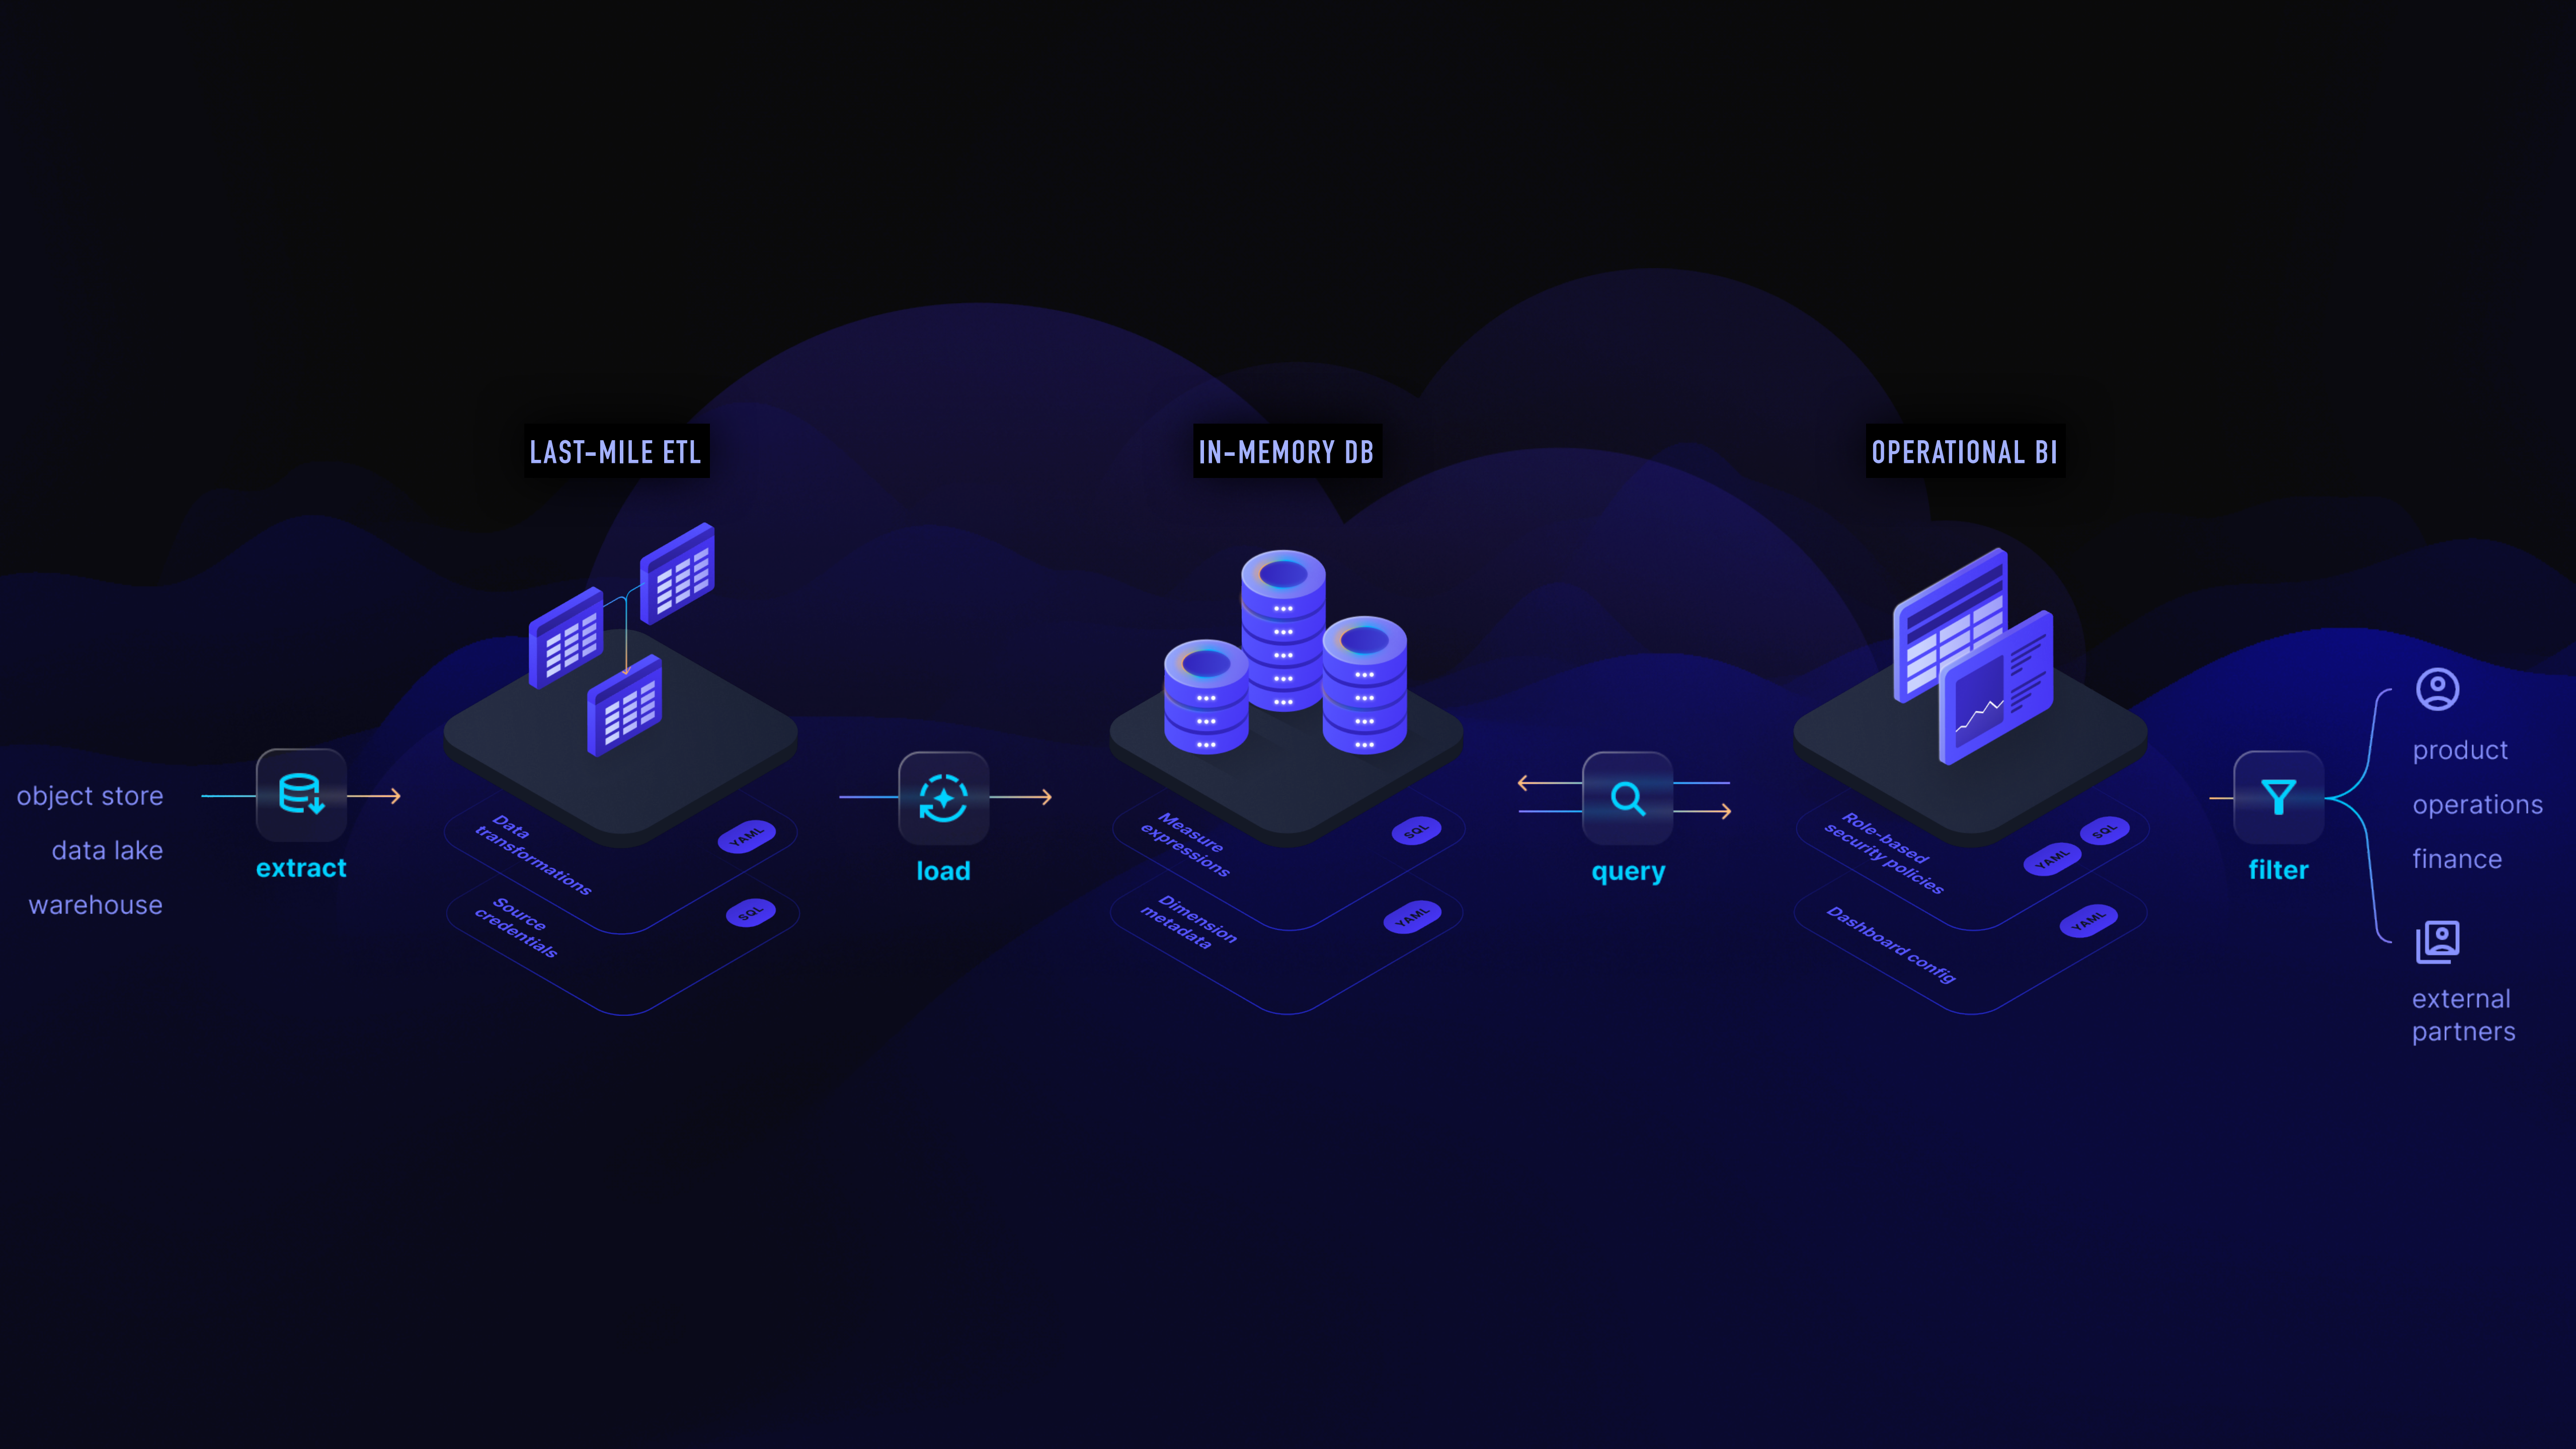
Develop locally, deploy globally with benefits of Git workflows

