

# Dynamic, partially-stateful data-flow for high-performance Web applications

Jon Gjengset

Malte Schwarzkopf

Jonathan Behrens

Lara Timbó Araújo

Martin Ek


Eddie Kohler

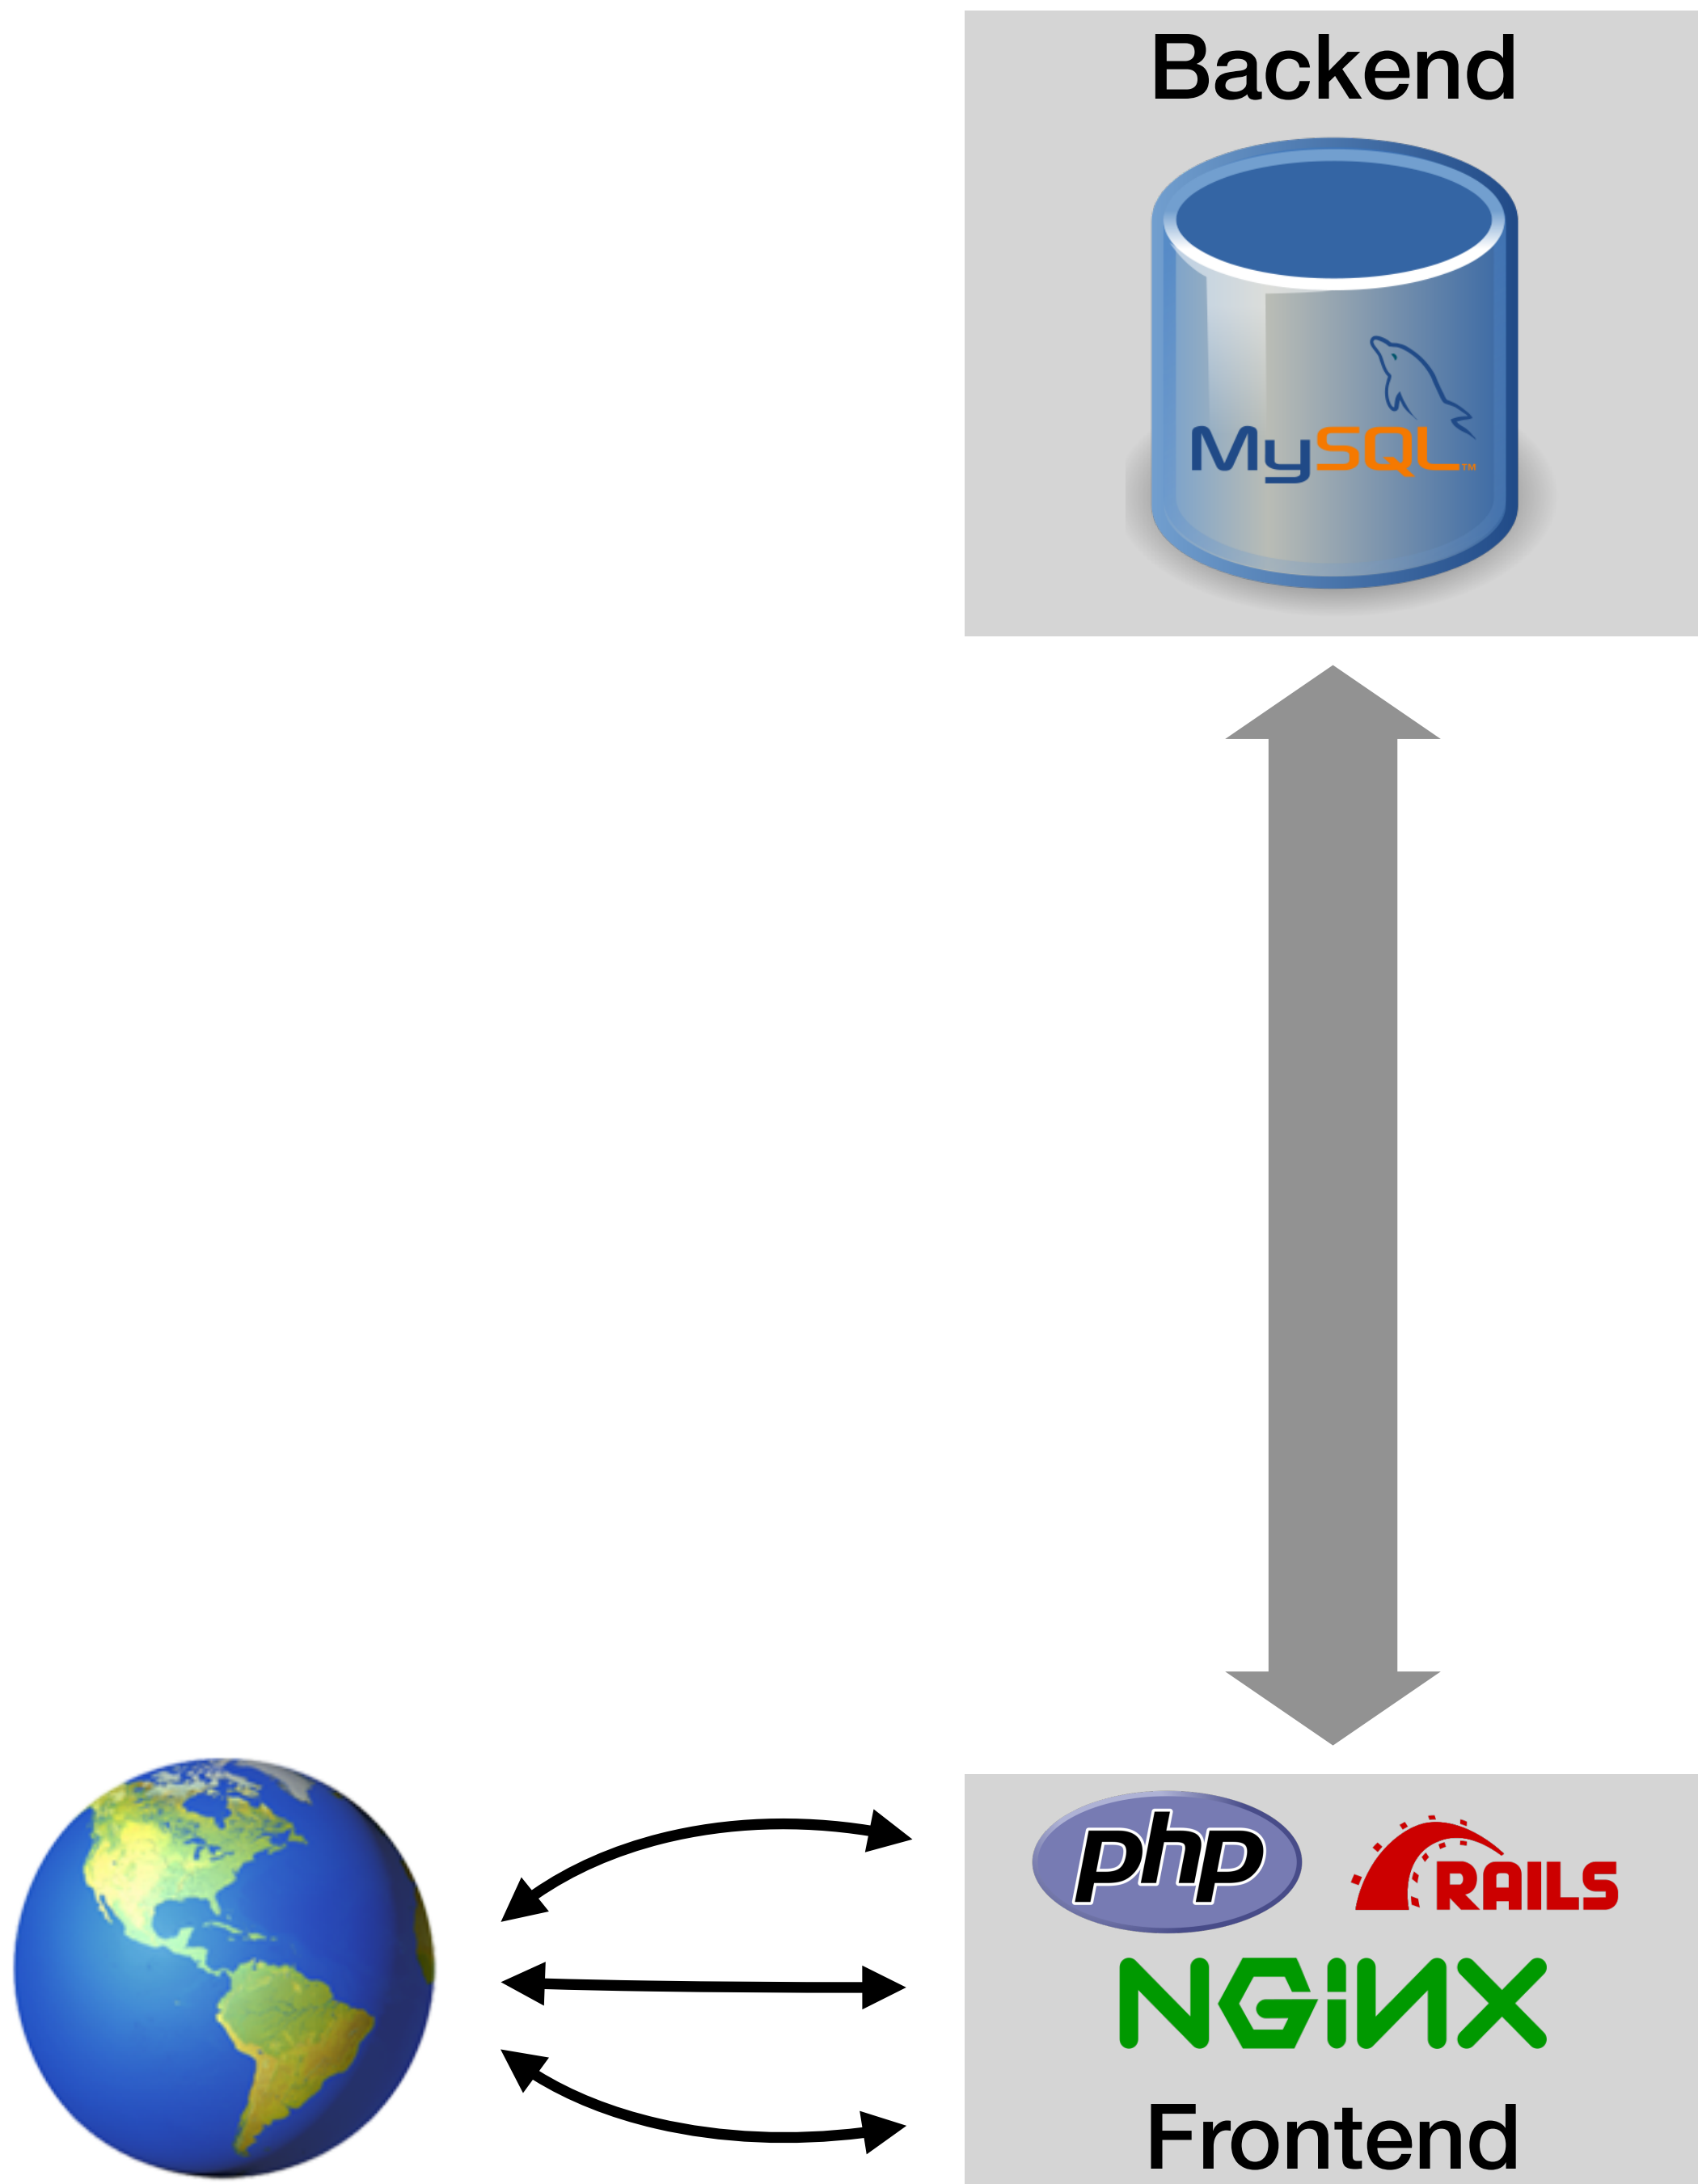
M. Frans Kaashoek

Robert Morris









**Backend**



**Frontend**

# Backend



Stories	
	Orange bar
	Green bar
	Blue bar

Votes	
	Orange, Green, Blue squares
	Blue, Orange squares


# Frontend

# Backend



Stories	
	Yellow bar
	Green bar
	Blue bar

Votes	
	Yellow, Green, Blue
	Blue, Yellow



# Frontend




# Backend



Stories	
	Yellow bar
	Green bar
	Blue bar

Votes	