Last-mile ETL with SQL

Join, transform, & aggregate data from dozens of sources

Metrics-first Philosophy

Focus on metrics, Rill automatically generates visualizations



LAST-MILE ETL

IN-MEMORY DB

OPERATIONAL BI

object store
data lake
warehouse

extract

Data transformations (YAML)
Source credentials (SQL)

load

Measure expressions (SQL)
Dimension metadata (YAML)

query

Role-based security policies (YAML, SQL)
Dashboard config (YAML)

filter

- product
- operations
- finance
- external partners

+ Add

- dashboards
 - auction.yaml
 - bids.yaml
 - duckdb_commits.yaml
 - offerings_issuers.yaml
- models
- sources
- .gitignore
- README.md
- rill.yaml

CONNECTORS

- duckdb OLAP
 - main_db
 - main
 - ISSUERS
 - OFFERINGS
 - RECIPIENTS
 - RELATEDPERSONS
 - SIGNATURES
 - auction_data_mo...
 - auction_data_raw
 - bids_data_model
 - bids_data_raw
 - duckdb_commits...
 - duckdb_commits...
 - duckdb_modified...
 - offerings_issuers...

```
1 # Visit https://docs.rilldata.com/ to learn more about Rill code files.
2
3 type: source
4 connector: "gcs"
5 uri: "gs://rilldata-public/github-commits"
6
7 extract:
8   files:
9     strategy: tail
10    size: 1
```

Save Revert changes

#	A	commit_hash							
1		ba75d81601913782d28a3878707d135							
2		82f1559651a20e7a46890ac5127f4d58							
3		4bbe14a209e1d9db965b0f4f9136993							
4		3db9596e75b3beaf1e0fe1b48b8847c							
5		2b8130d335ab4afef2ae641633fdcf7b							
6		dea8a2fd666823af58199d75ae70bca							
7		c14647dd6575972e8d5d81c95d38083ced4d...	Add logical query plan generator.		Mark Raasveldt	mark.raasveldt@gmail.com	2018-07-19 09:41:57 Z		
8		990aeef2d1302e632230ca3fbd40ef66fc670...	Fix GCC compilation.		Mytherin	mark.raasveldt@gmail.com	2018-07-19 13:52:13 Z		
9		f66100070cf8cf367a86d0b6b2e3cd8e25e27...	Use vector + map instead of just map to store...		Mark Raasveldt	mark.raasveldt@gmail.com	2018-07-19 15:45:22 Z		

Add a source

- Google Cloud Storage
- amazon S3
- Microsoft Azure Blob Storage
- Google BigQuery
- amazon ATHENA
- amazon REDSHIFT
- DuckDB
- MotherDuck
- PostgreSQL
- MySQL
- SQLite
- snowflake
- salesforce
- Local file
- https://

Connect an OLAP engine

- ClickHouse
- druid
- pinot

Don't see what you're looking for? [Request a new connector](#)

Untitled Rill Project

offerings_issuers_model.sql

Computed on Aug 15, 2024, 7:33 AM Export Generate dashboard

- dashboards
- models
 - auction_data_model.sql
 - bids_data_model.sql
 - offerings_issuers_model.sql
- sources
 - .gitignore
 - README.md
 - rill.yaml

- CONNECTORS
- duckdb OLAP
 - main_db
 - main
 - ISSUERS
 - OFFERINGS
 - RECIPIENTS
 - RELATEDPERSONS
 - SIGNATURES
 - auction_data_mo...
 - auction_data_raw
 - bids_data_model
 - bids_data_raw
 - offerings_issuers...

```

1
2 SELECT
3   SALE_DATE AS offering_saledate,
4   INDUSTRYGROUPTYPE AS offering_IndustryGroupType,
5   INVESTMENTFUNDTYPE AS offering_InvestmentFundType,
6   ISSUERS.ENTITYNAME AS issuer_EntityName,
7   RECIPIENTS.RECIPIENTNAME AS recipient_name,
8   CAST(
9     CASE
10      WHEN TOTALOFFERINGAMOUNT == 'Indefinite' THEN 0
11      ELSE TOTALOFFERINGAMOUNT
12    END
13    AS BIGINT) AS TotalOfferingAmount,
14  *
15 FROM
16  OFFERINGS
17  LEFT JOIN ISSUERS ON OFFERINGS.ACCESSIONNUMBER = ISSUERS.ACCESSIONNUMBER
18  LEFT JOIN RECIPIENTS ON OFFERINGS.ACCESSIONNUMBER = RECIPIENTS.ACCESSIONNUMBER
19 WHERE SALE_DATE BETWEEN DATE '2020-01-01' AND current_date()
20
21 ORDER BY SALE_DATE DESC
    
```