	Yellow, Green, Blue
	Blue, Yellow



# Frontend

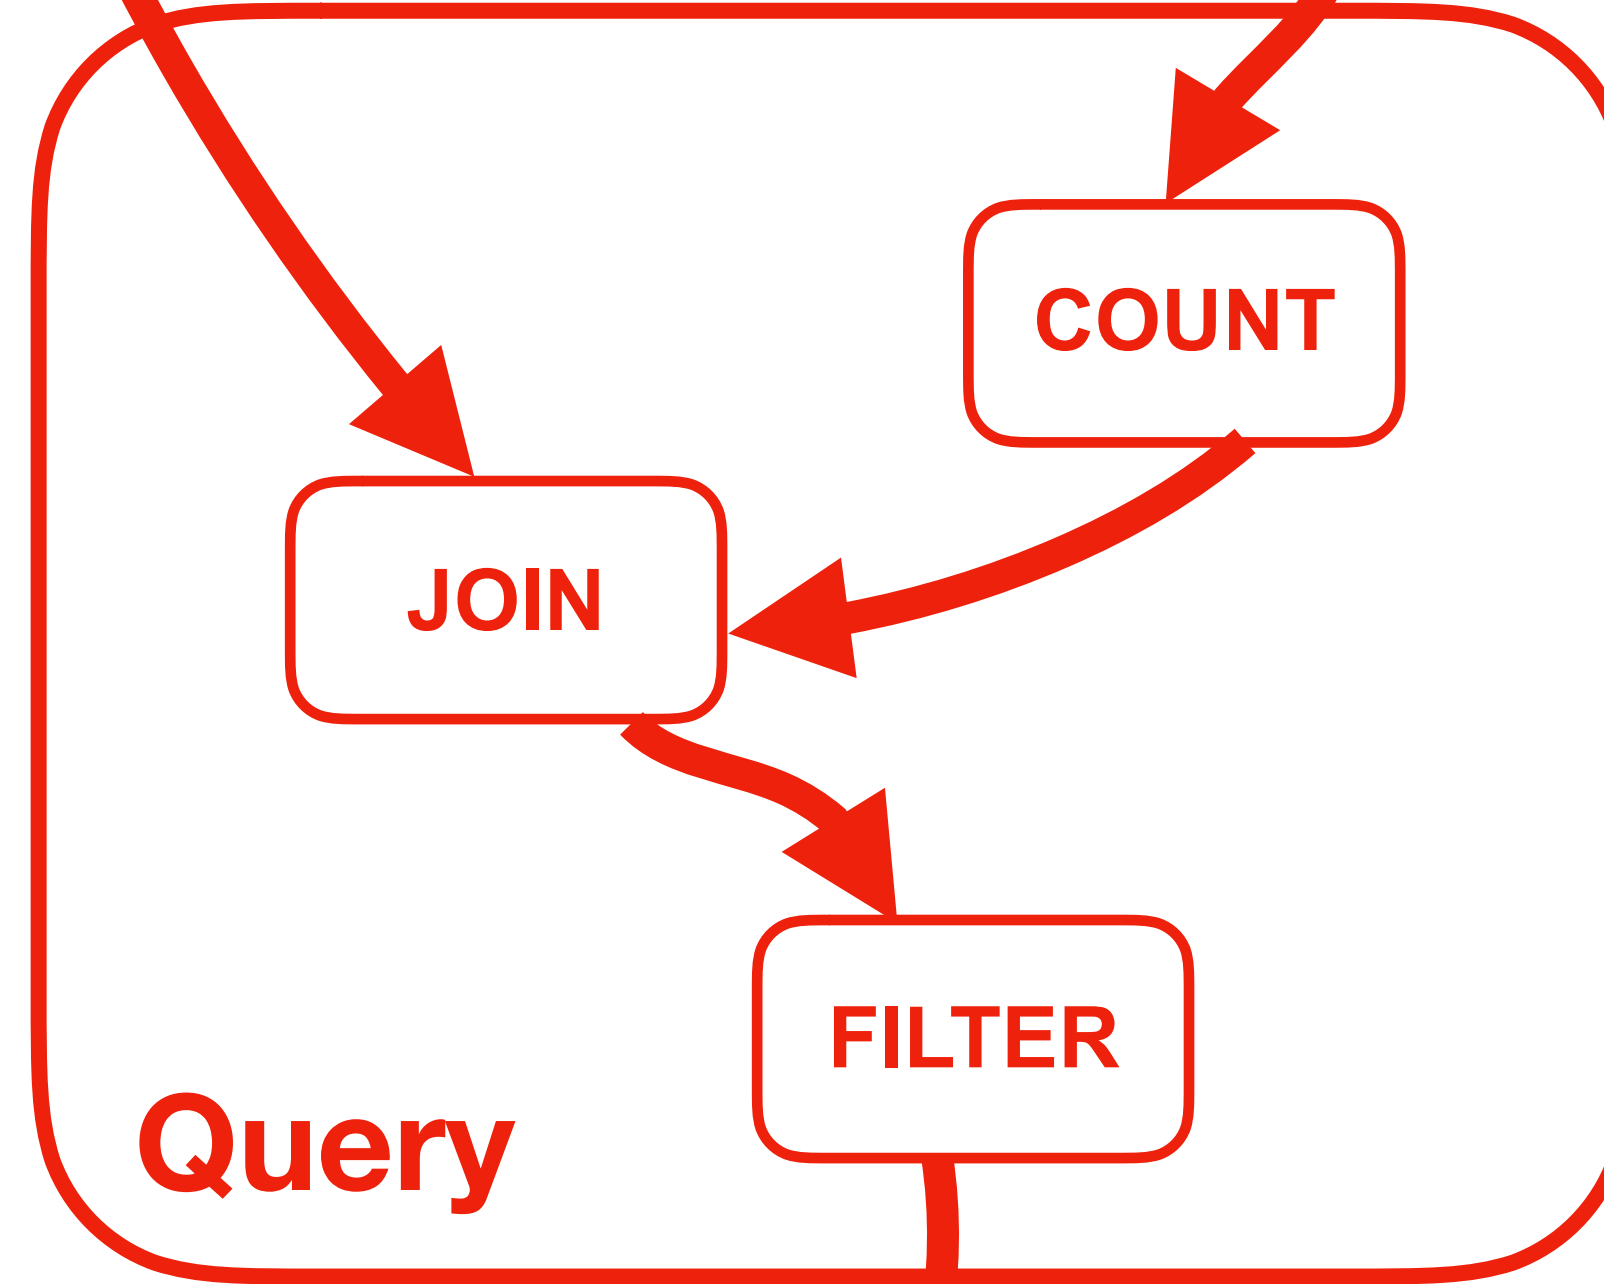
# Backend



Stories	
	Yellow bar
	Green bar
	Blue bar

Votes	
	Yellow, Green, Blue squares
	Blue, Yellow squares

**90% reads**  
10% writes




# Frontend

# Backend



Stories	
	Yellow bar
	Green bar
	Blue bar

Votes	
	Yellow, Green, Blue
	Blue, Yellow

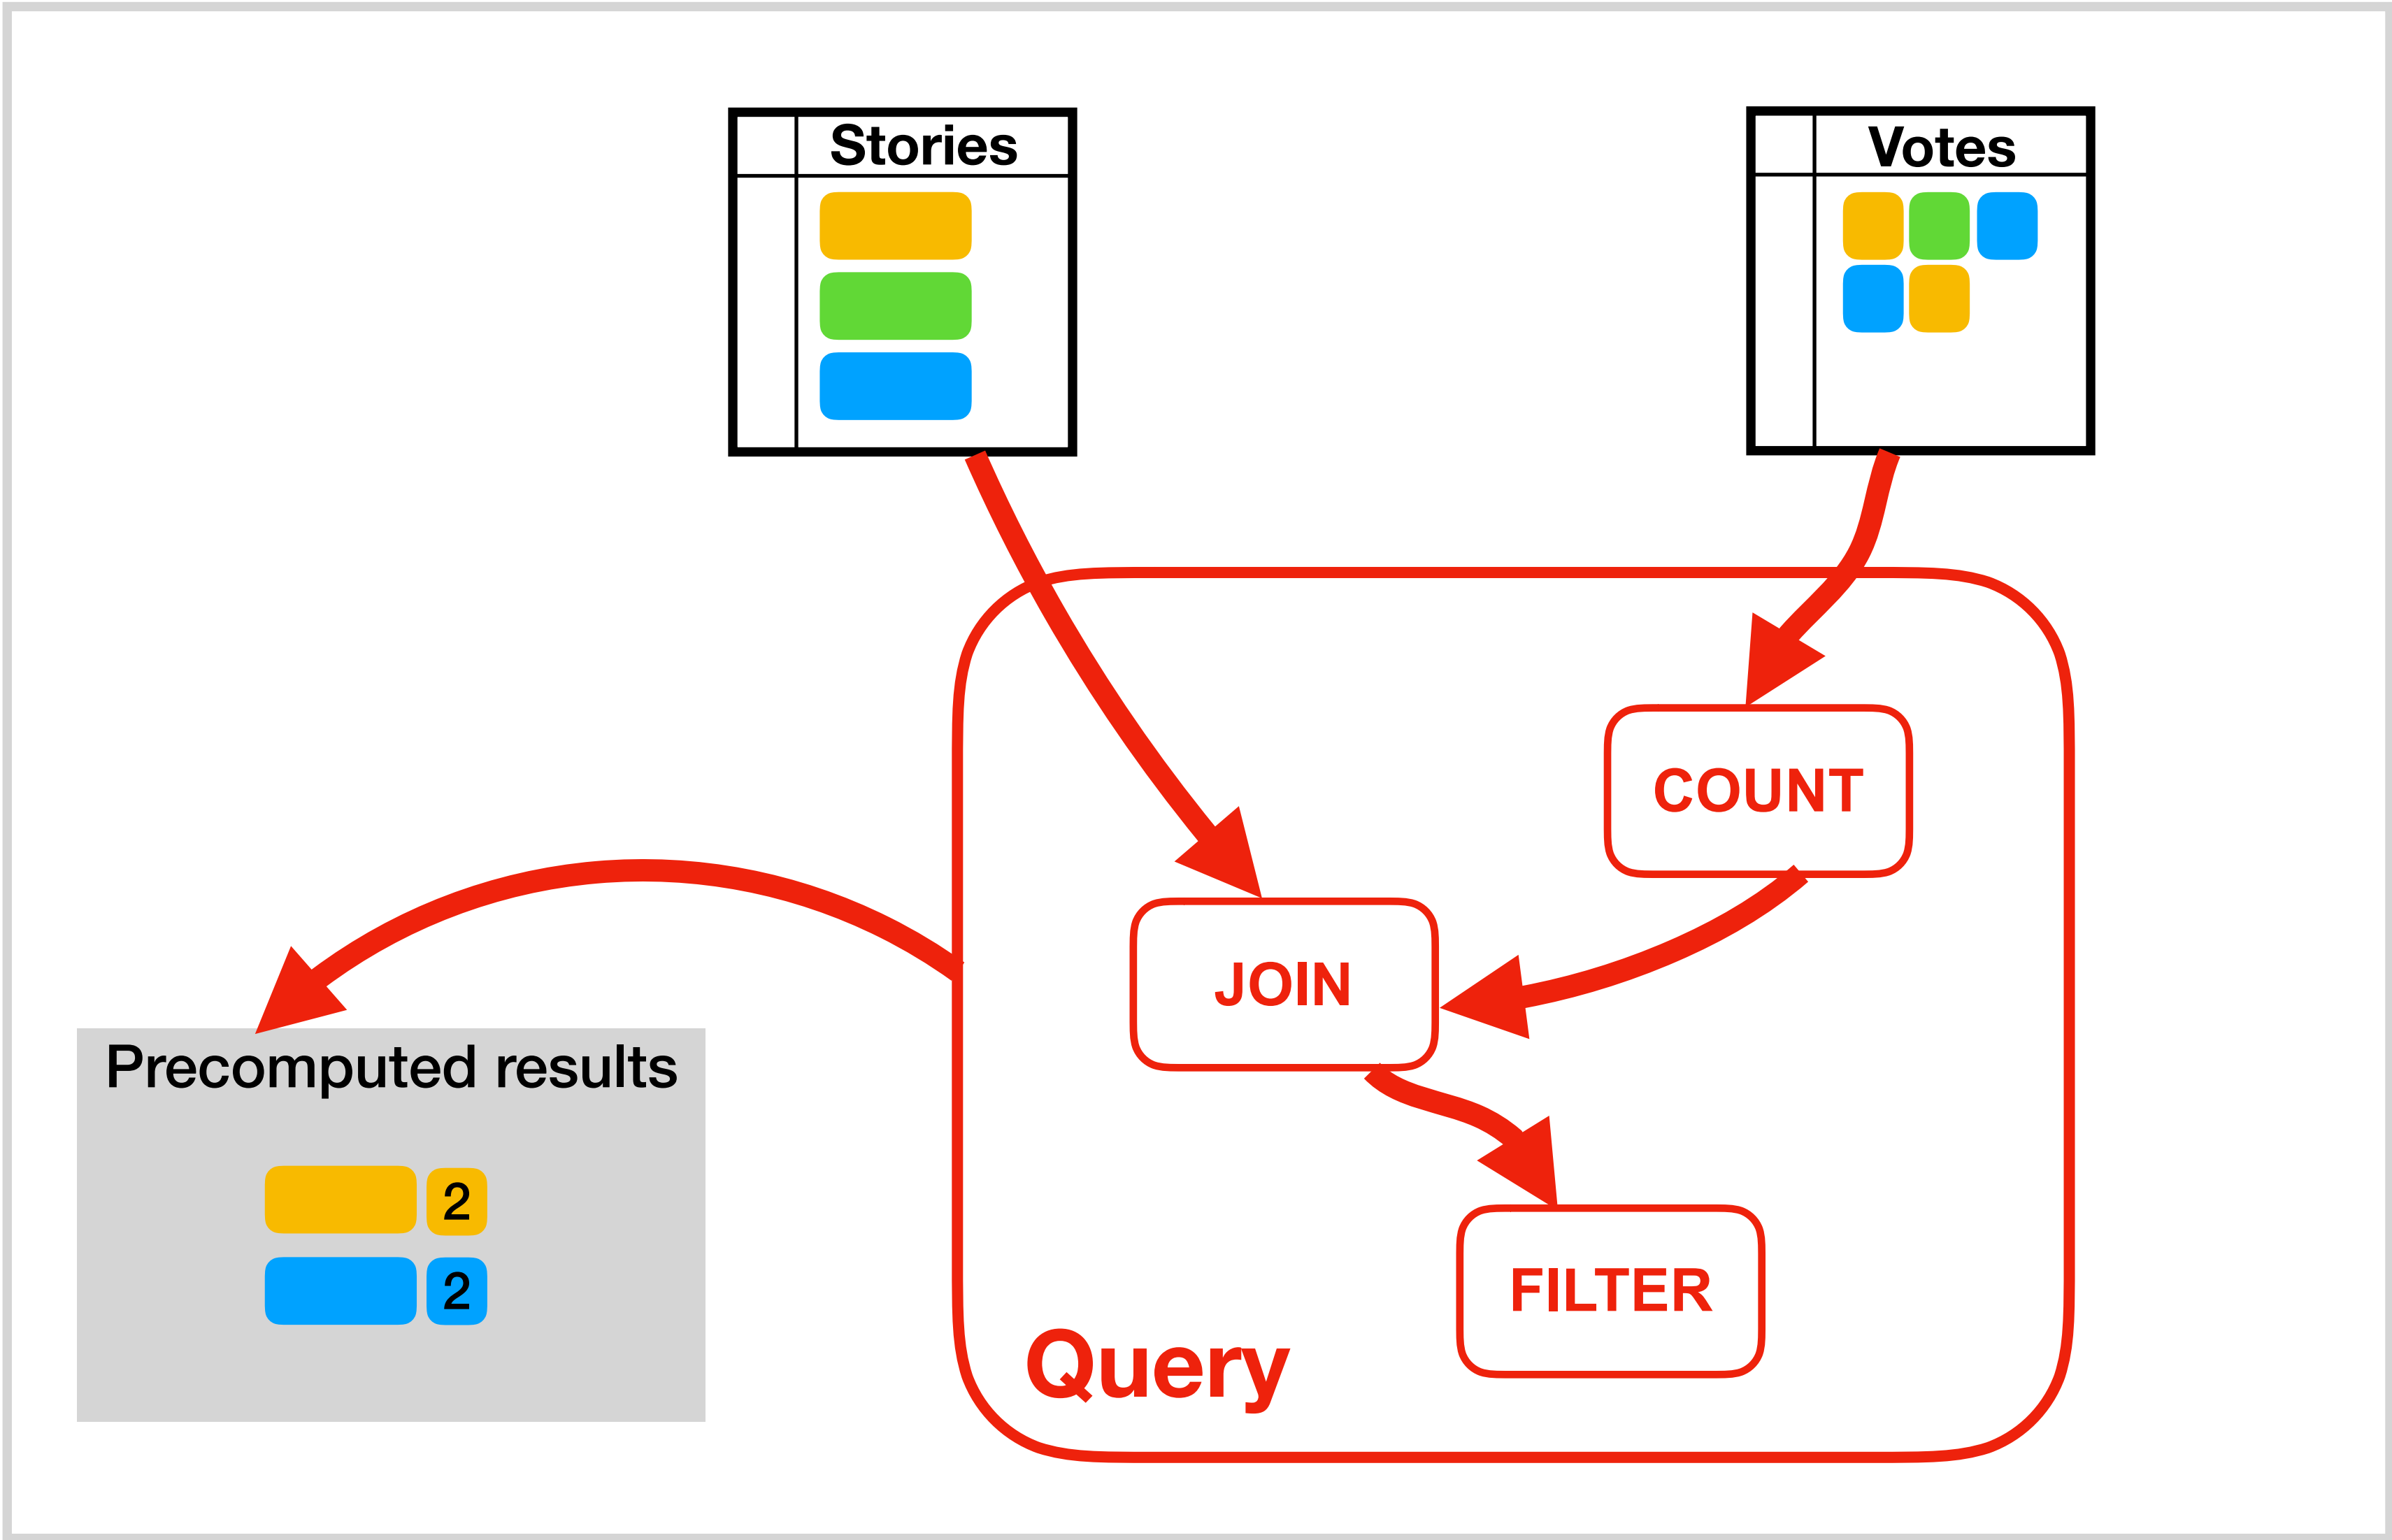


**90% reads**  
10% writes

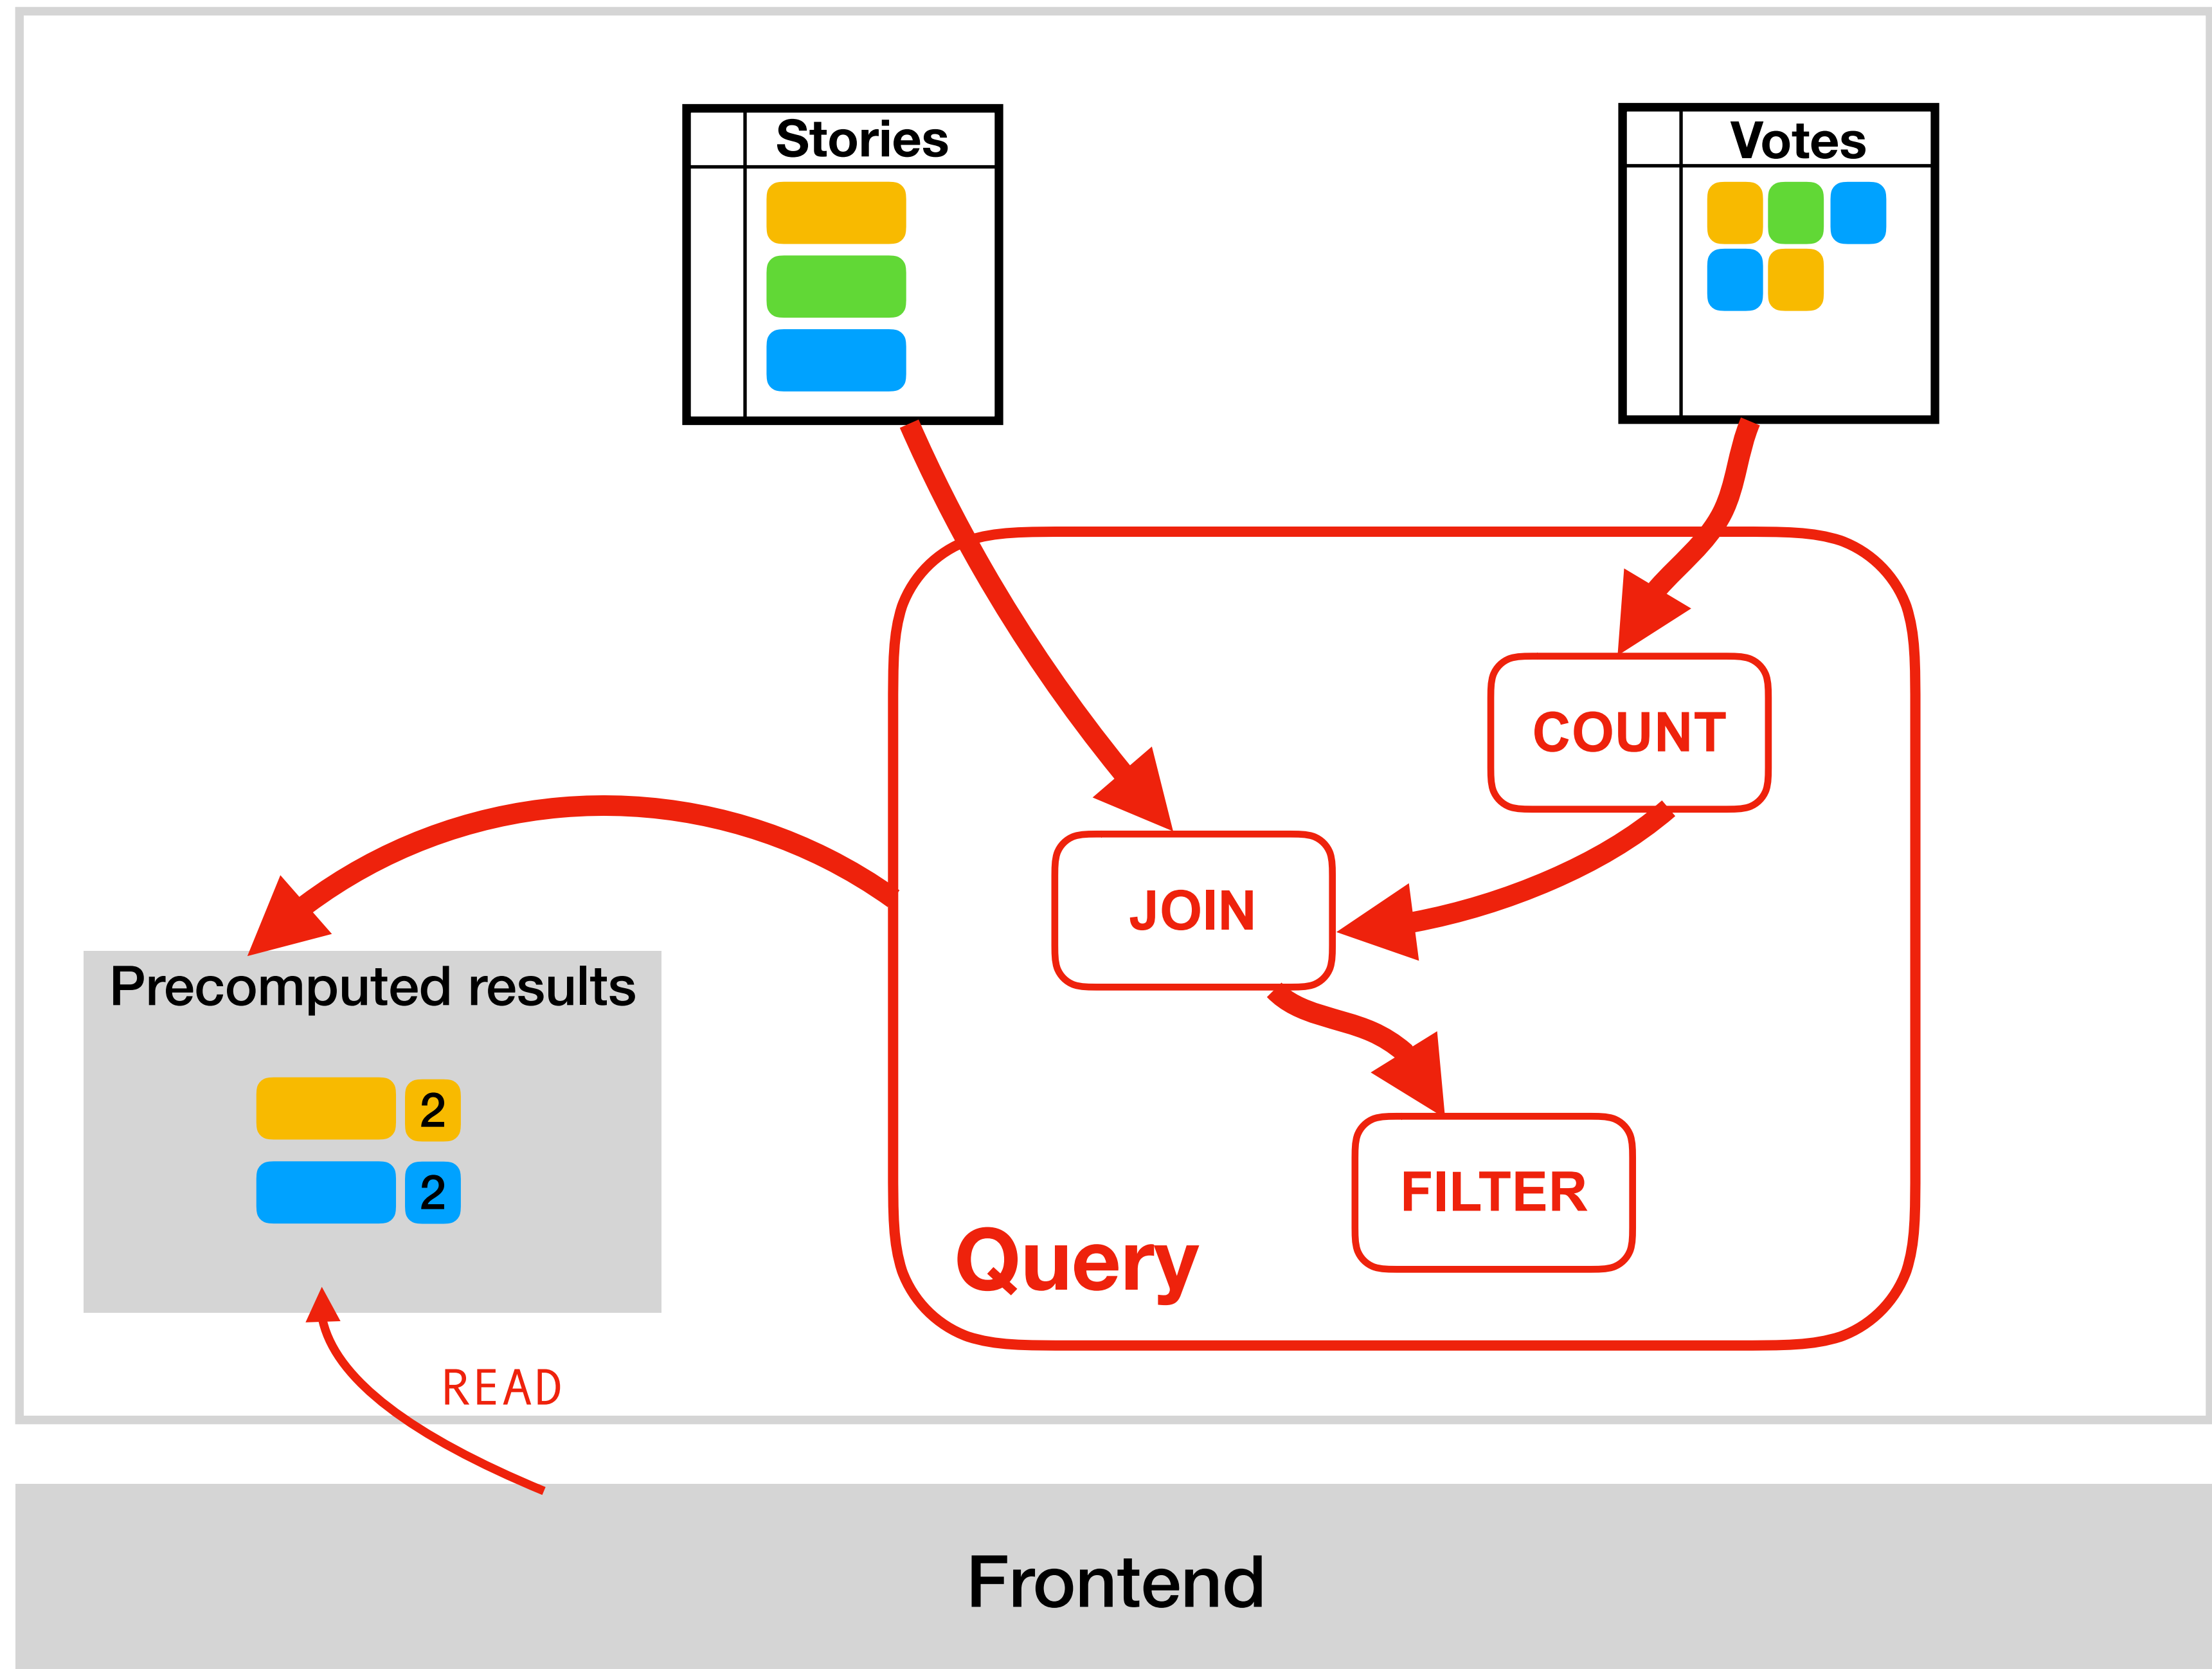
**Slow reads,  
repeated work!**



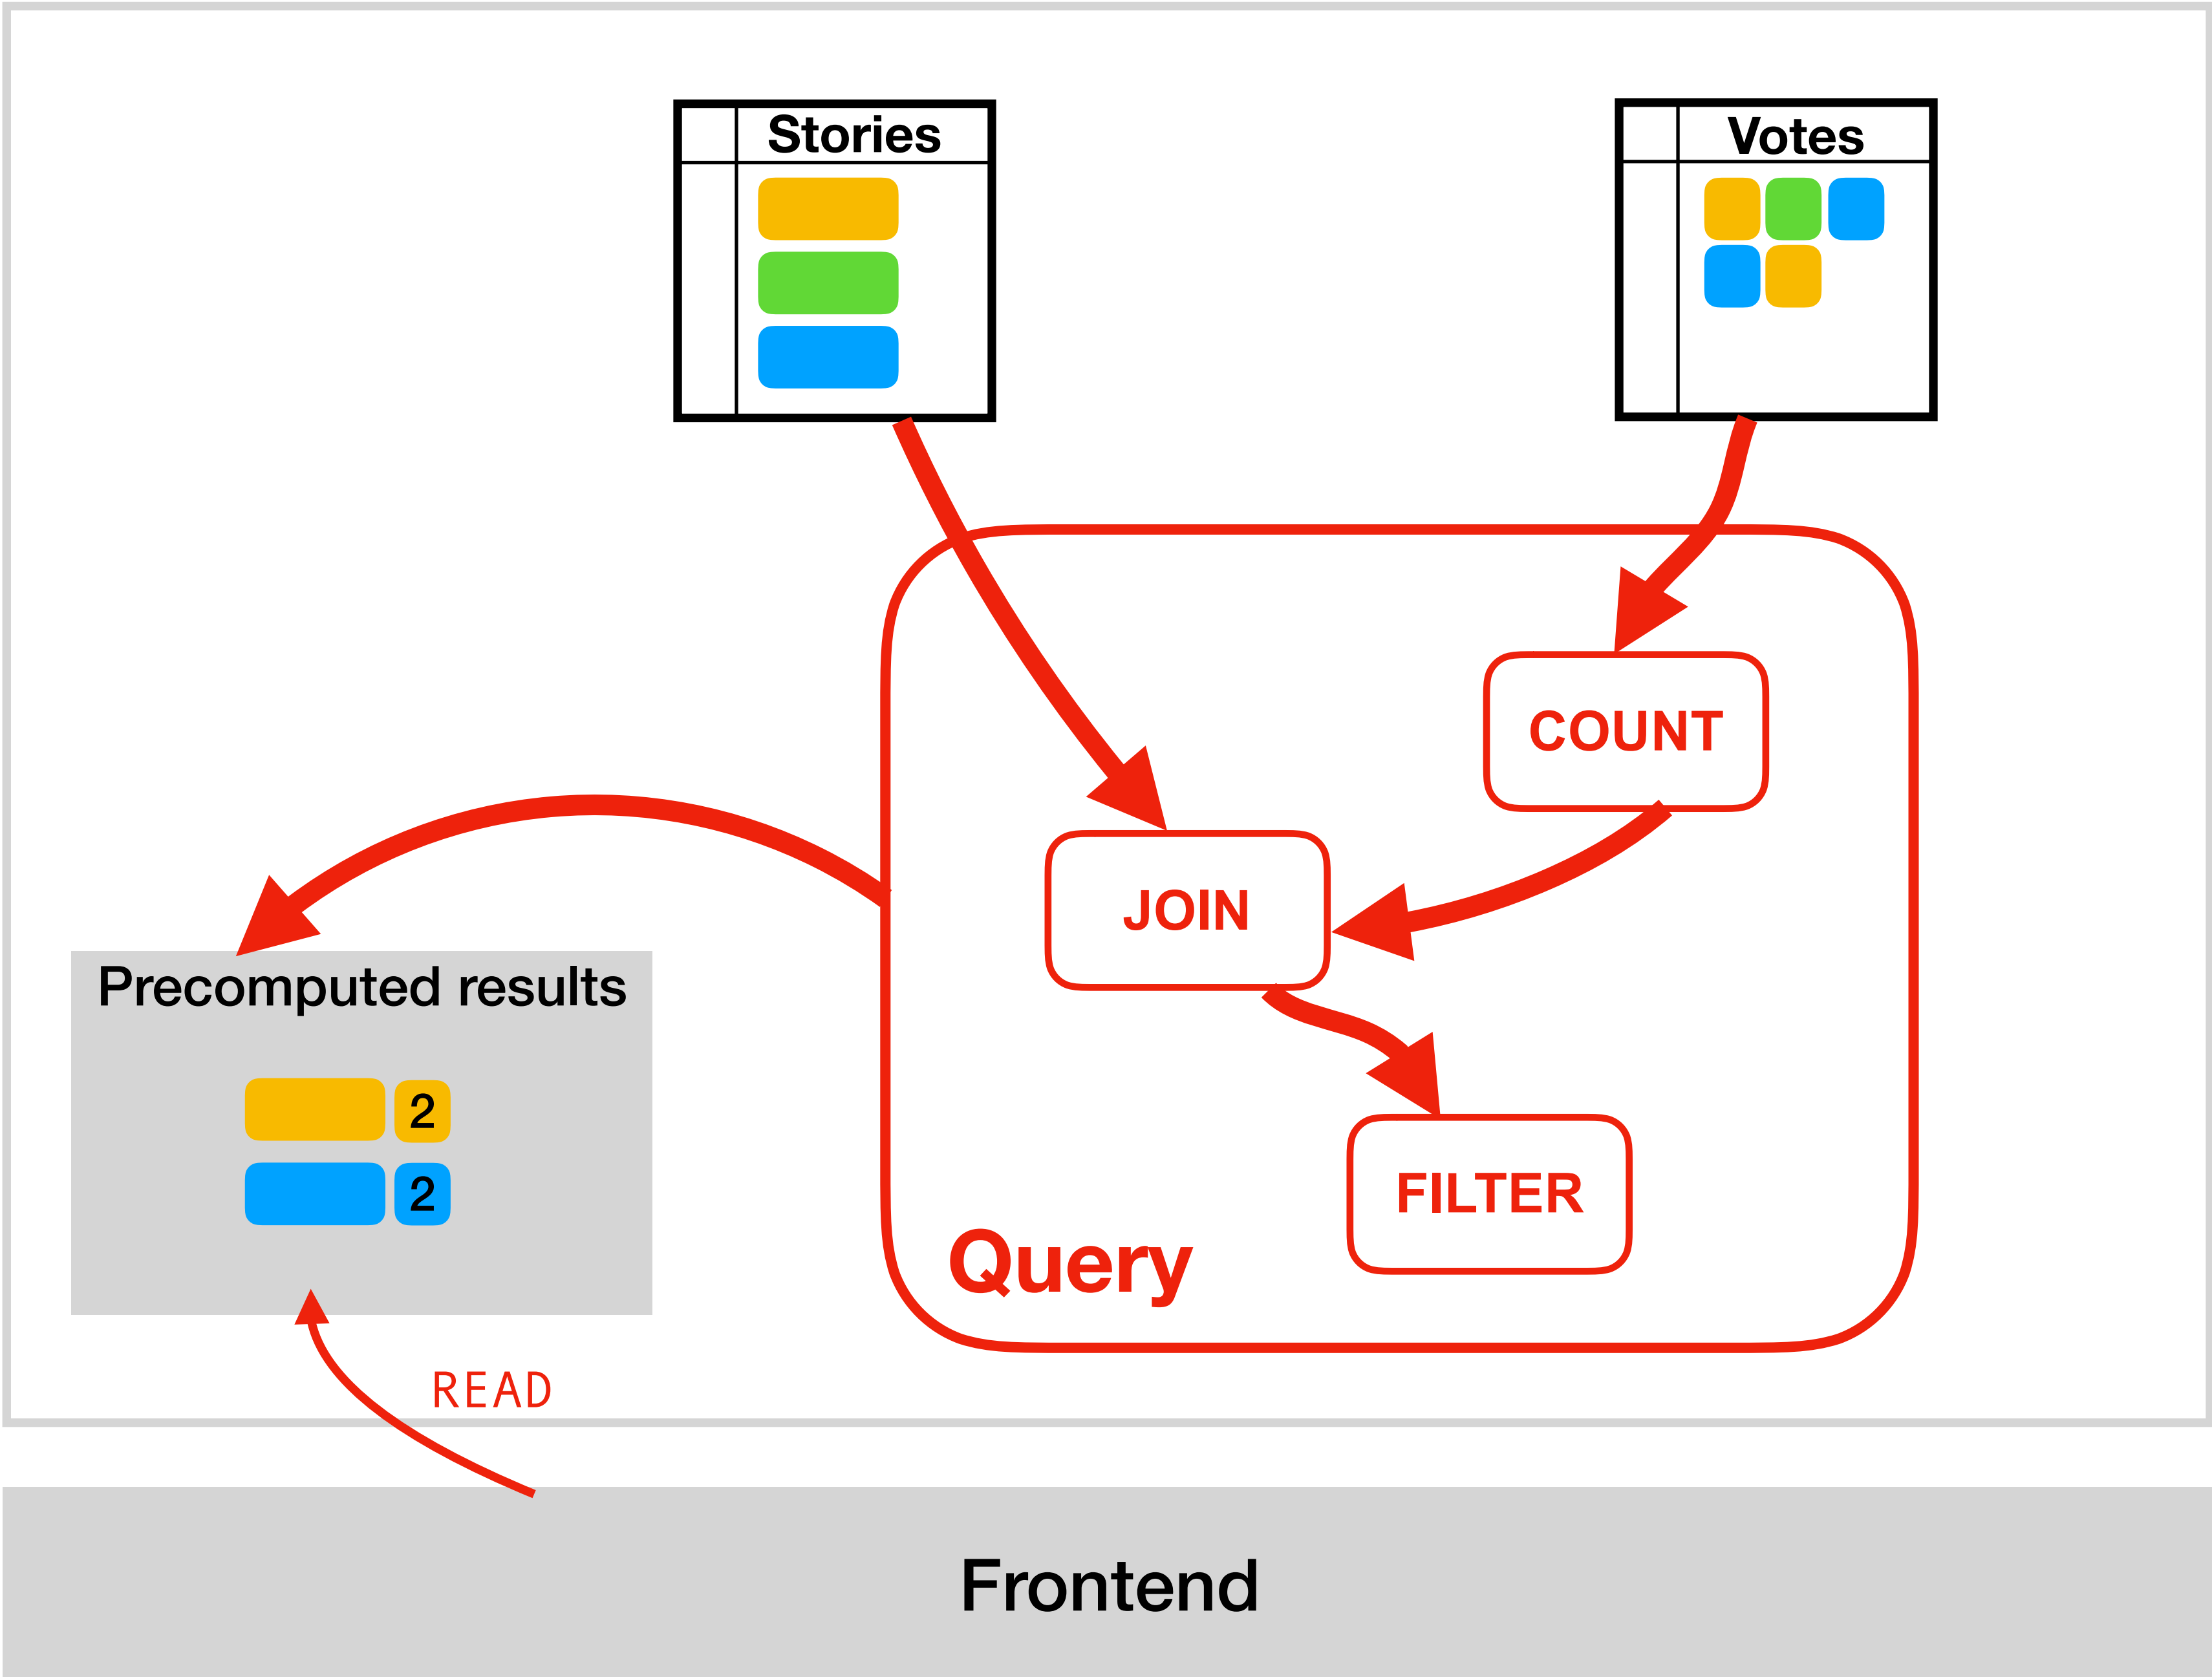
# Frontend



Frontend




Store in base table?  
– manual, slow.



Store in base table?  
– manual, slow.

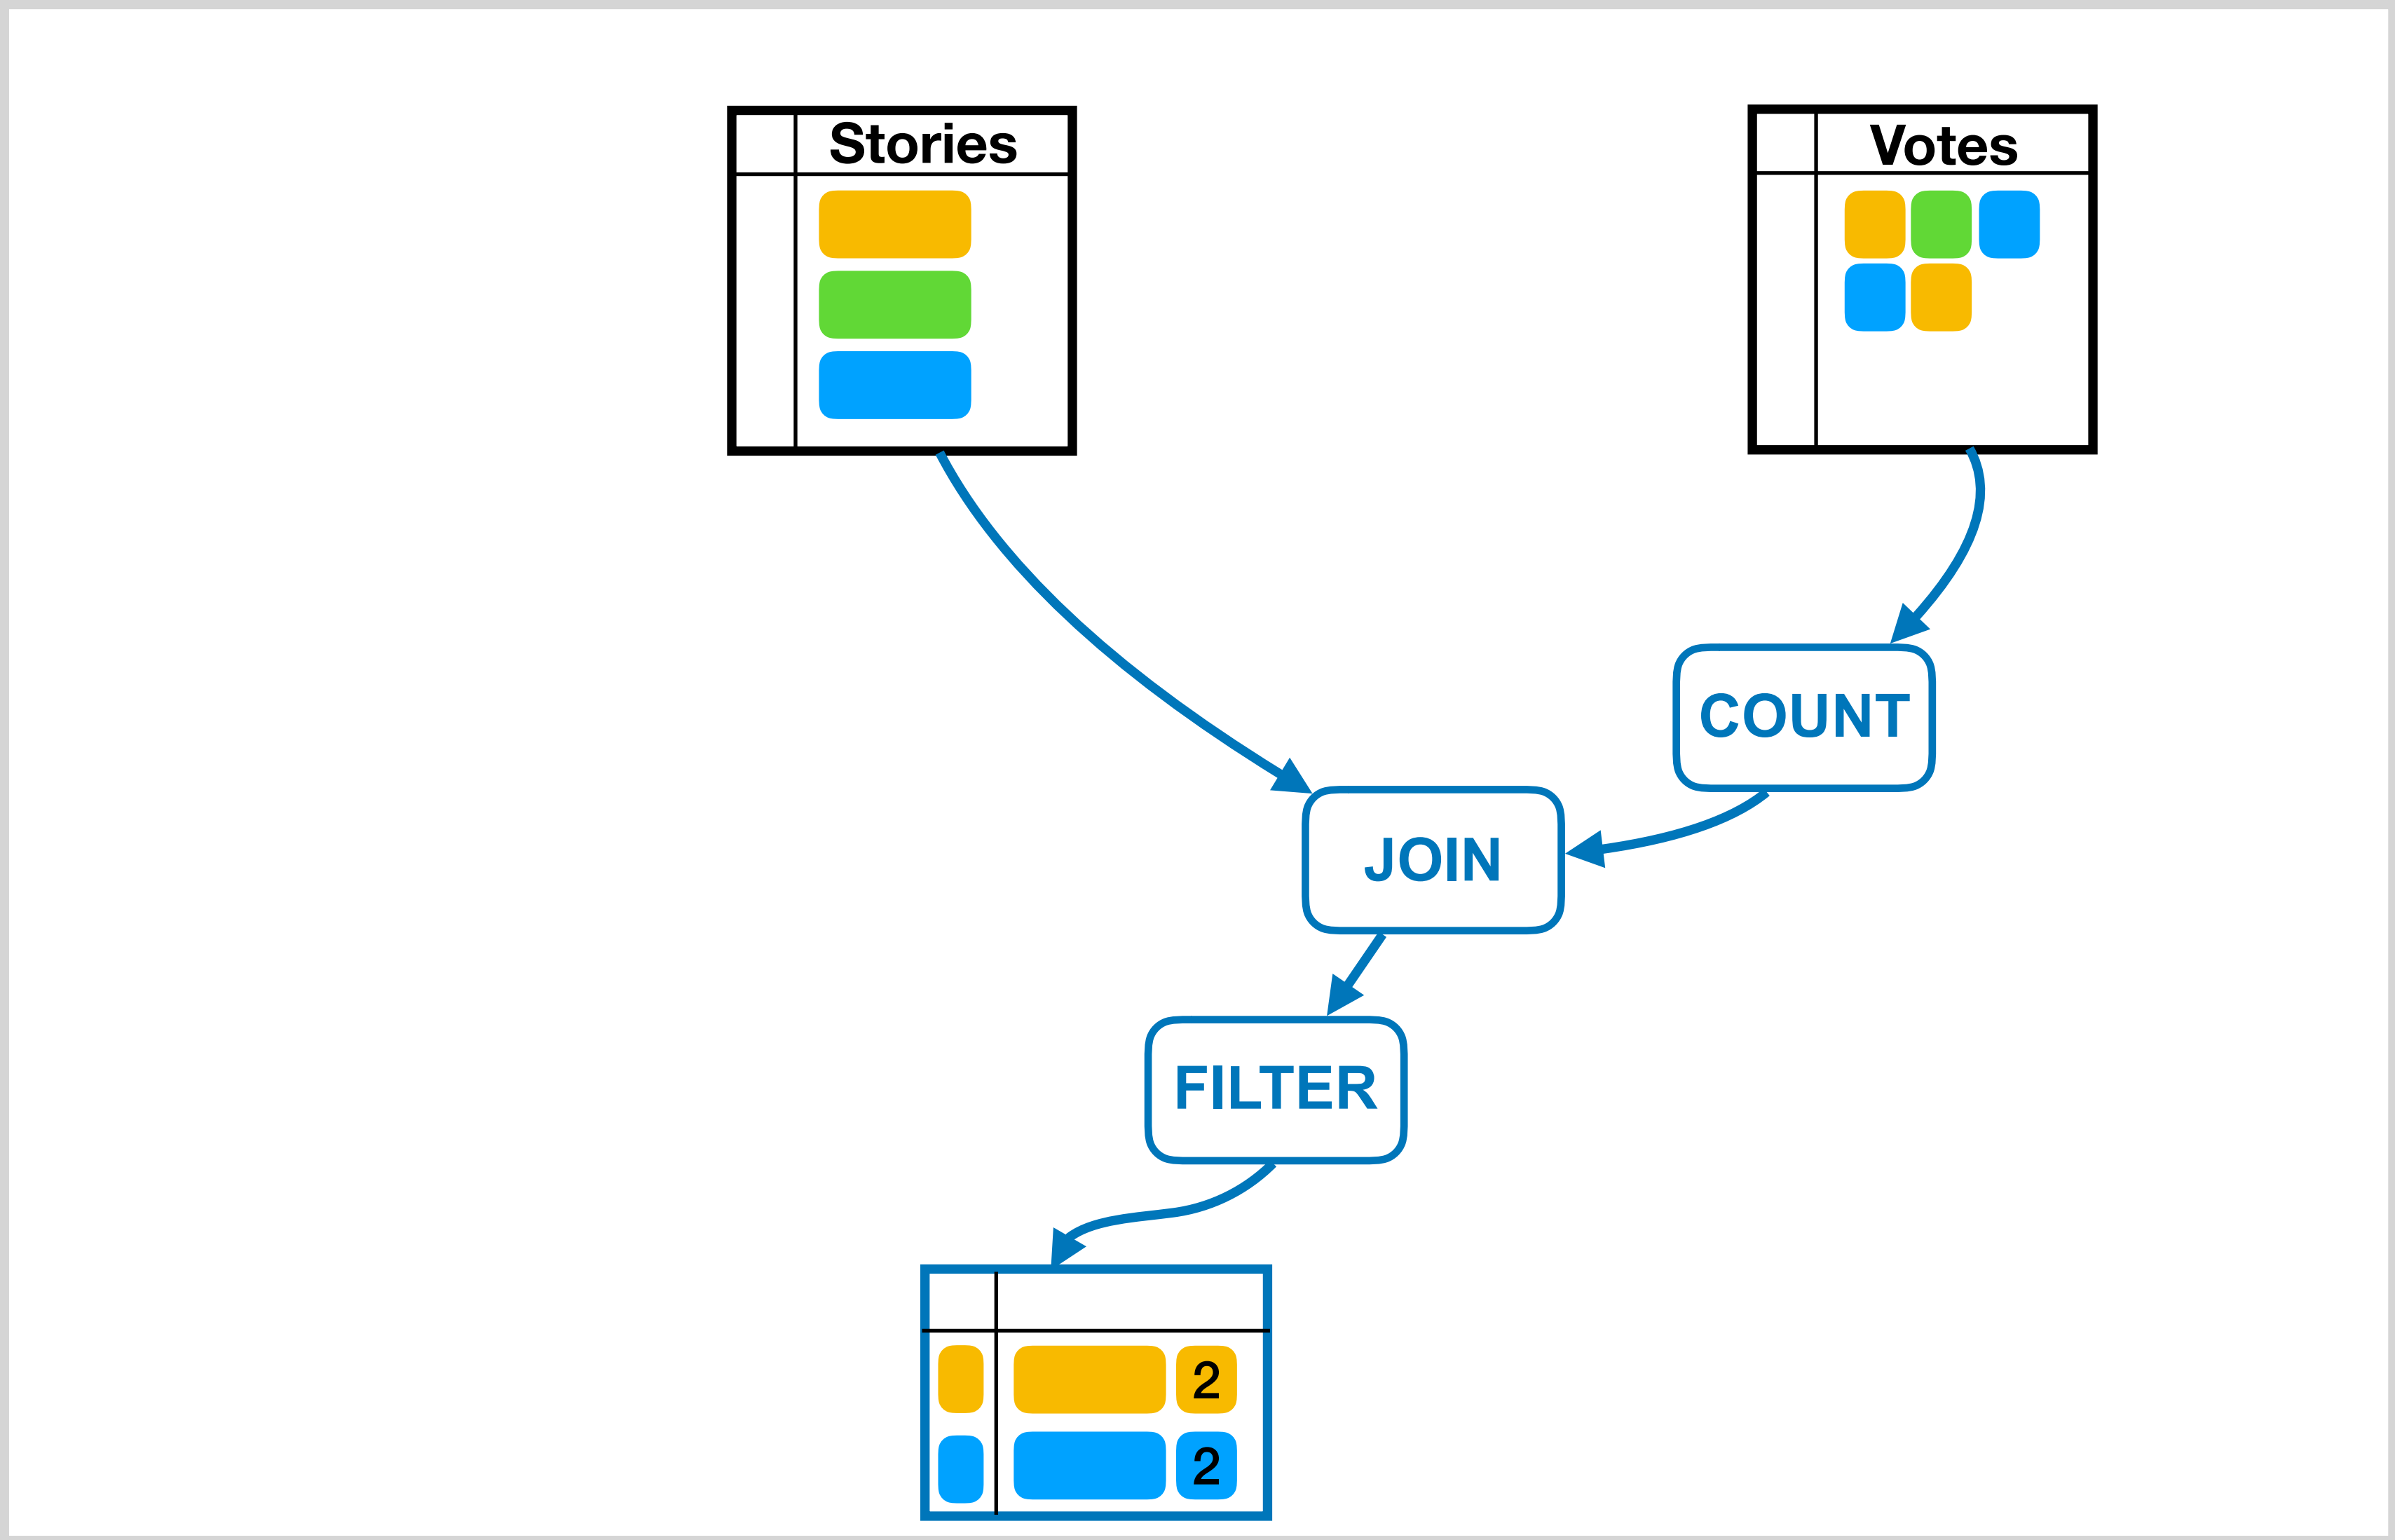
memcached?  
– complex  
[Facebook NSDI'13].



Store in base table?  
– manual, slow.

memcached?  
– complex  
[Facebook NSDI'13].

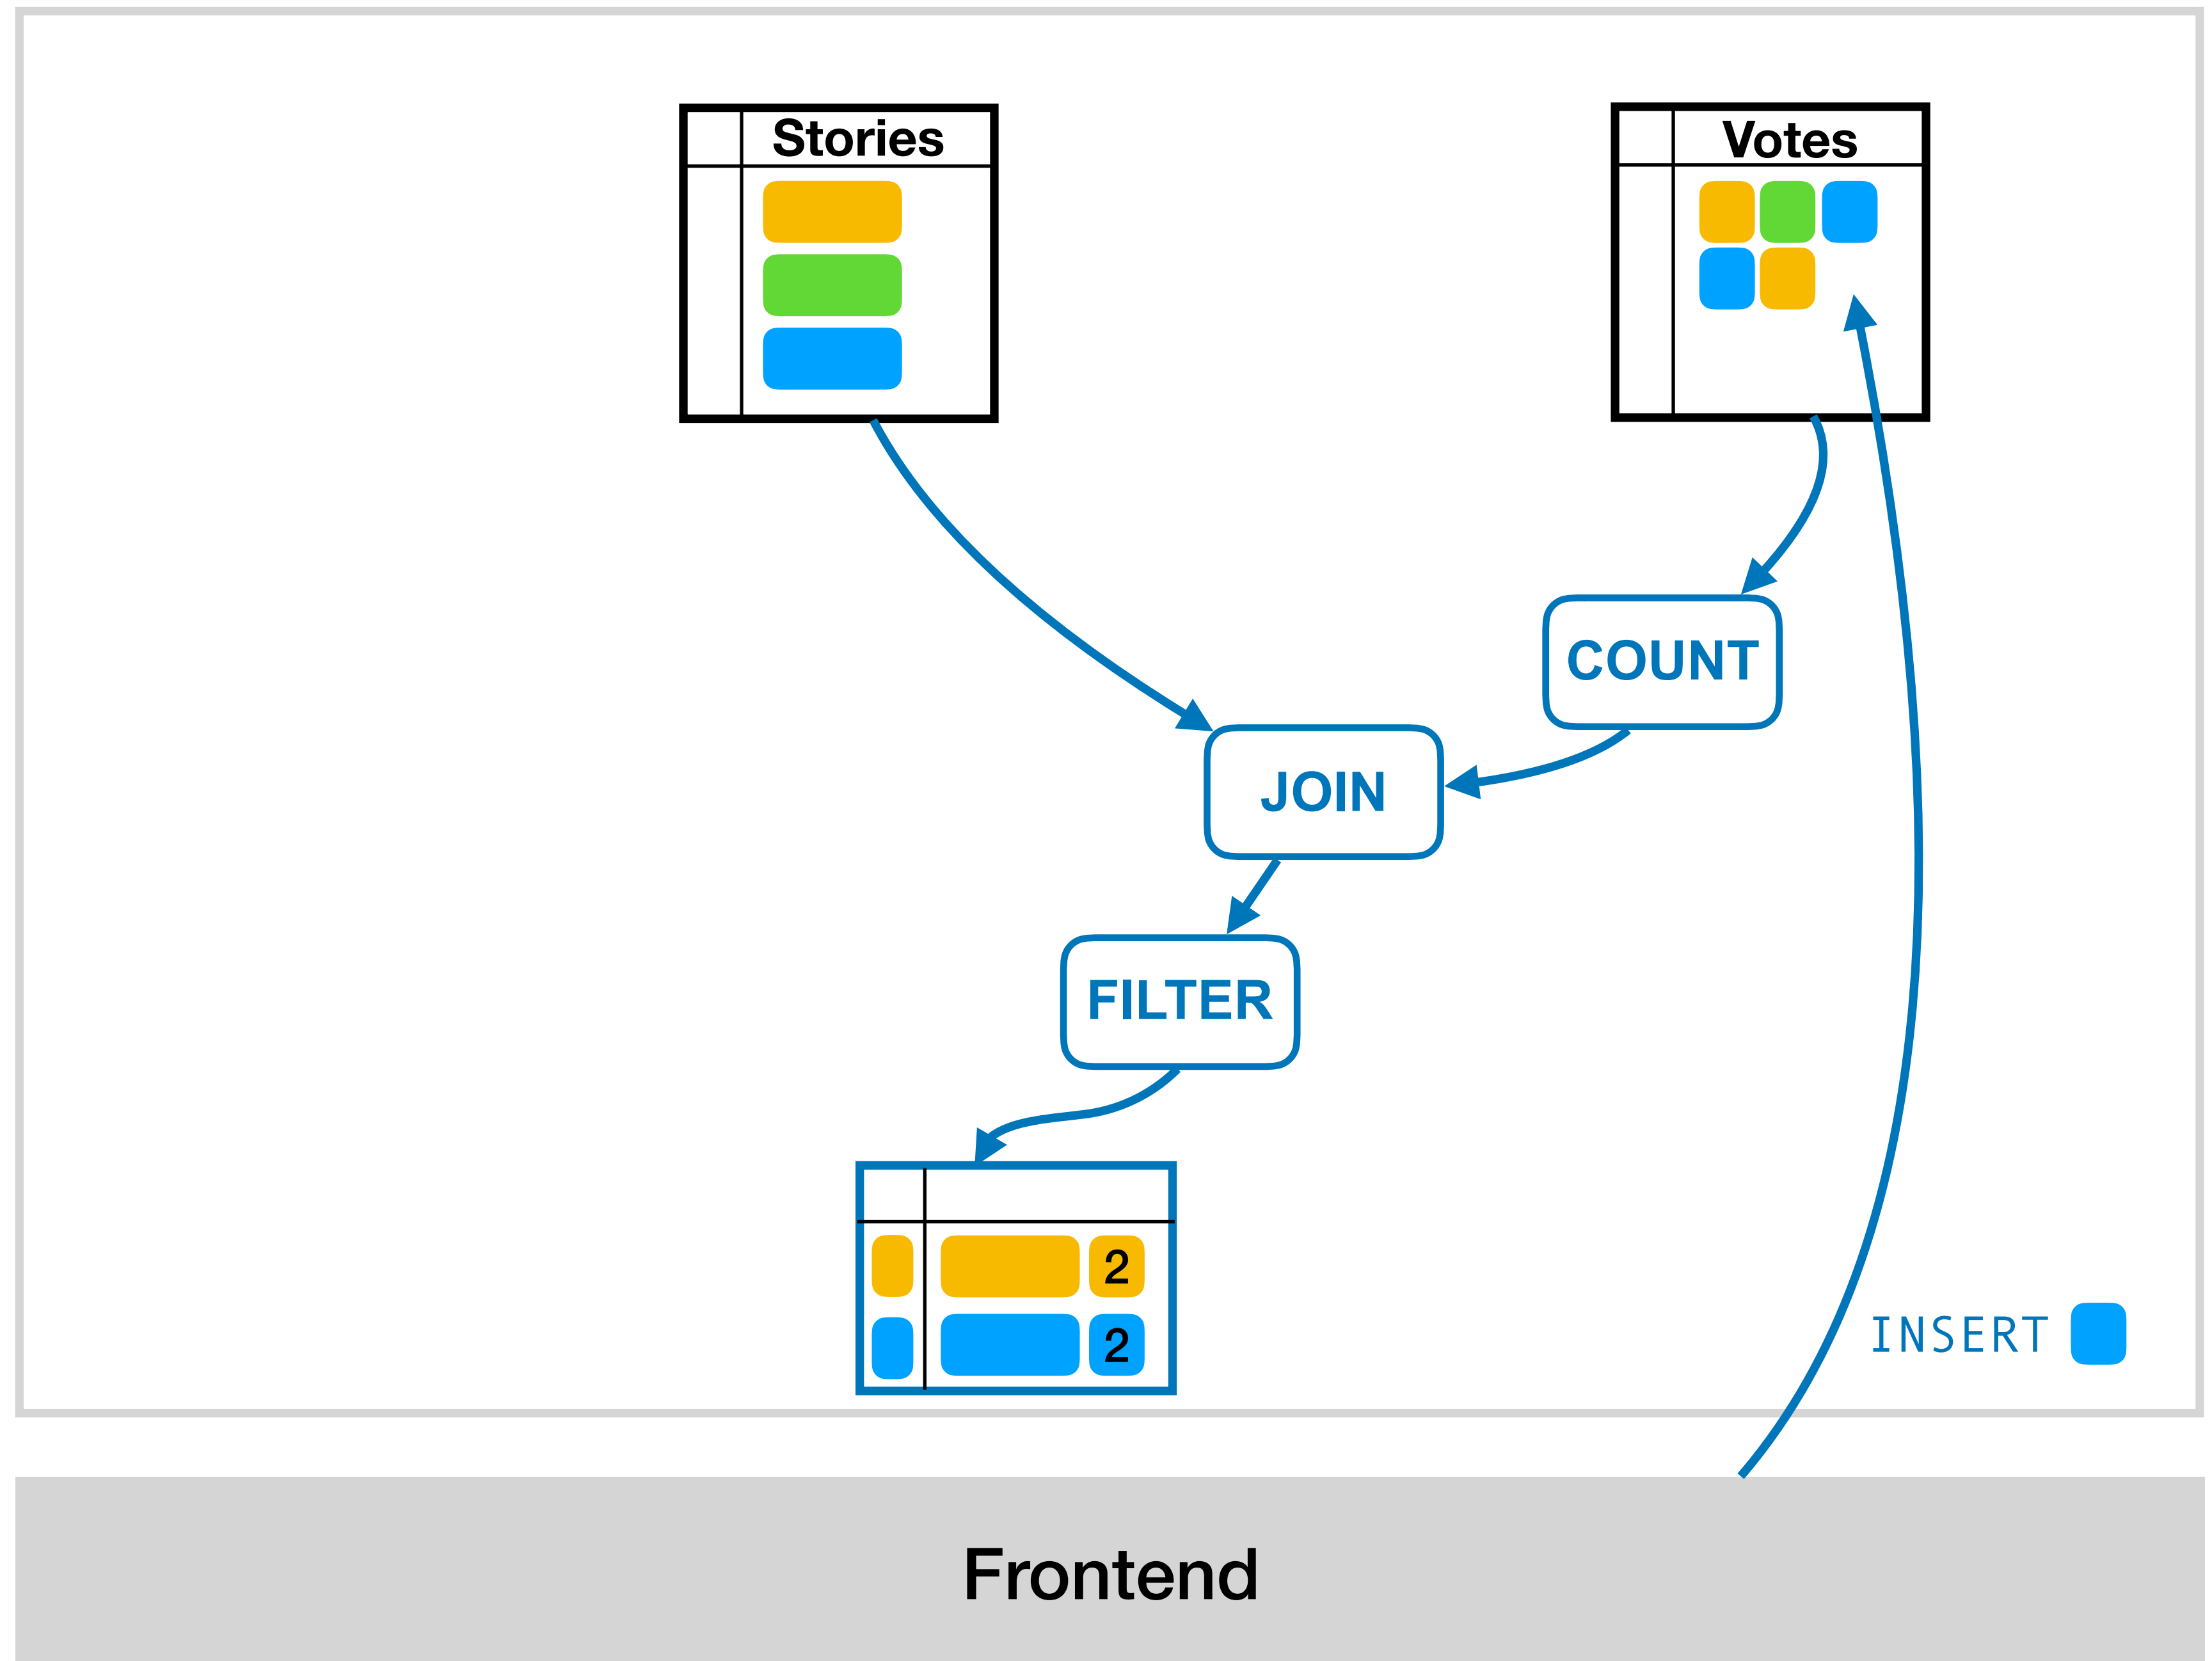
# Streaming data-flow?