9.71% of source rows **137,340 rows**
6 columns added **85 columns**

Referenced in this model

- ISSUERS 552k rows
- OFFERINGS 545k rows
- RECIPIENTS 318k rows

Model columns

sort by type show summary

Column Name	Count	Percentage
SALE_DATE		0%
offering_saledate		0%
TotalOfferingAmount		0%
TOTALNUMBERALREADYINV...		0%
TOTALAMOUNTSOLD		0%
SALESCOMM_DOLLARAMO...		0%
MINIMUMINVESTMENTACCE...		0%
ISSUER_SEQ_KEY		0%
GROSSPROCEEDSUSED_DO...		0%
FINDERSFEE_DOLLARAMOU...		0%
YEAROFINC_VALUE_ENTERED		11%
RECIPIENT_SEQ_KEY		63%
NUMBERNONACCREDITEDI...		96%
A ACCESSIONNUMBER	103k	0%
A ACCESSIONNUMBER_1	103k	0%
A issuer_EntityName	77.9k	0%
A ENTITYNAME	77.9k	0%
A CIK	76.8k	0%
A STREET1	42.2k	0%
A ISSUERPHONENUMBER	39.3k	0%
A TOTALOFFERINGAMOUNT_1	33.3k	0%
A PREVIOUSACCESSIONNUMB...	18.1k	68%
A TOTALREMAINING	17.5k	0%
A ACCESSIONNUMBER_2	17.3k	63%
A STREET2	8.95k	40%
A ZIPCODE	8.86k	0%
A SALESAMTCLARIFICATIONO...	7.62k	78%
A GROSSPROCEEDSUSED_CLA...	7.06k	58%
A STATES_OR_VALUE_LIST	5.85k	65%
A DESCRIPTIONS_LIST	5.78k	65%
A STREET1_1	5.67k	63%
A CITY	5.26k	0%
A recipient_name	5.03k	63%
A RECIPIENTNAME	5.03k	63%
A EDGAR_PREVIOUSNAME_1	4.45k	94%
A ISSUER_PREVIOUSNAME_1	3.48k	86%
A DESCRIPTIONOFOTHERTYPE	3.44k	92%
A FINDERFEECLARIFICATIONO...	2.97k	78%
A RECIPIENTCRDNUMBER	2.50k	63%
A ZIPCODE_1	2.07k	63%
A ASSOCIATEDBDNAME	1.79k	63%

#	offering_saledate	offering_IndustryGroupType	offering_InvestmentFundType	issuer_EntityName
1	2022-12-30	Other Health Care	null	IzoMed, Inc.
2	2022-12-30	Pooled Investment Fund	Venture Capital Fund	TH71 Fund I, a series of Roll Up Vehicles, LP
3	2022-12-30	Pooled Investment Fund	Venture Capital Fund	ED23 Fund I, a series of Roll Up Vehicles, LP
4	2022-12-30	Pooled Investment Fund	Venture Capital Fund	UP19 Fund I, a series of Roll Up Vehicles, LP
5	2022-12-29	Pooled Investment Fund	Other Investment Fund	IRE U.S. Fund VI Blocker, LLC
6	2022-12-29	Computers	null	Mini Mines II, LLC
7	2022-12-29	Pooled Investment Fund	Venture Capital Fund	PR80 Fund I, a series of Roll Up Vehicles, LP
8	2022-12-29	Pooled Investment Fund	Venture Capital Fund	TY Fund I, a series of Climate Capital Deals, ...
9	2022-12-29	Pooled Investment Fund	Venture Capital Fund	RE-1117 Fund I, a series of Roll Up Vehicles, LP

UR Untitled Rill Project

+ Add

- dashboards
 - auction.yaml
 - bids.yaml
 - offerings_issuers.yaml
- models
 - auction_data_model.sql
 - bids_data_model.sql
 - offerings_issuers_model.sql
- sources
 - .gitignore
 - README.md
 - rill.yaml