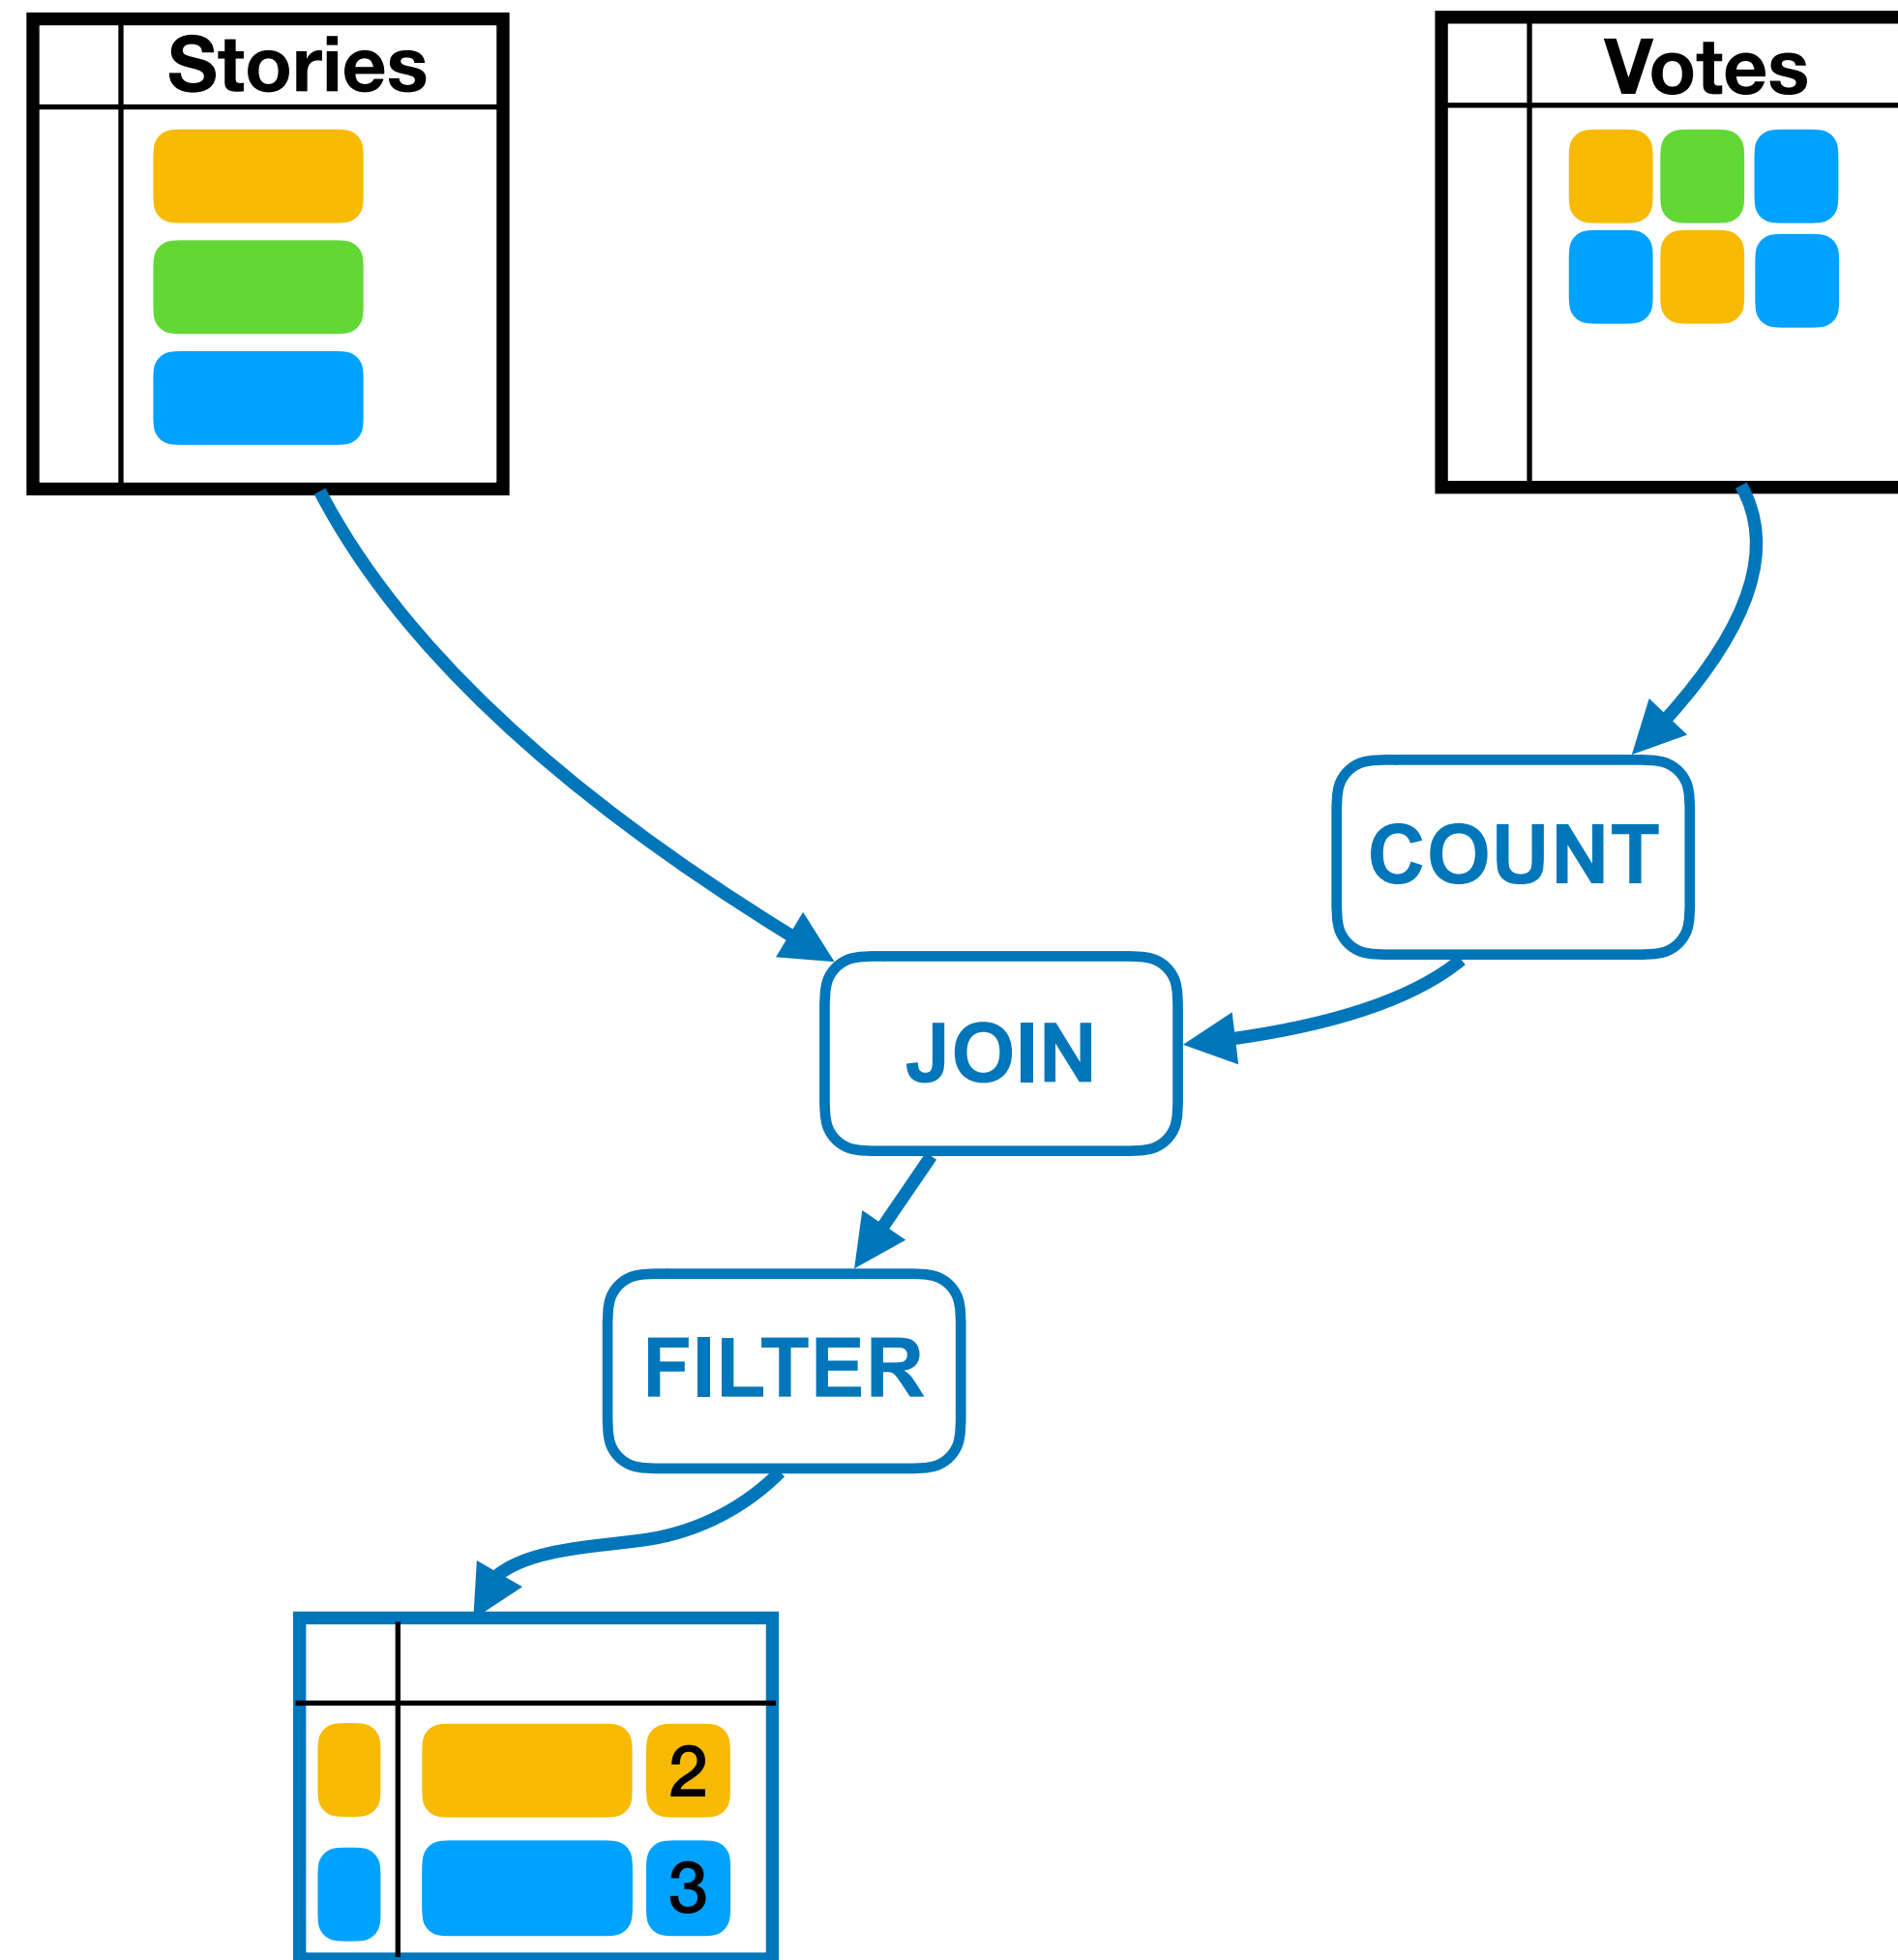
Frontend



# Streaming data-flow?

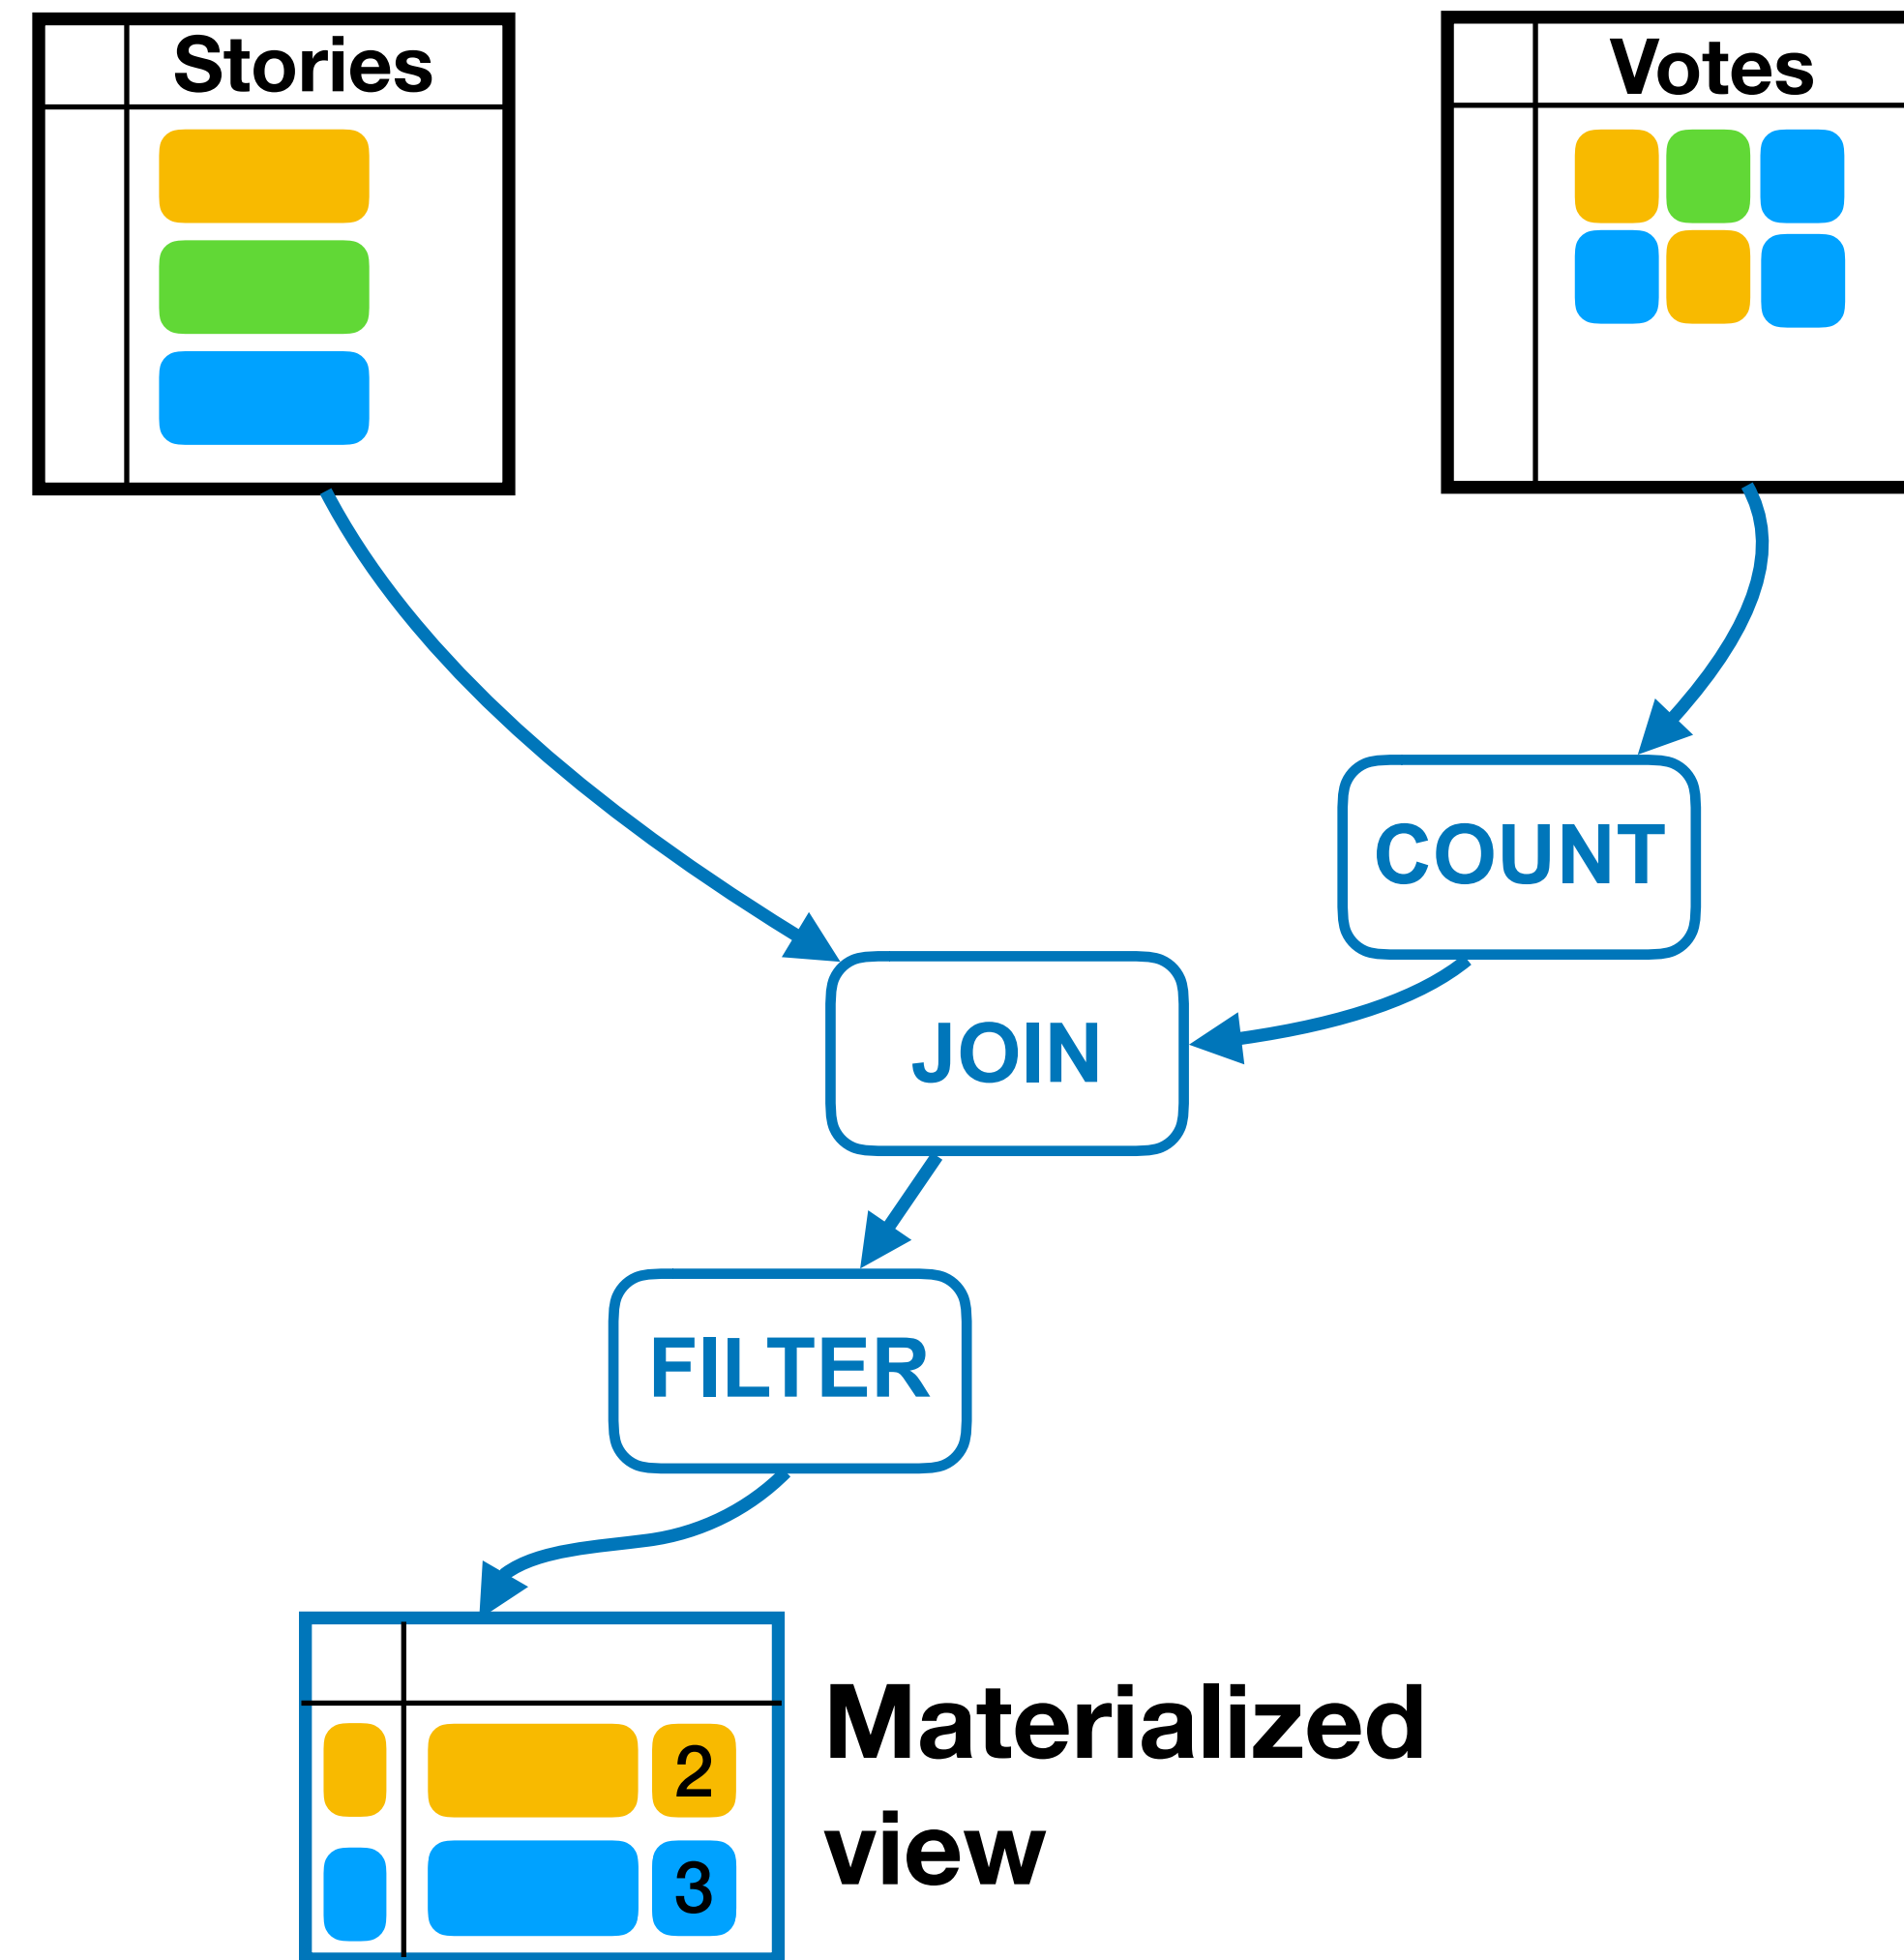


# Streaming data-flow?



Frontend

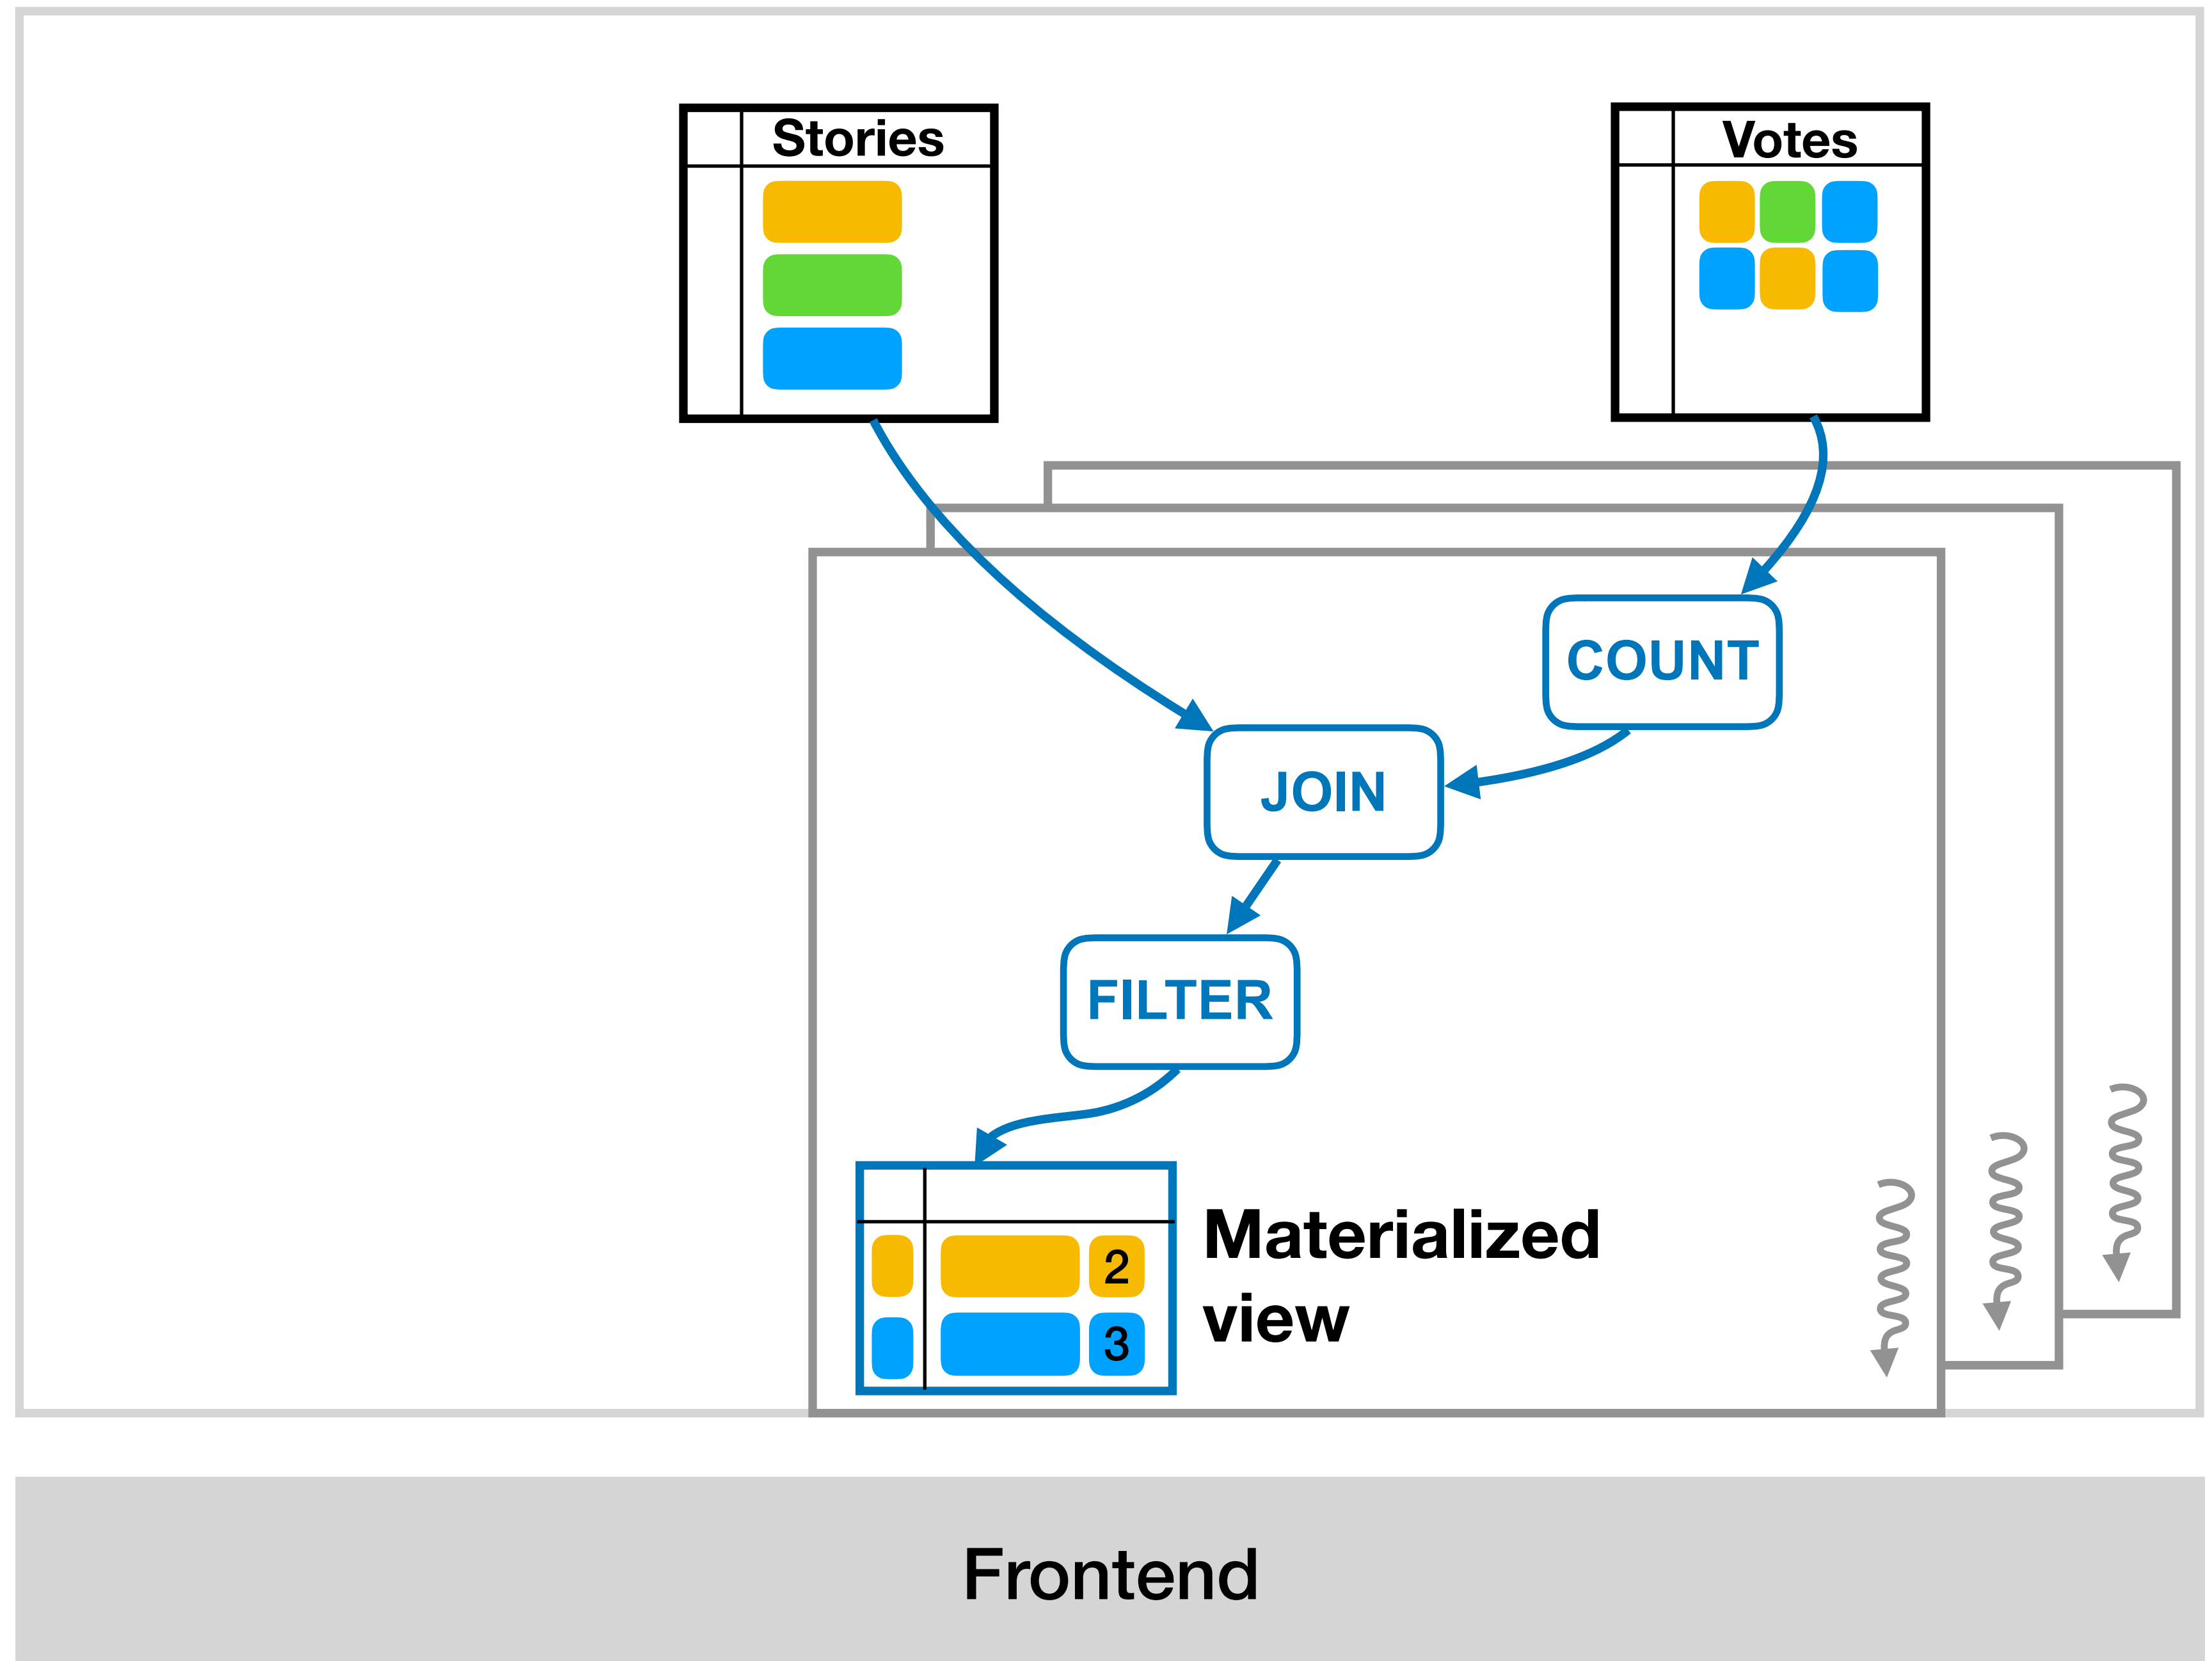
# Streaming data-flow?



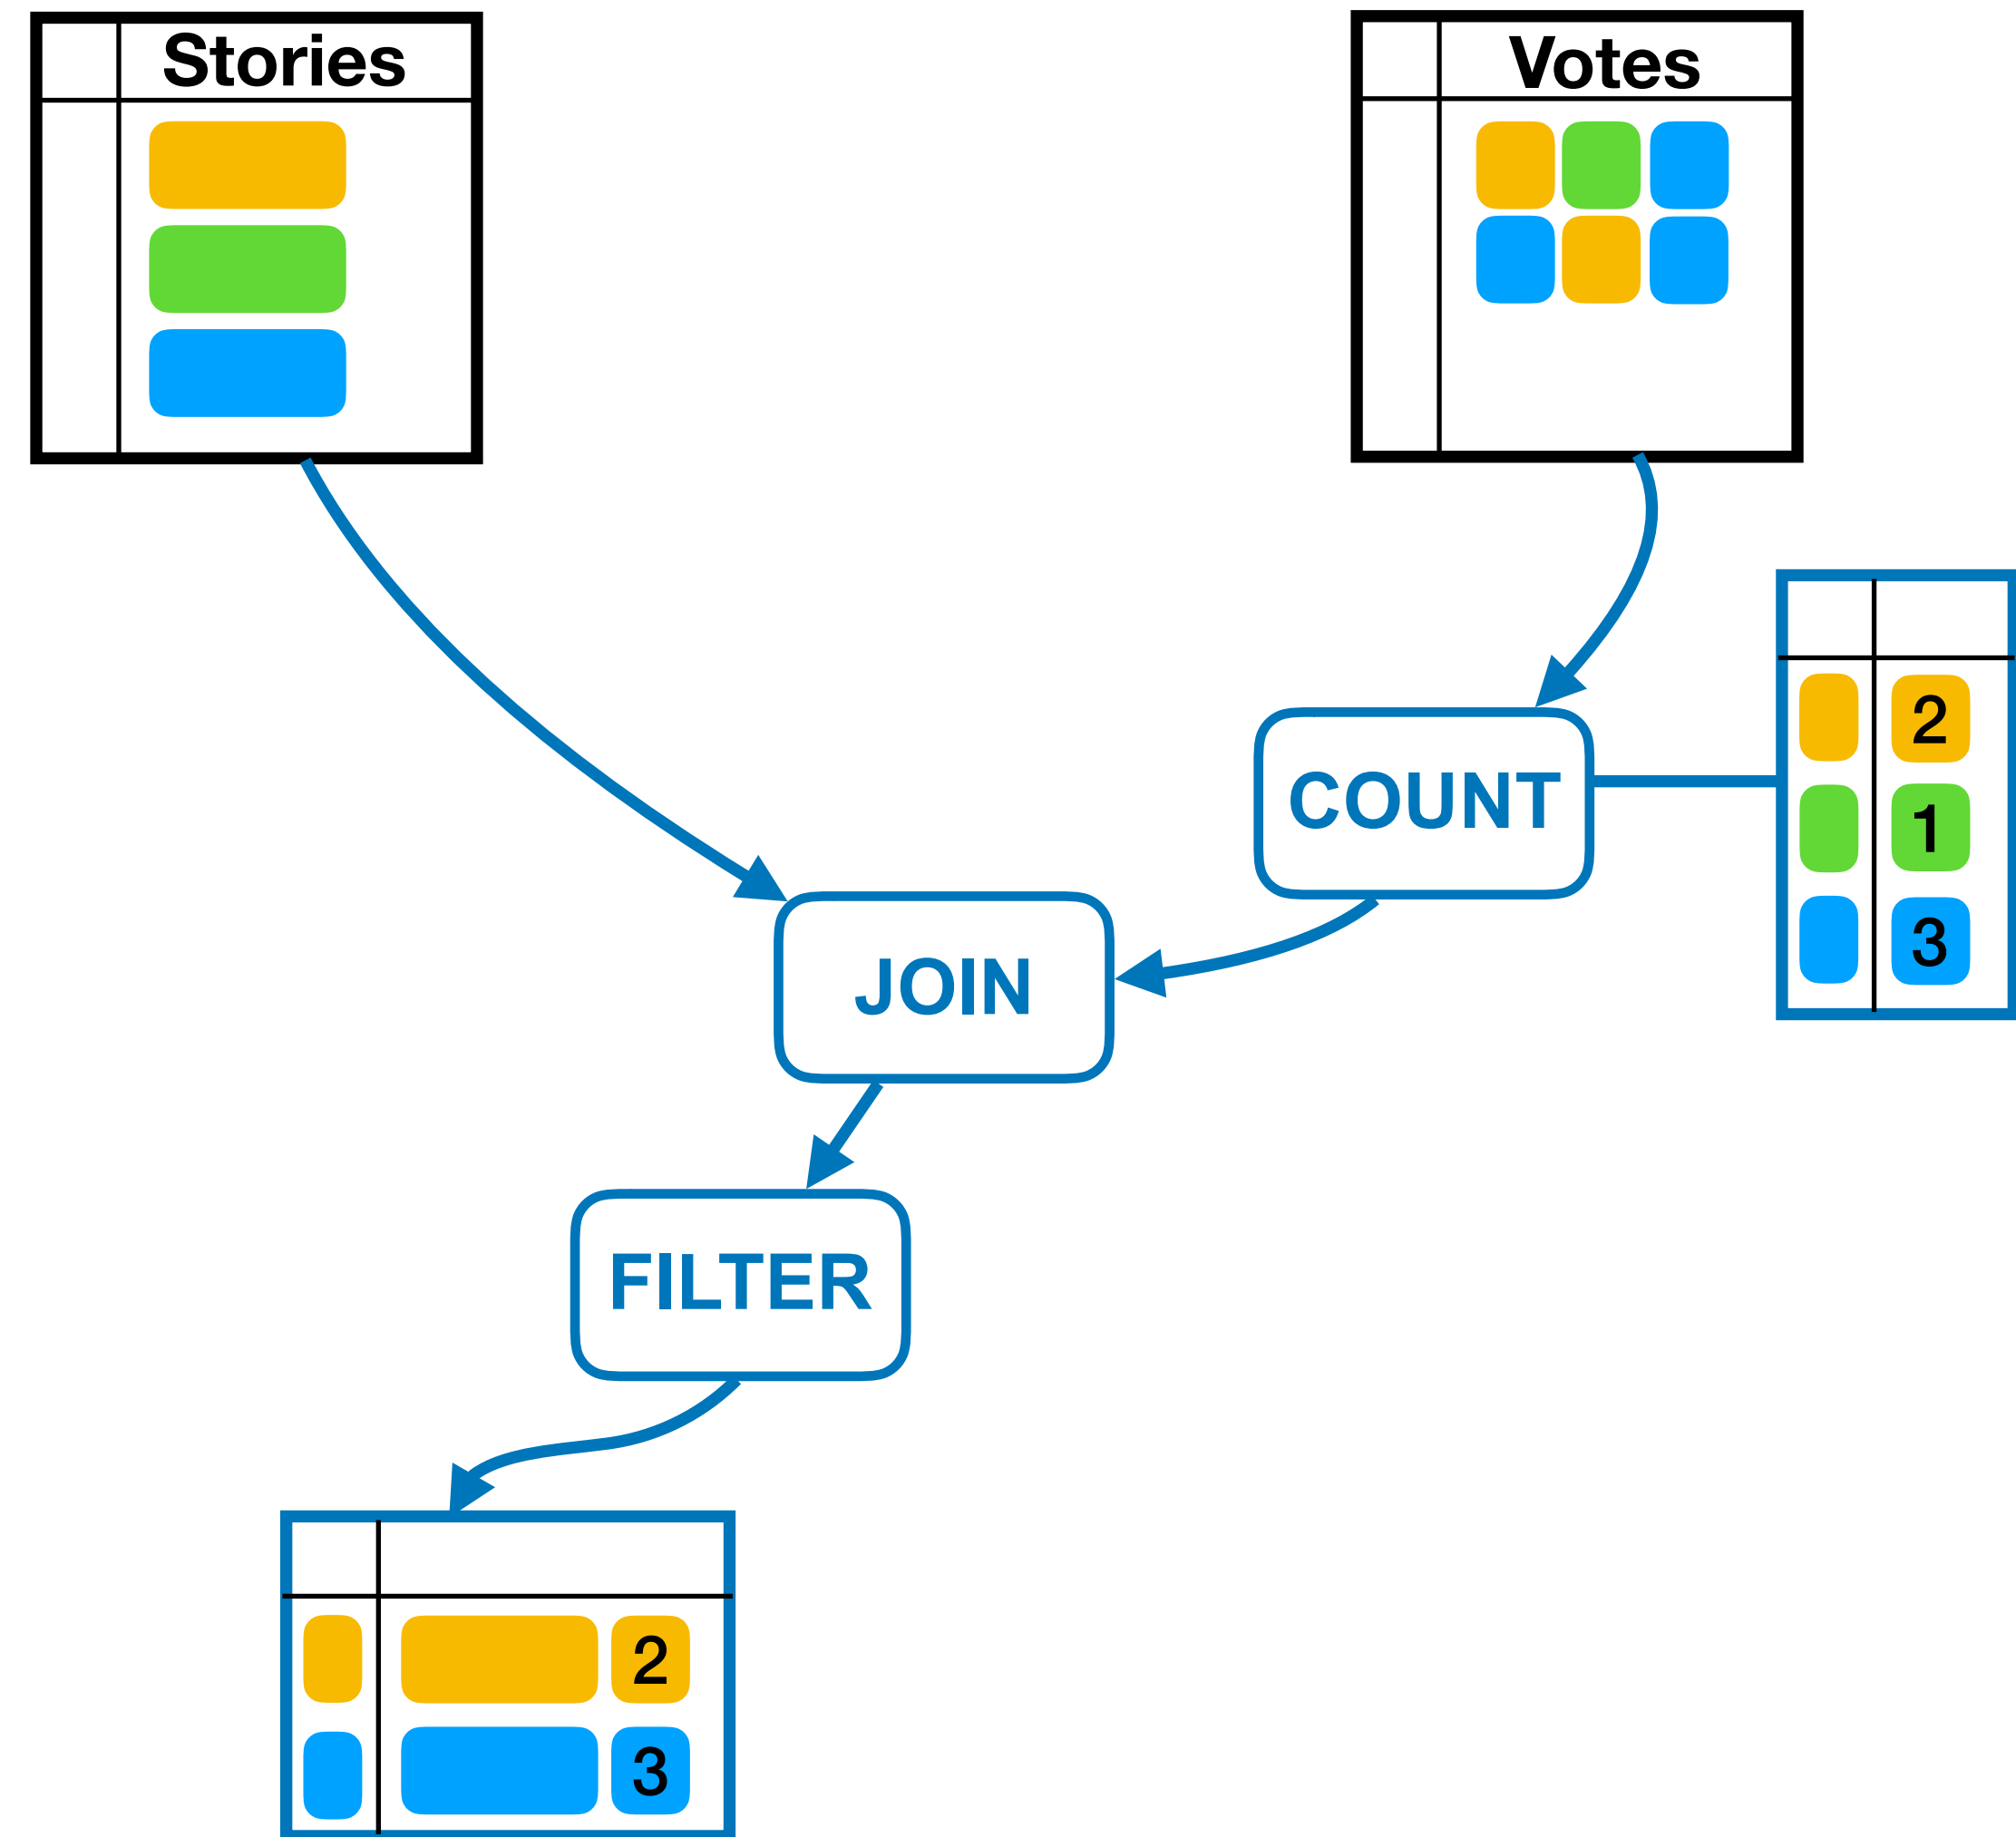
Frontend

# Streaming data-flow?

Fast reads.  
Efficient writes.  
Parallelizes well.



# Challenges




Frontend

# Challenges

State-of-the-art  
data-flow systems:

- **Change queries? Restart!**

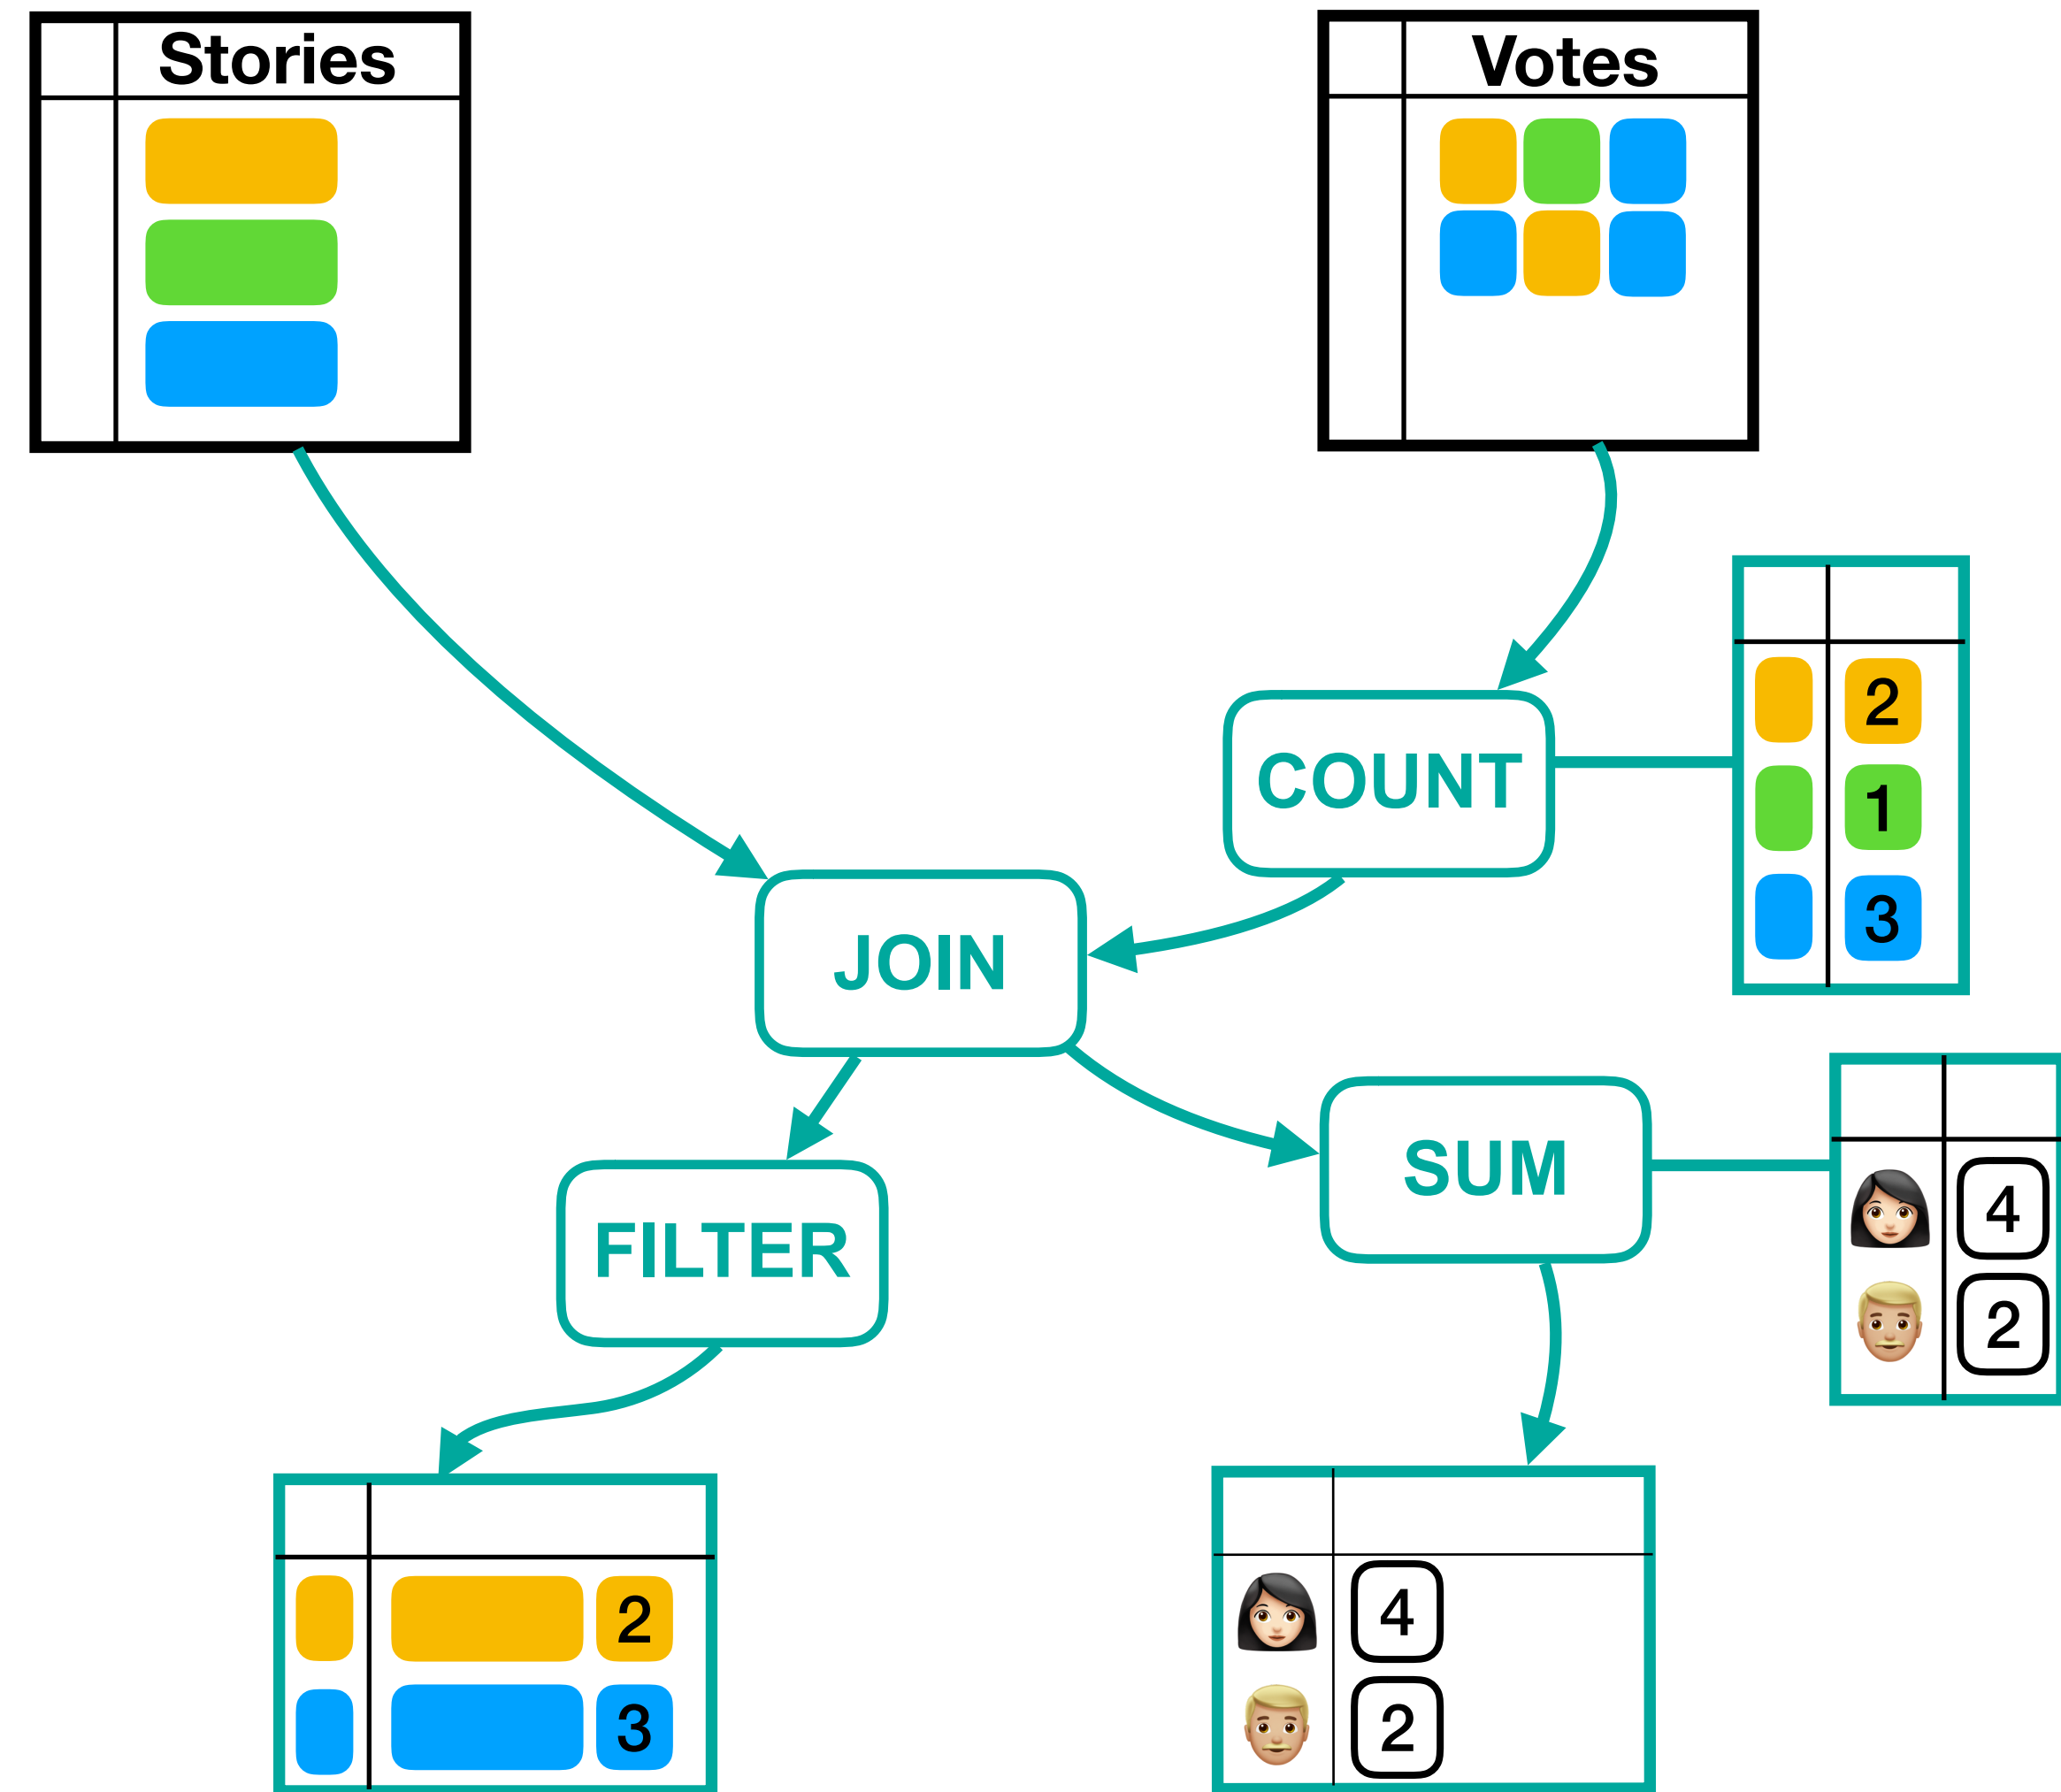


Frontend

# Challenges

State-of-the-art data-flow systems:

► Change queries? **Restart!**

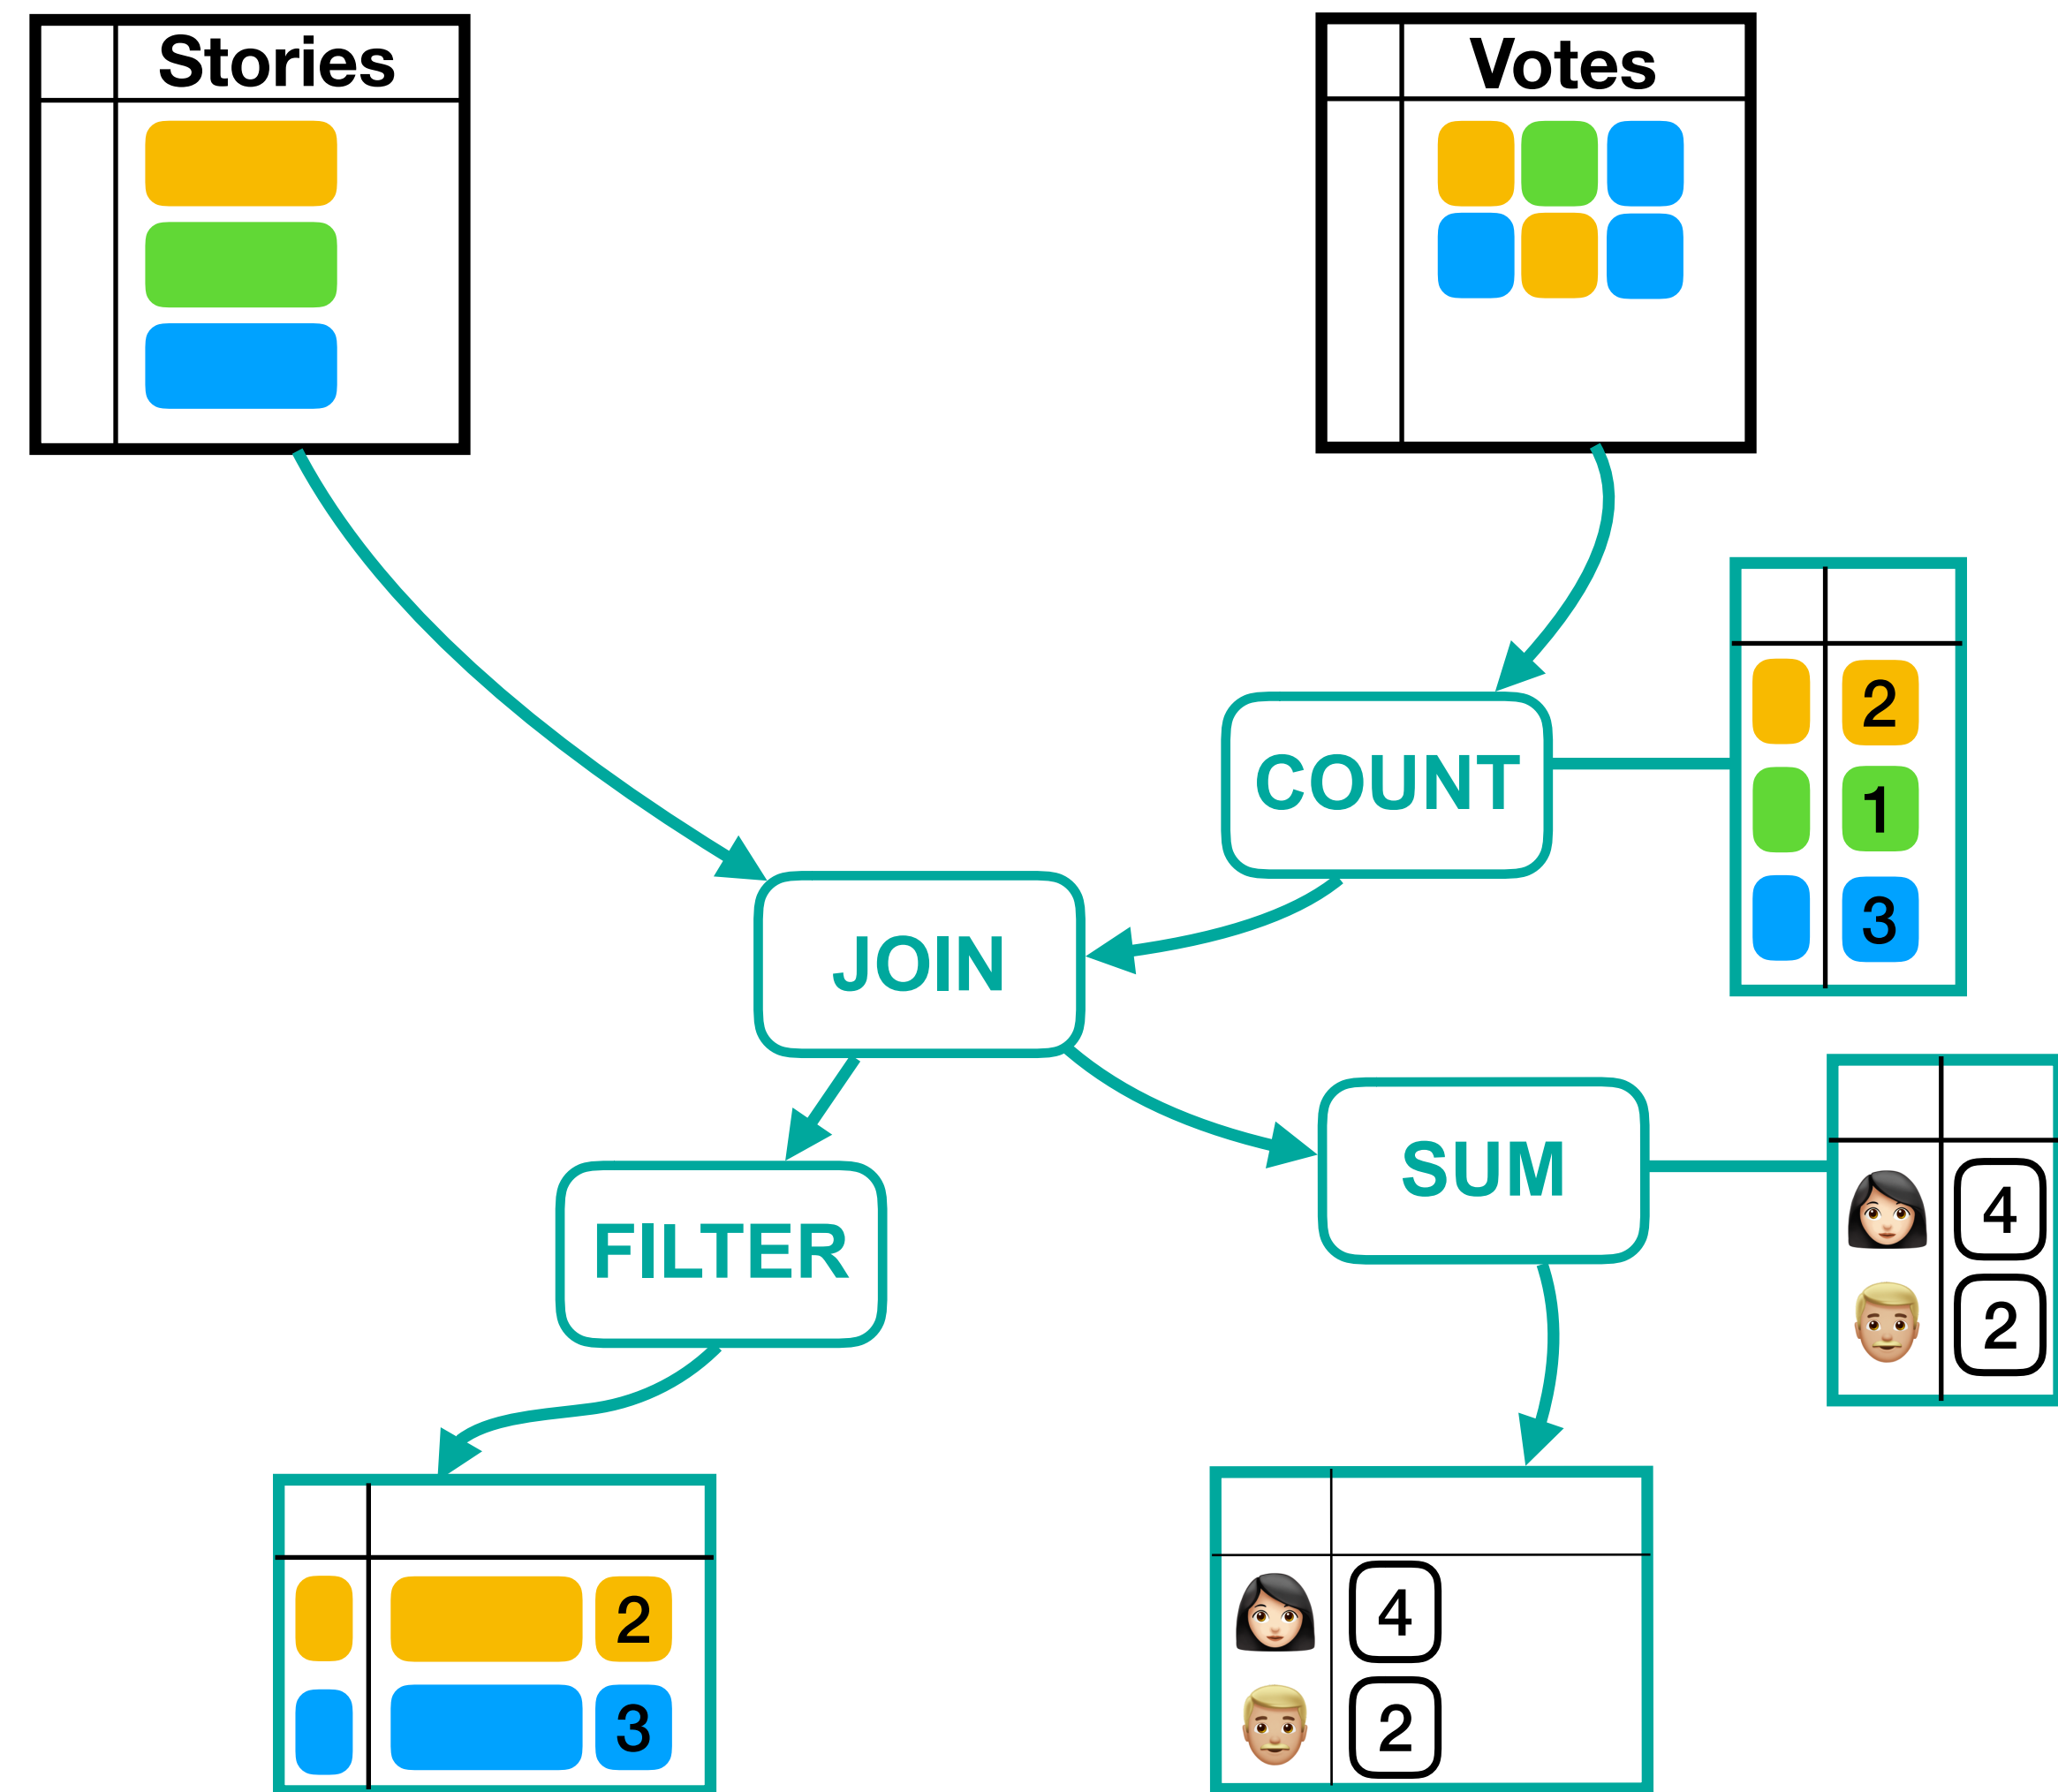


Frontend

# Challenges

State-of-the-art  
data-flow systems:




- ▶ Change queries? **Restart!**
- ▶ Memory footprint? **Grows!**

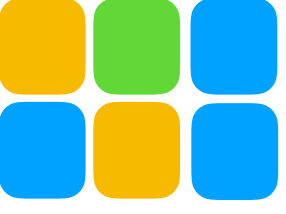


Frontend



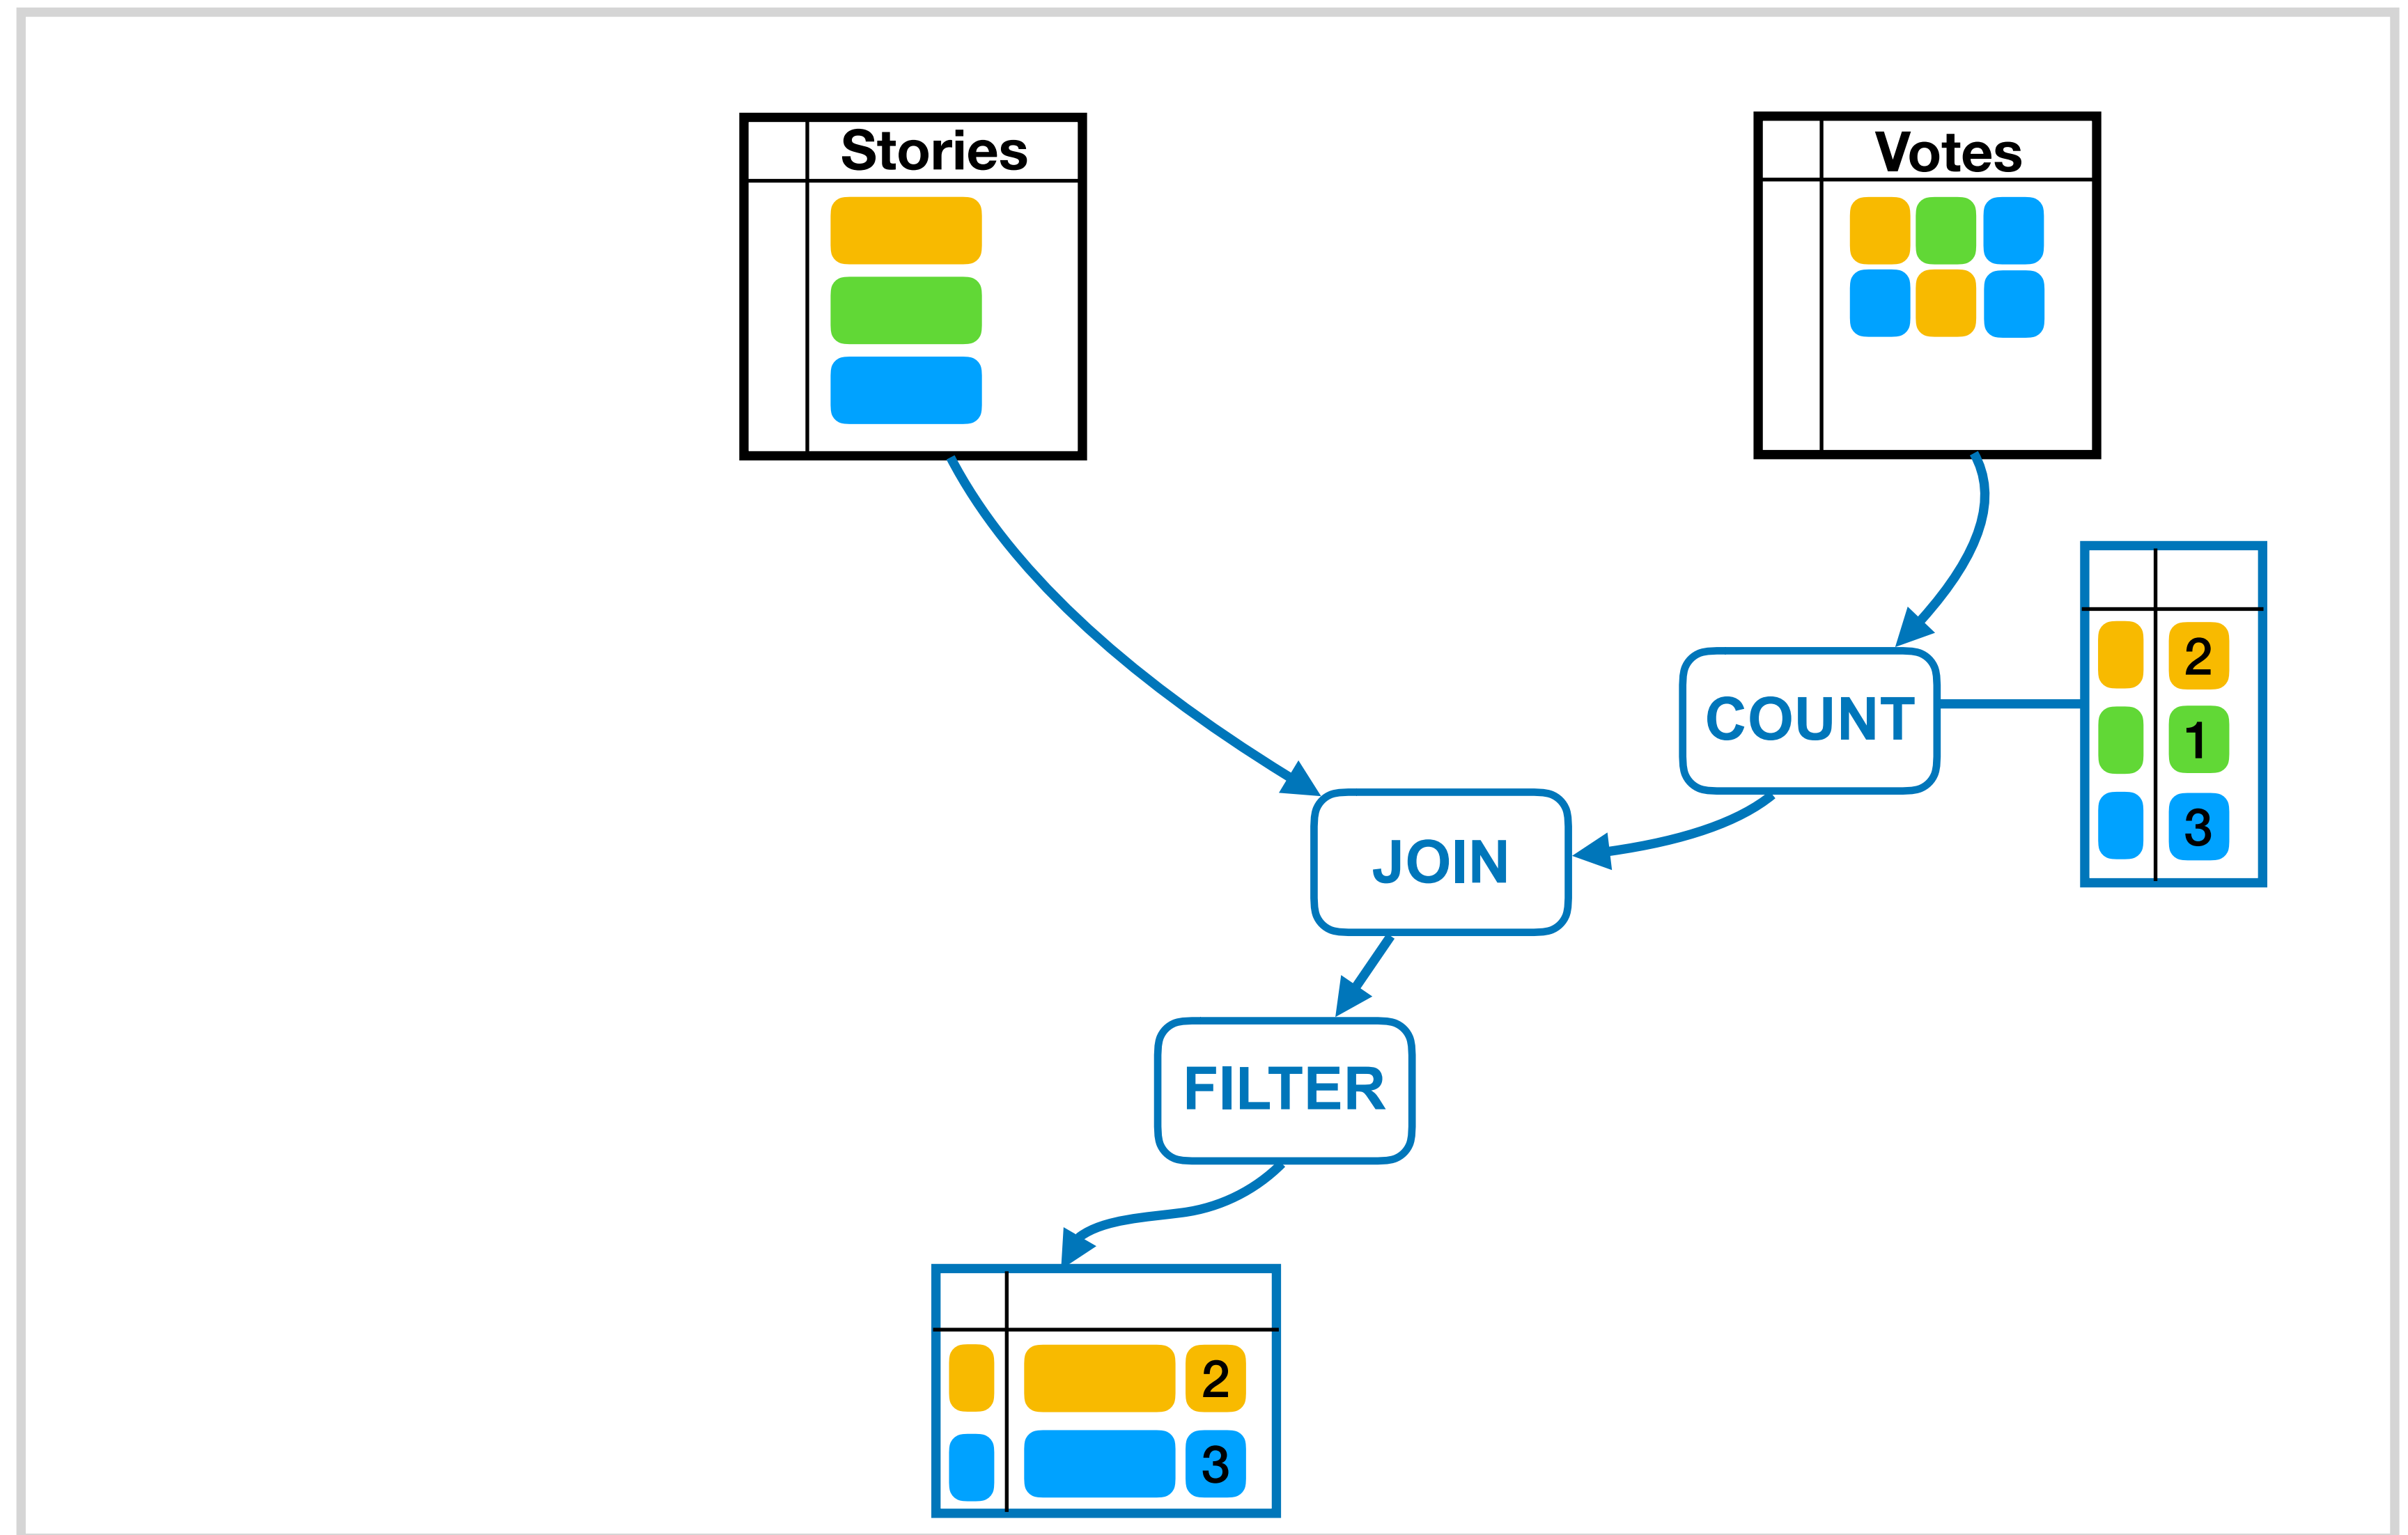
# Noria

	Stories
	
	
	

	Votes
	

Frontend




# Noria









Frontend




# Noria

► Change queries? **Live.**

Stories	
	
	
	





Votes		
		
		