CONNECTORS

- duckdb OLAP
 - main_db
 - main
 - ISSUERS
 - OFFERINGS
 - RECIPIENTS
 - RELATEDPERSONS
 - SIGNATURES
 - auction_data_mo...
 - auction_data_raw
 - bids_data_model
 - bids_data_raw
 - offerings_issuers...

```
1 # Visit https://docs.rilldata.com/reference/project-files to learn more about Rill project files.
2
3 title: "Programmatic Ads Auction"
4 model: "auction_data_model"
5 type: metrics_view
6 timeseries: "_time"
7 measures:
8   - label: Requests
9     name: requests
10    expression: sum(bid_request_cnt)
11    description: Total Requests
12    format_preset: humanize
13   - label: "Avg Bid Floor"
14     name: avg_bid_floor
15     expression: "sum(bid_floor) / sum(has_bid_floor_cnt) "
16     description: "Average Bid Floor"
17     format_preset: currency_usd
18   - label: "1D QPS"
19     name: 1d_qps
20     expression: "sum(bid_request_cnt) / 86400 "
21     description: "1D QPS"
22     format_preset: humanize
23 dimensions:
24   - label: App Site Name
25     expression: app_site_name
26     description: ""
27   - label: App Site Domain
28     expression: app_site_domain
29     description: ""
30   - label: Pub Name
31     expression: pub_name
32     description: ""
33   - label: Ad Size
34     expression: ad_size
35     description: ""
36 smallest_time_grain: "hour"
37 default_time_range: "P1W"
38 available_time_zones:
39   - America/Los_Angeles
40   - America/Chicago
41   - America/New_York
42   - Europe/London
43   - Europe/Paris
44   - Asia/Jerusalem
45   - Europe/Moscow
46   - Asia/Kolkata
47   - Asia/Shanghai
48   - Asia/Tokyo
49   - Australia/Sydney
50
```

Preview Deploy to share

duckdb 1,391,572 rows
auction_data_model 22 columns

Table columns

sort by type show summary

Column	Value	Percentage
_time		0%
has_bid_floor_cnt		0%
bid_request_cnt		0%
bid_floor		0%
app_site_name	19.5k	0%
app_site_domain	15.5k	0%
pub_name	7.05k	0%
app_site_cat	2.08k	11%
ad_size	532	0%
device_osv	110	0%
bid_floor_bucket	63	0%
device_state	54	0%
platform_browser	39	0%
device_os	28	0%
device_type	9	0%
ad_position	7	0%
video_max_duration_bucket	7	0%
video_min_duration_bucket	7	0%
placement_type	5	0%
device_country	4	0%
auction_type	3	0%
app_or_site	2	0%

Why We Chose DuckDB

Speed

Profiling 10s of GBs of datasets with sub-second performance

Simplicity

Lightweight, embeddable, low maintenance overhead

Scalability

Serve 100s of concurrent queries and scales to 100s of GBs

Open source

Permissive license with a vibrant developer community

Untitled Rill Project

bids.yaml

Add

- dashboards
 - auction.yaml
 - bids.yaml
- models
- sources
- .gitignore
- README.md
- rill.yaml