COUNT

	2
	1
	3

JOIN

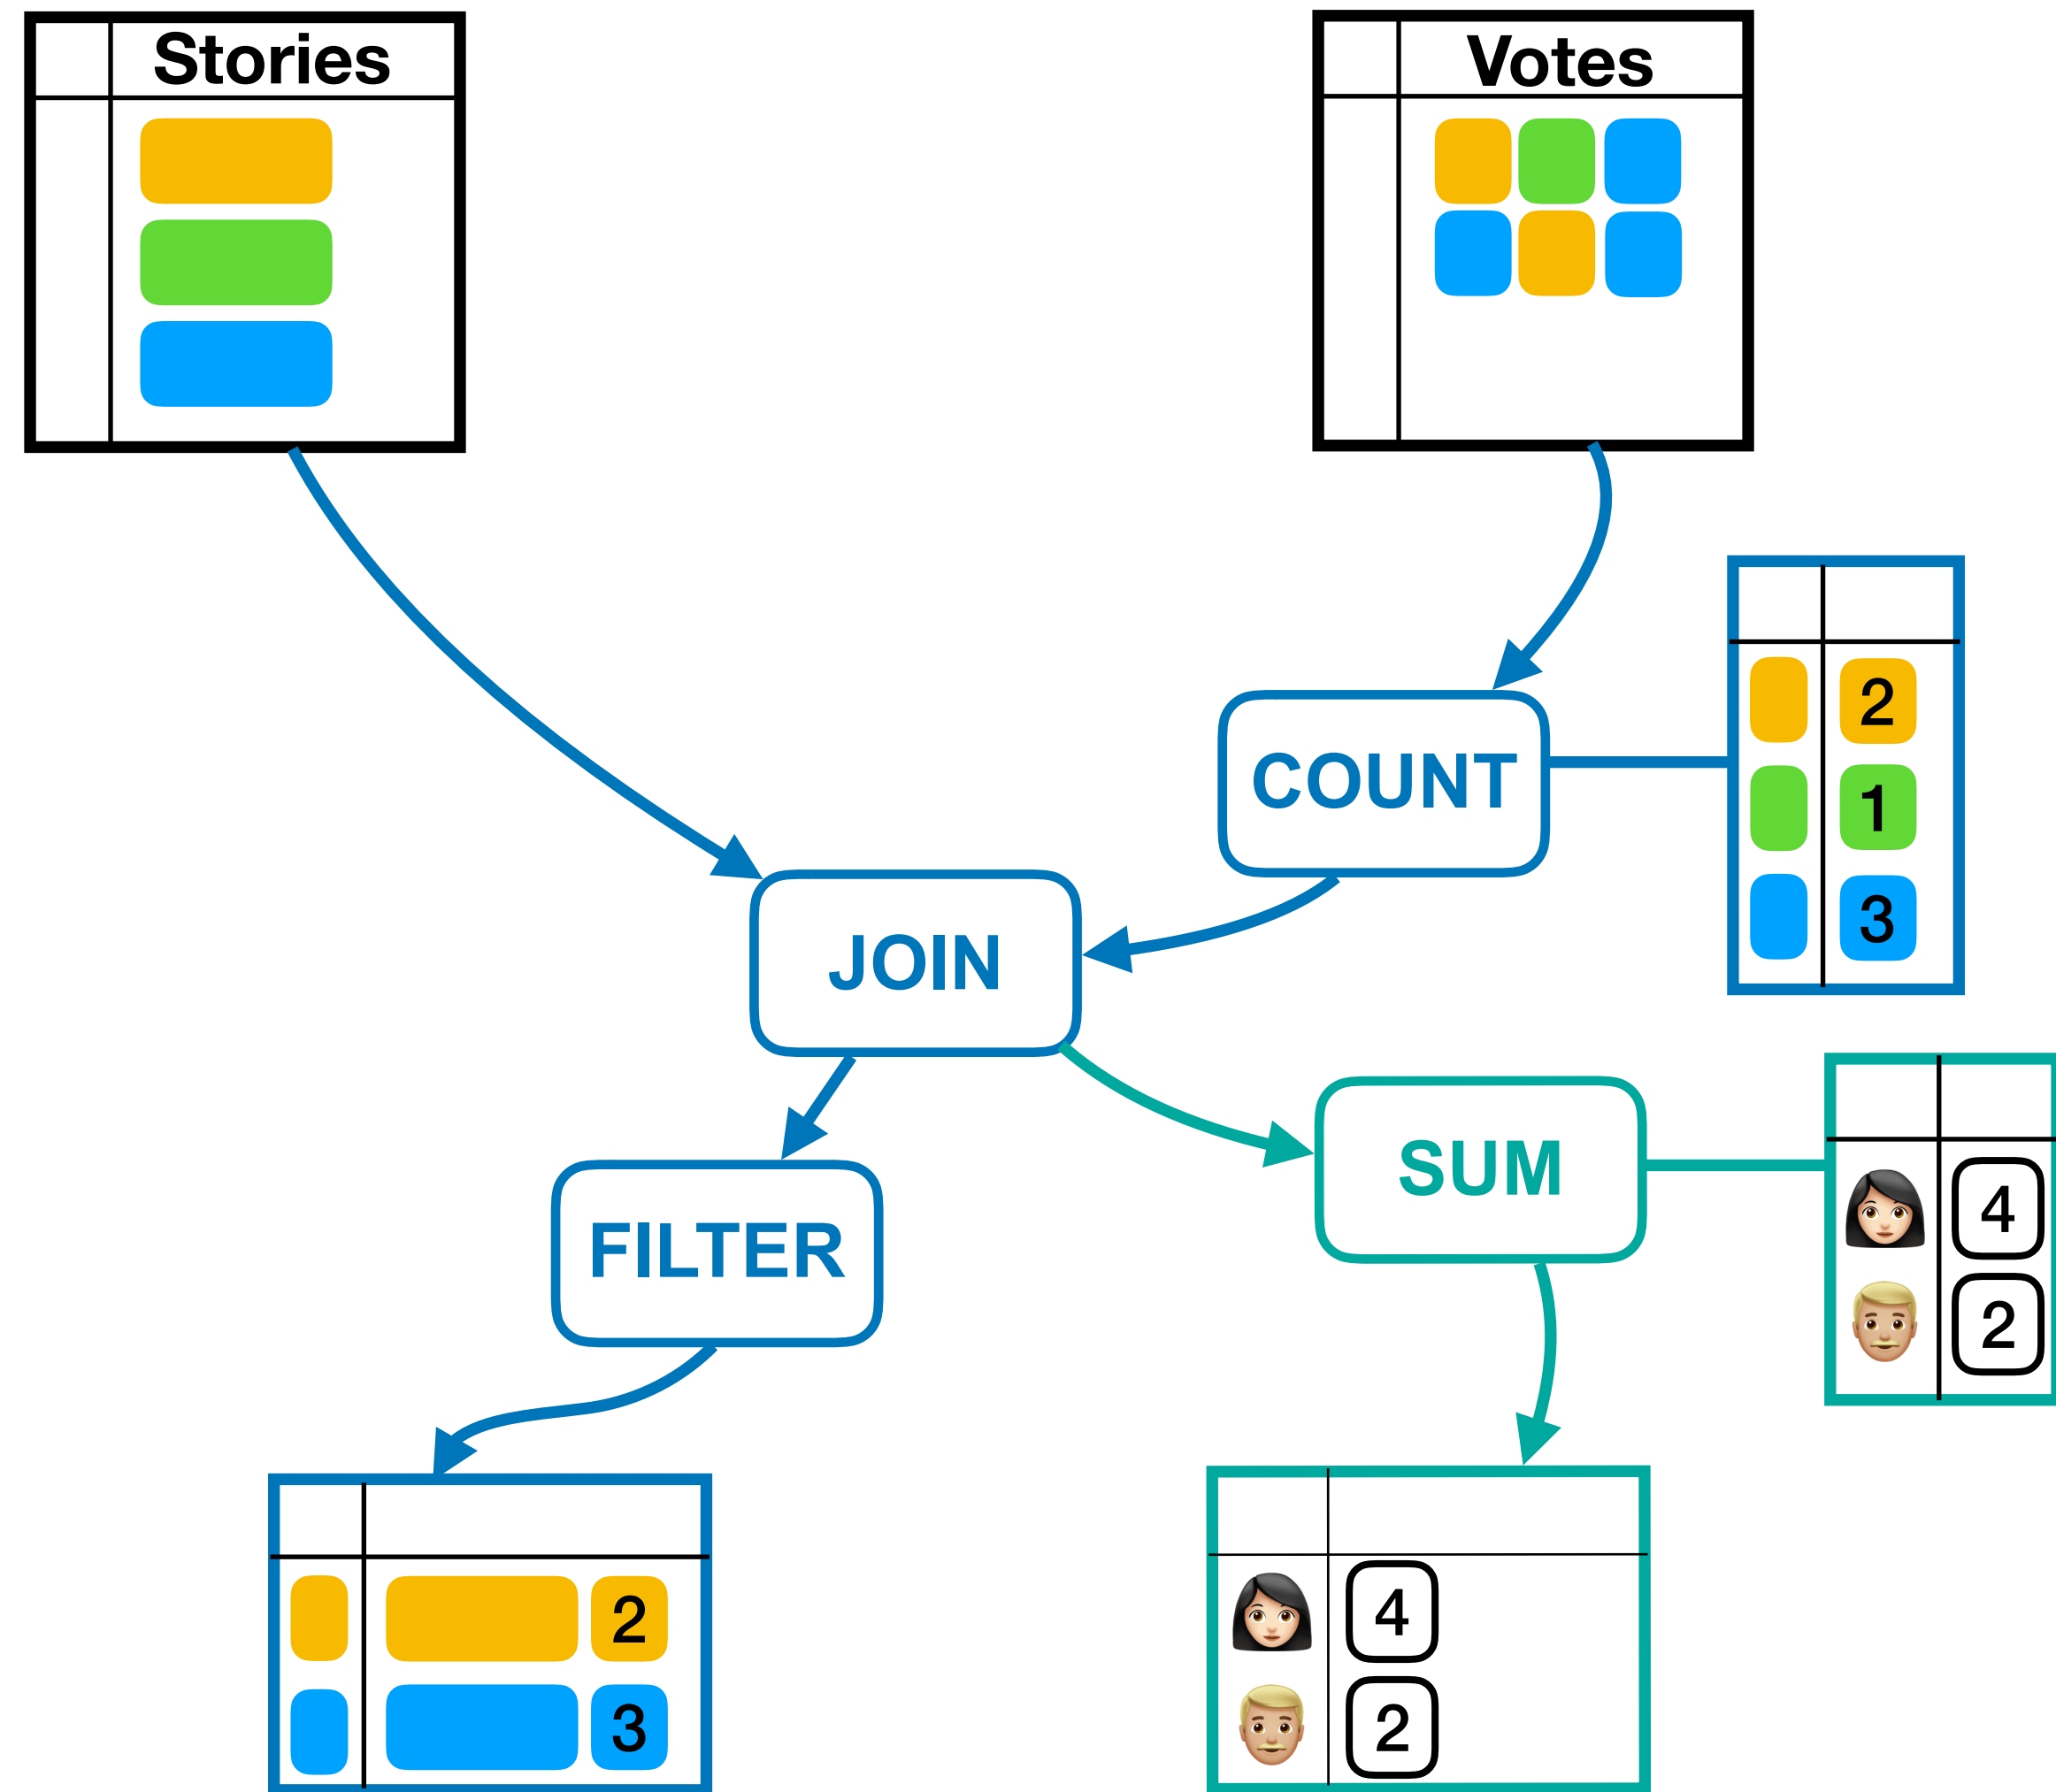
FILTER

		2
		3

Frontend

# Noria

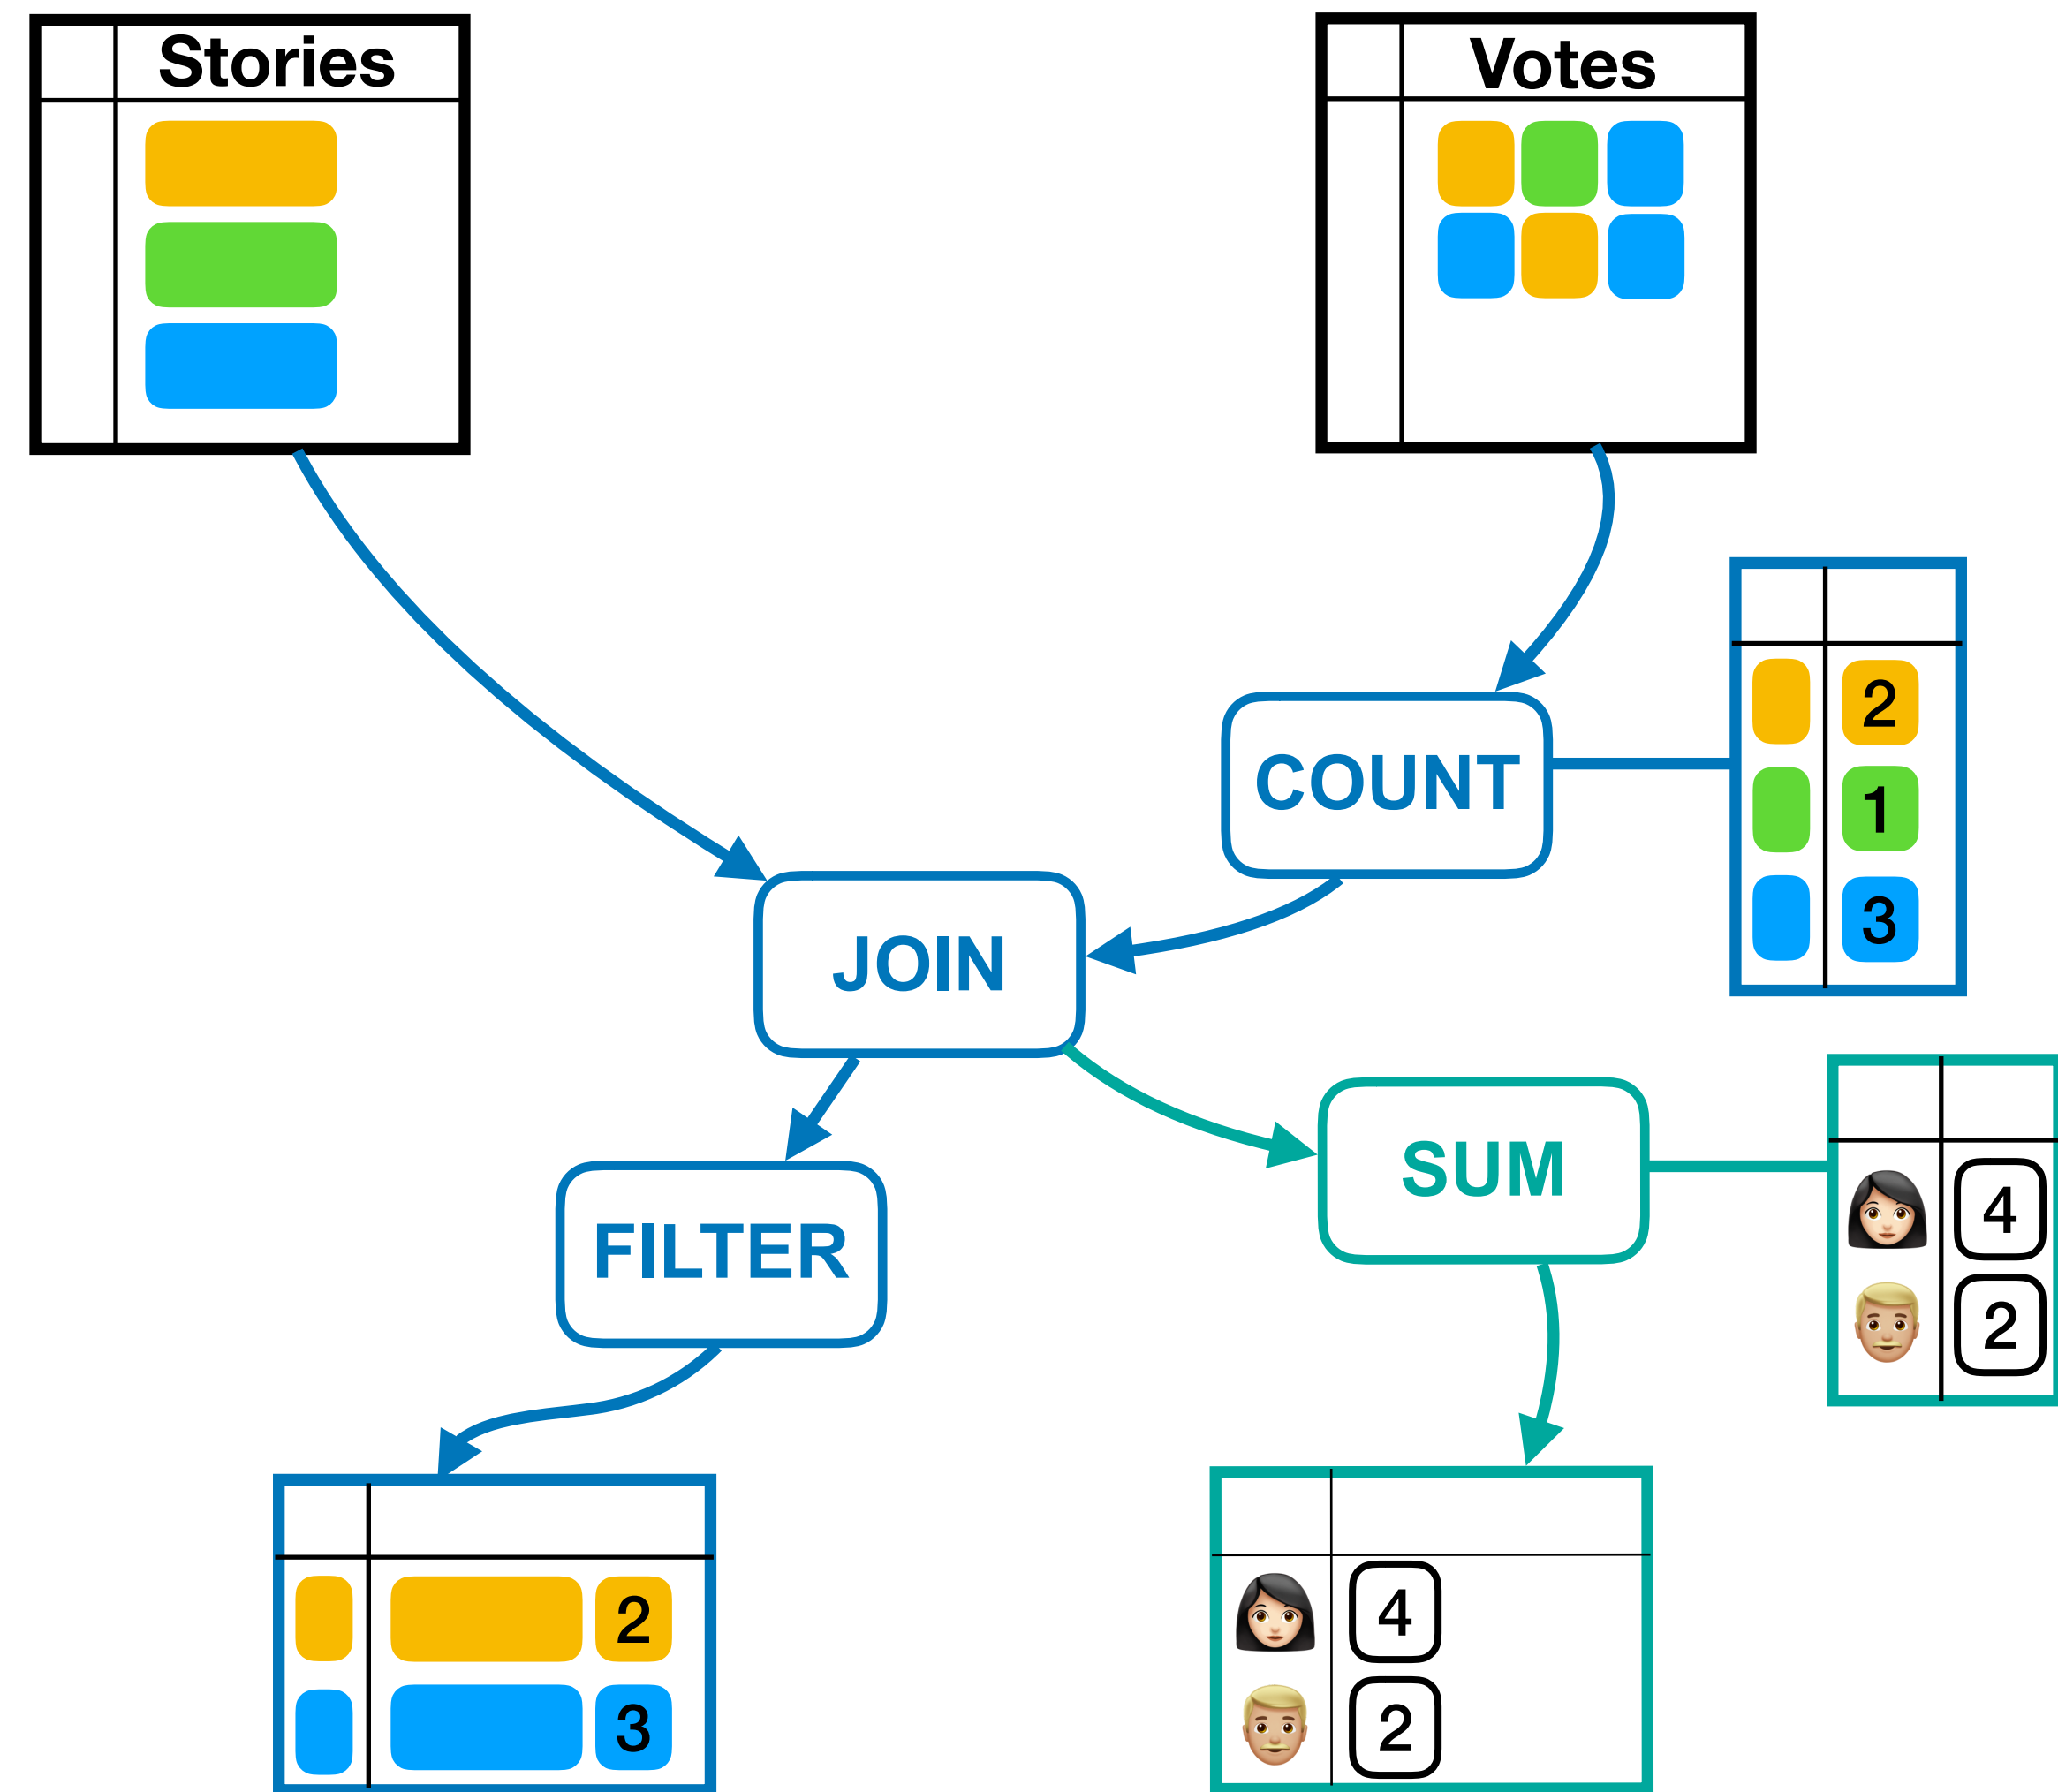
► Change queries? **Live.**



Frontend

# Noria

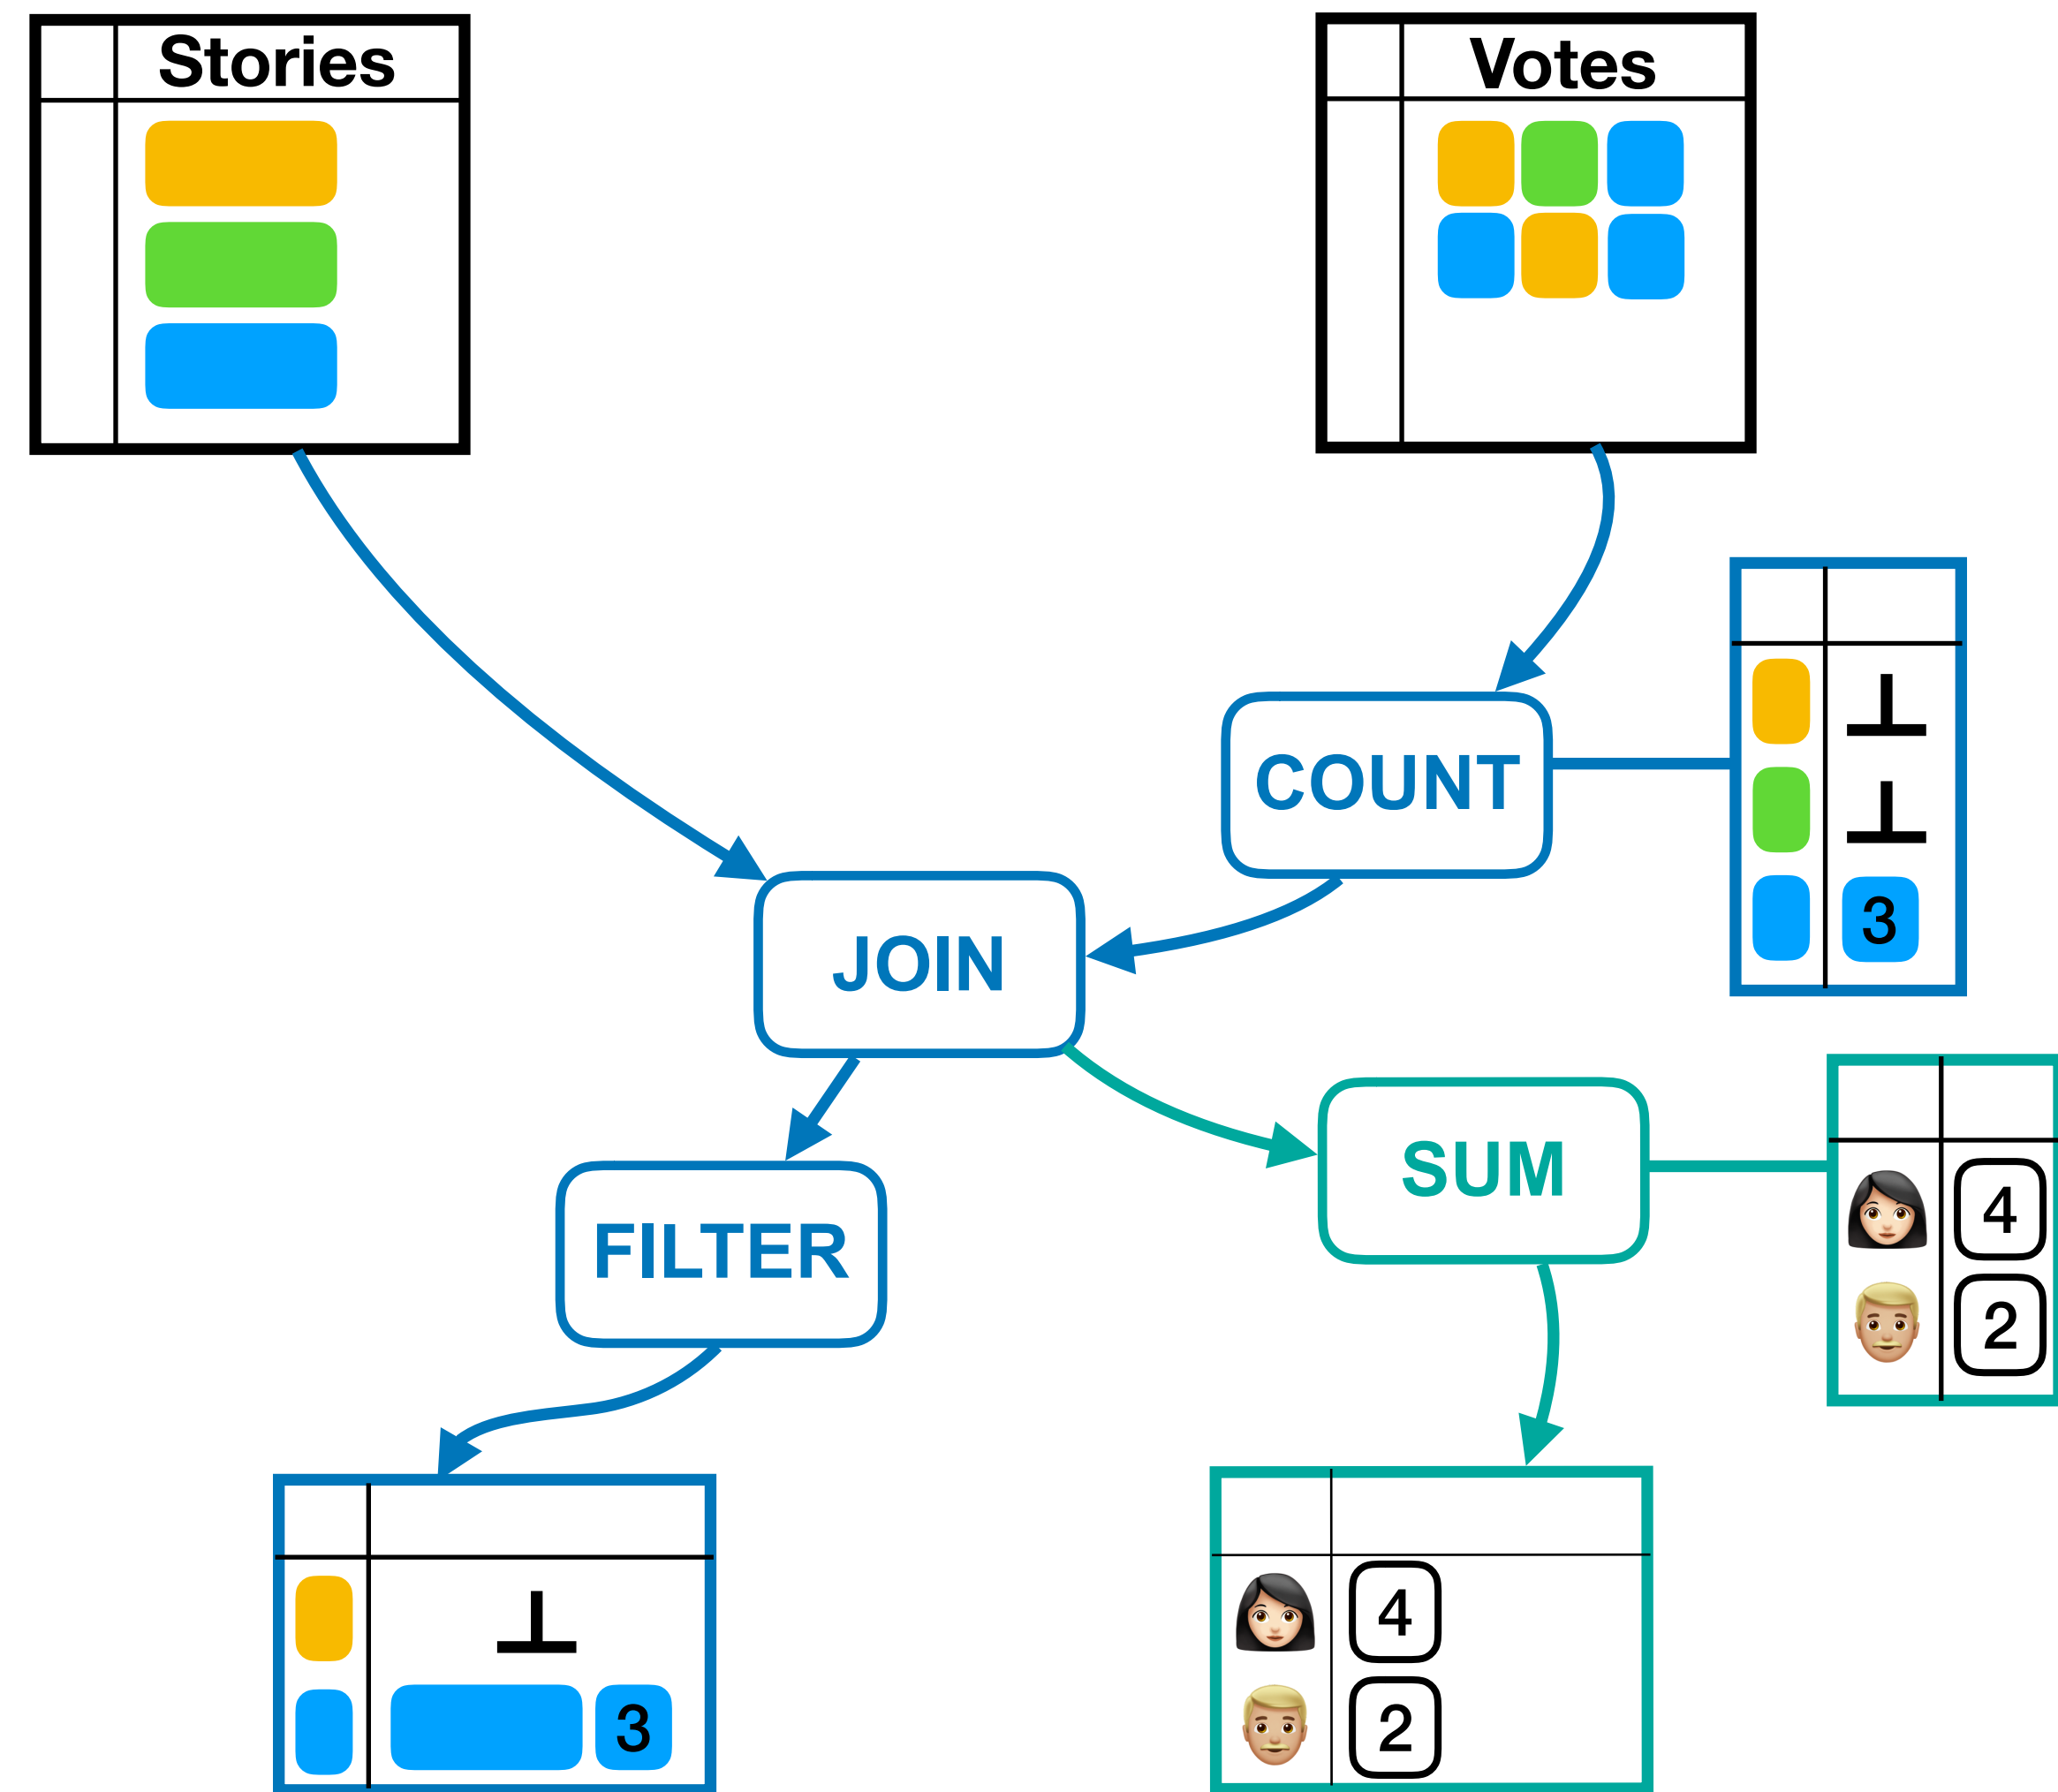
- ▶ Change queries? **Live.**