CONNECTORS

- duckdb (OLAP)
 - main_db
 - main
 - auction_data_mo...
 - auction_data_raw
 - bids_data_model
 - bids_data_raw

```

1 # Visit https://docs.rilldata.com/reference/project-files to learn more about Rill project files.
2
3 title: "Programmatic Ads Bids"
4 model: "bids_data_model"
5 type: metrics_view
6 timeseries: "__time"
7 smallest_time_grain: "hour"
8 default_time_range: "PT24H"
9 measures:
10   - label: Advertising Spend Overall
11     name: overall_spend
12     expression: sum(media_spend_usd)/1000
13     description: Total Spend
14     format_preset: currency_usd
15   - label: "Bids"
16     name: total_bids
17     expression: "sum(bid_cnt)"
18     description: "Total Bids"
19     format_preset: humanize
20   - label: "Impressions"
21     name: impressions
22     expression: "sum(imp_cnt)"
23     description: "Total Impressions"
24     format_preset: humanize
25   - label: "Win Rate"
26     name: win_rate
27     expression: "sum(imp_cnt)*1.0/sum(bid_cnt)"
28     description: "Impressions / Bids"
29     format_preset: percentage
30   - label: "Clicks"
31     name: clicks
32     expression: "sum(click_reg_cnt)"
33     description: "Total Clicks"
34     format_preset: humanize
35   - label: "CTR"
36     name: ctr
37     expression: "sum(click_reg_cnt)*1.0/sum(imp_cnt)"
38     description: "Click Through Rate"
39     format_preset: percentage
40   - label: "Video Starts"
41     name: video_starts
42     expression: "sum(video_start_cnt)"
43     description: "Total Video Starts"
44     format_preset: humanize
45   - label: "Video Completes"
46     name: video_completes
47     expression: "sum(video_complete_cnt)"
48     description: "Total Video Completes"
49     format_preset: humanize
50   - label: "Video Completion Rate"
51     name: video_completion_rate

```

Auto-save

Report an issue | version unknown (built from source)

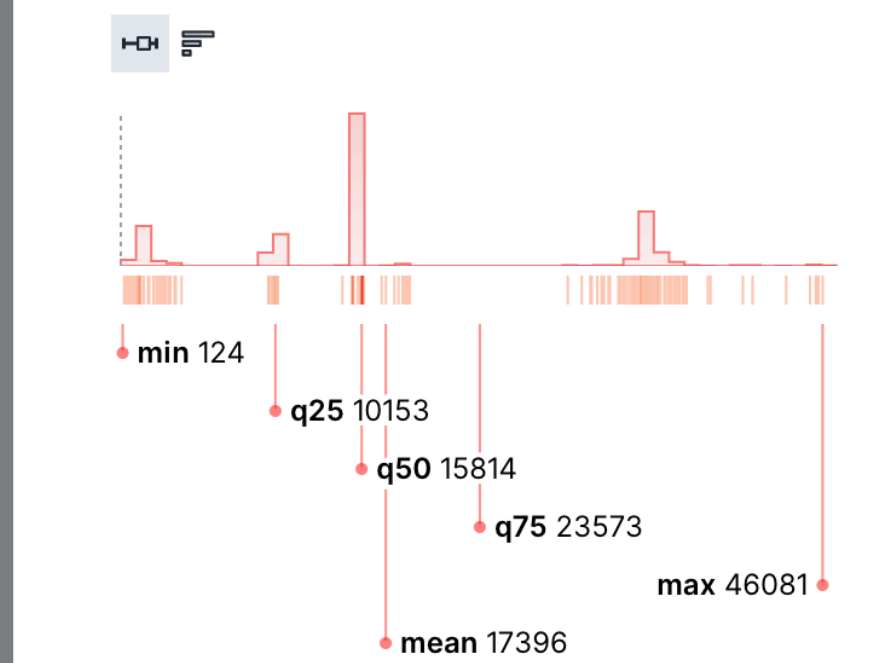
Preview | Deploy to share

duckdb 1,569,315 rows | bids_data_model 49 columns

Table columns

sort by type | show summary

__time	0%
winning_bid_price_usd	0%
vpaid	0%
video_third_quartile_cnt	0%
video_start_cnt	0%
video_skip_cnt	0%
video_midpoint_cnt	0%
video_first_quartile_cnt	0%
video_complete_cnt	0%
video_activity_cnt	0%
mraid	0%
media_spend_usd	0%
line_item_id	0%



imp_cnt	0%	
has_imp	0%	
has_bid_floor_cnt	0%	
click_vid_comp_cnt	0%	
click_reg_cnt	0%	
click_cnt	0%	
bid_price_usd	0%	
bid_cnt	0%	
advertiser_id	0%	
# bid_floor	0%	
A pub_name	7.46k	0%
A line_item_name	1.53k	0%
A device_region	810	0%
A app_site_name	455	0%
A app_site_domain	215	0%
A campaign_name	210	0%
A adomain	171	0%

Demo


```
$ curl https://rill.sh | sh
```



Copy command at rilldata.com

Thank you, DuckCon!

[@rilldata](#)

rilldata.com