- ▶ Memory footprint? **Bounded.**



Frontend

# Noria

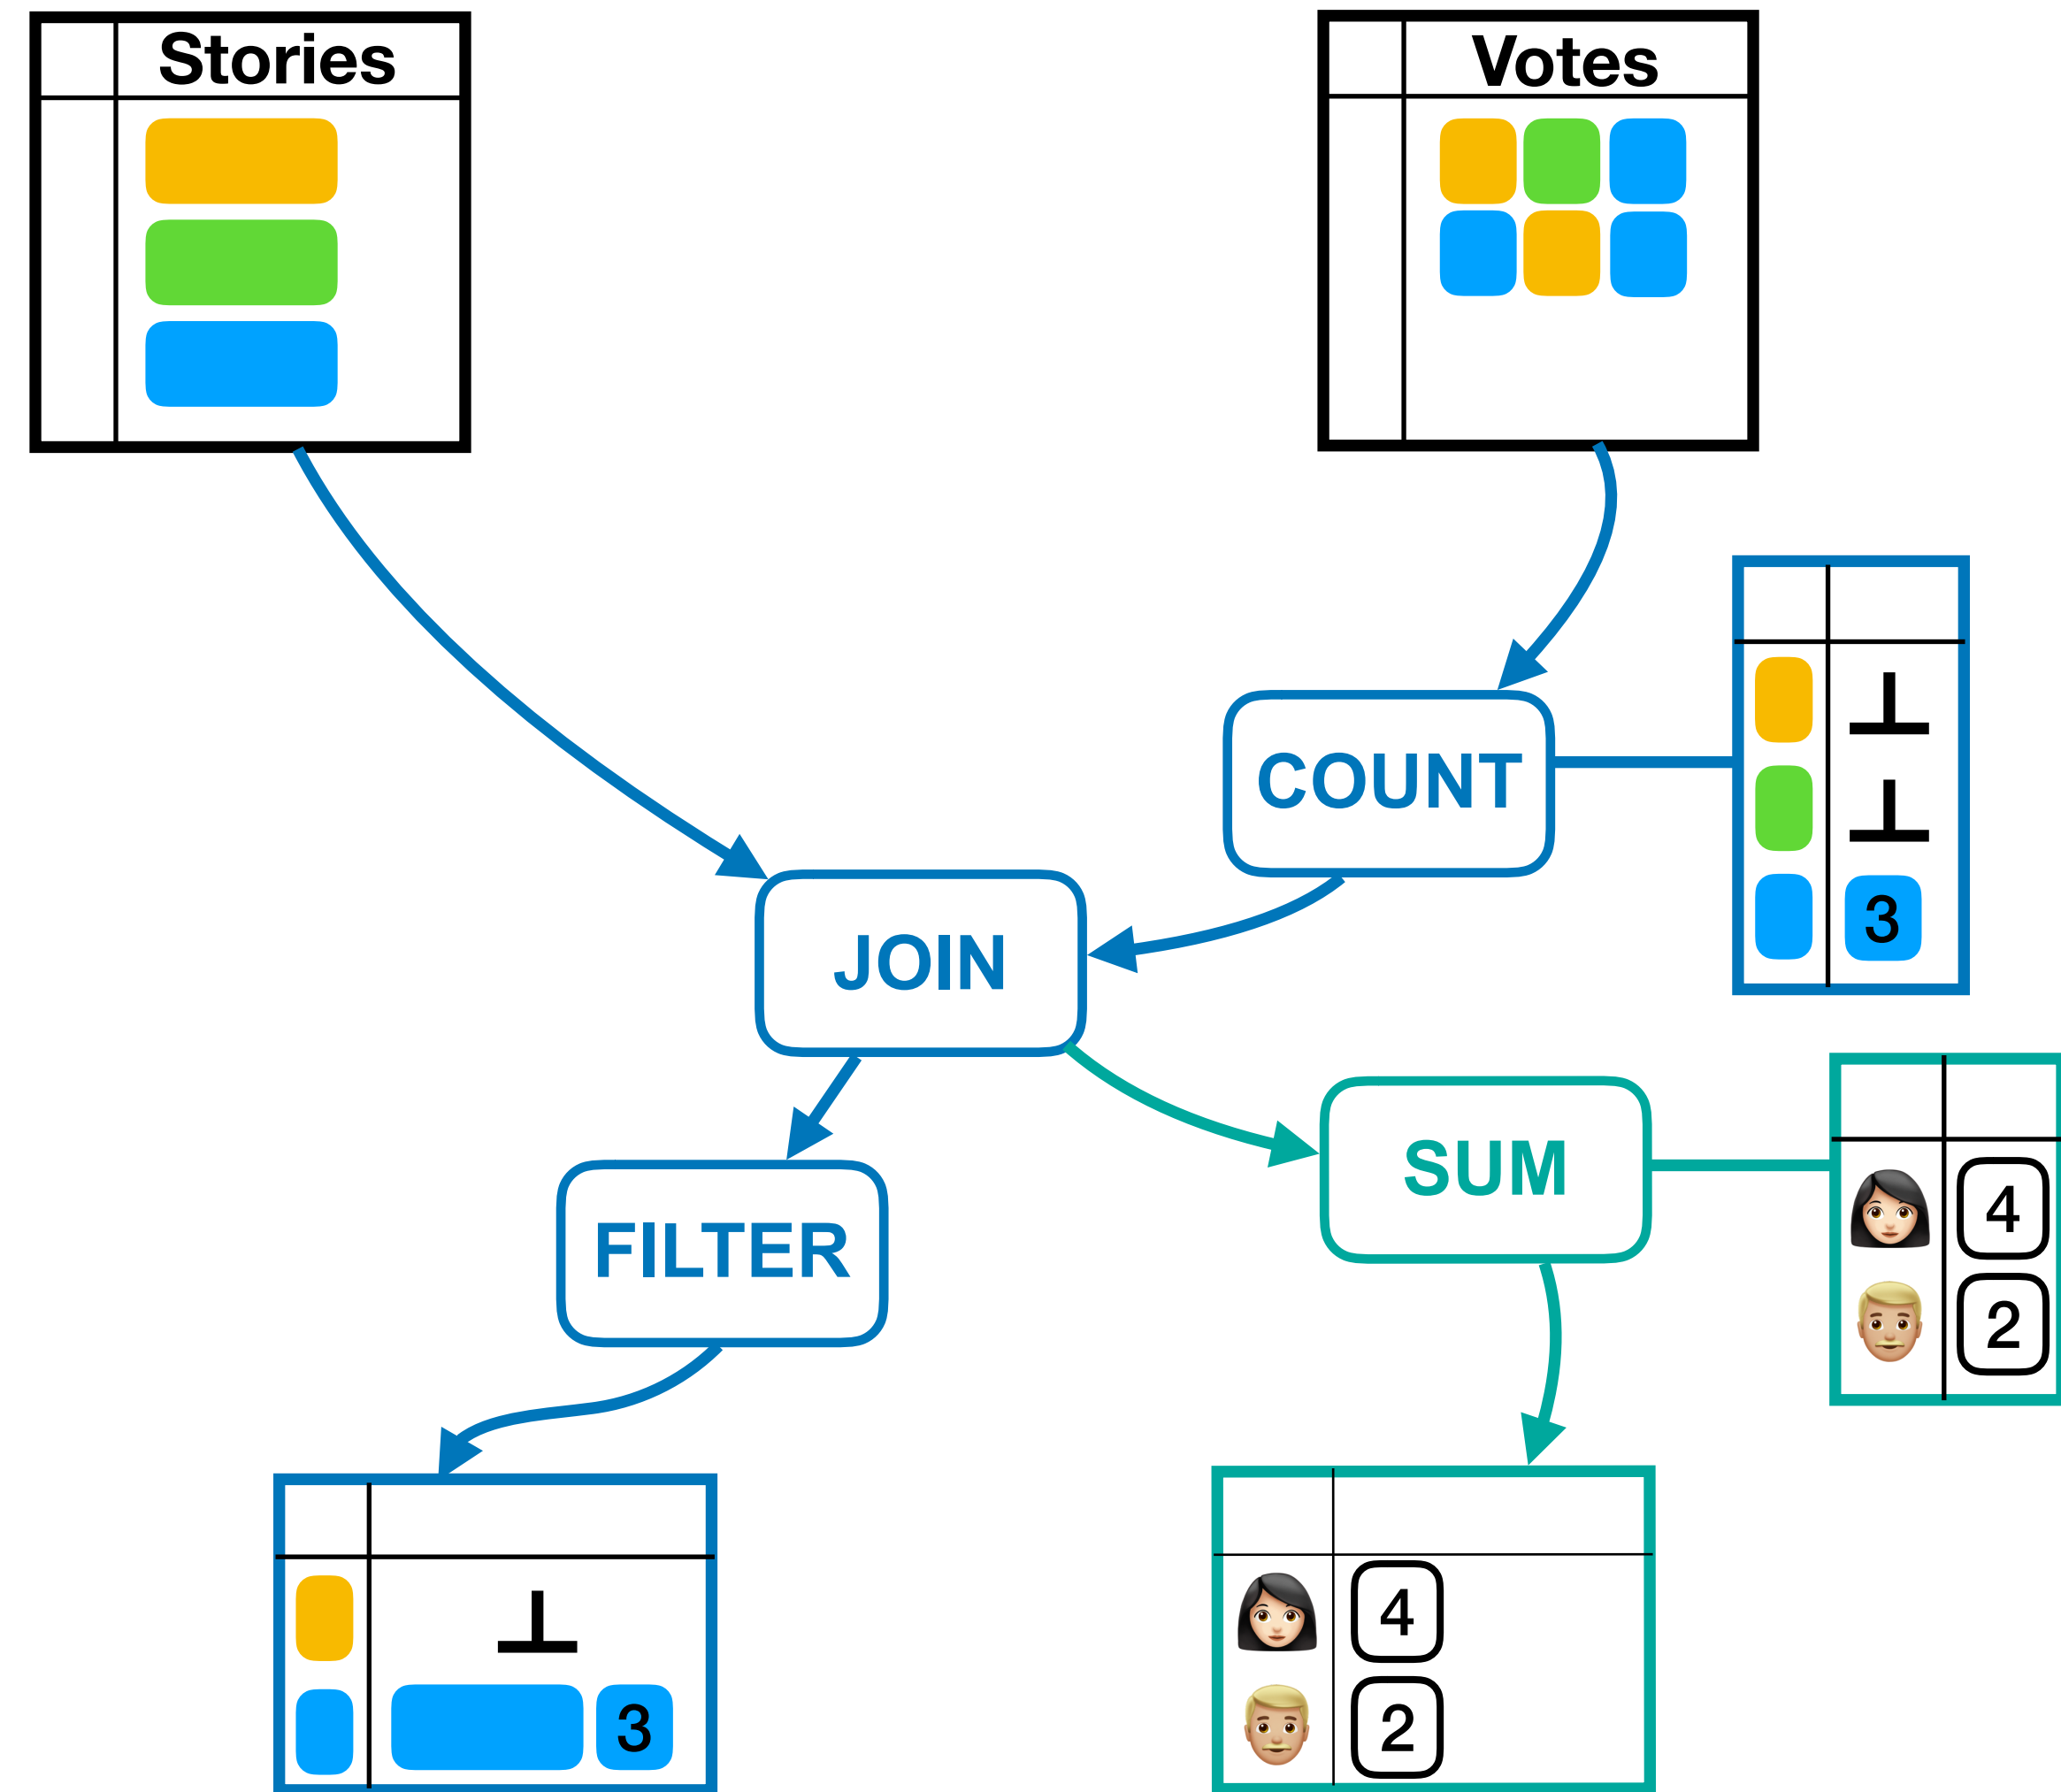
- ▶ Change queries? **Live.**
- ▶ Memory footprint? **Bounded.**



Frontend

# Noria

- ▶ Change queries? **Live.**
- ▶ Memory footprint? **Bounded.**
- ▶ No global coordination.



Frontend

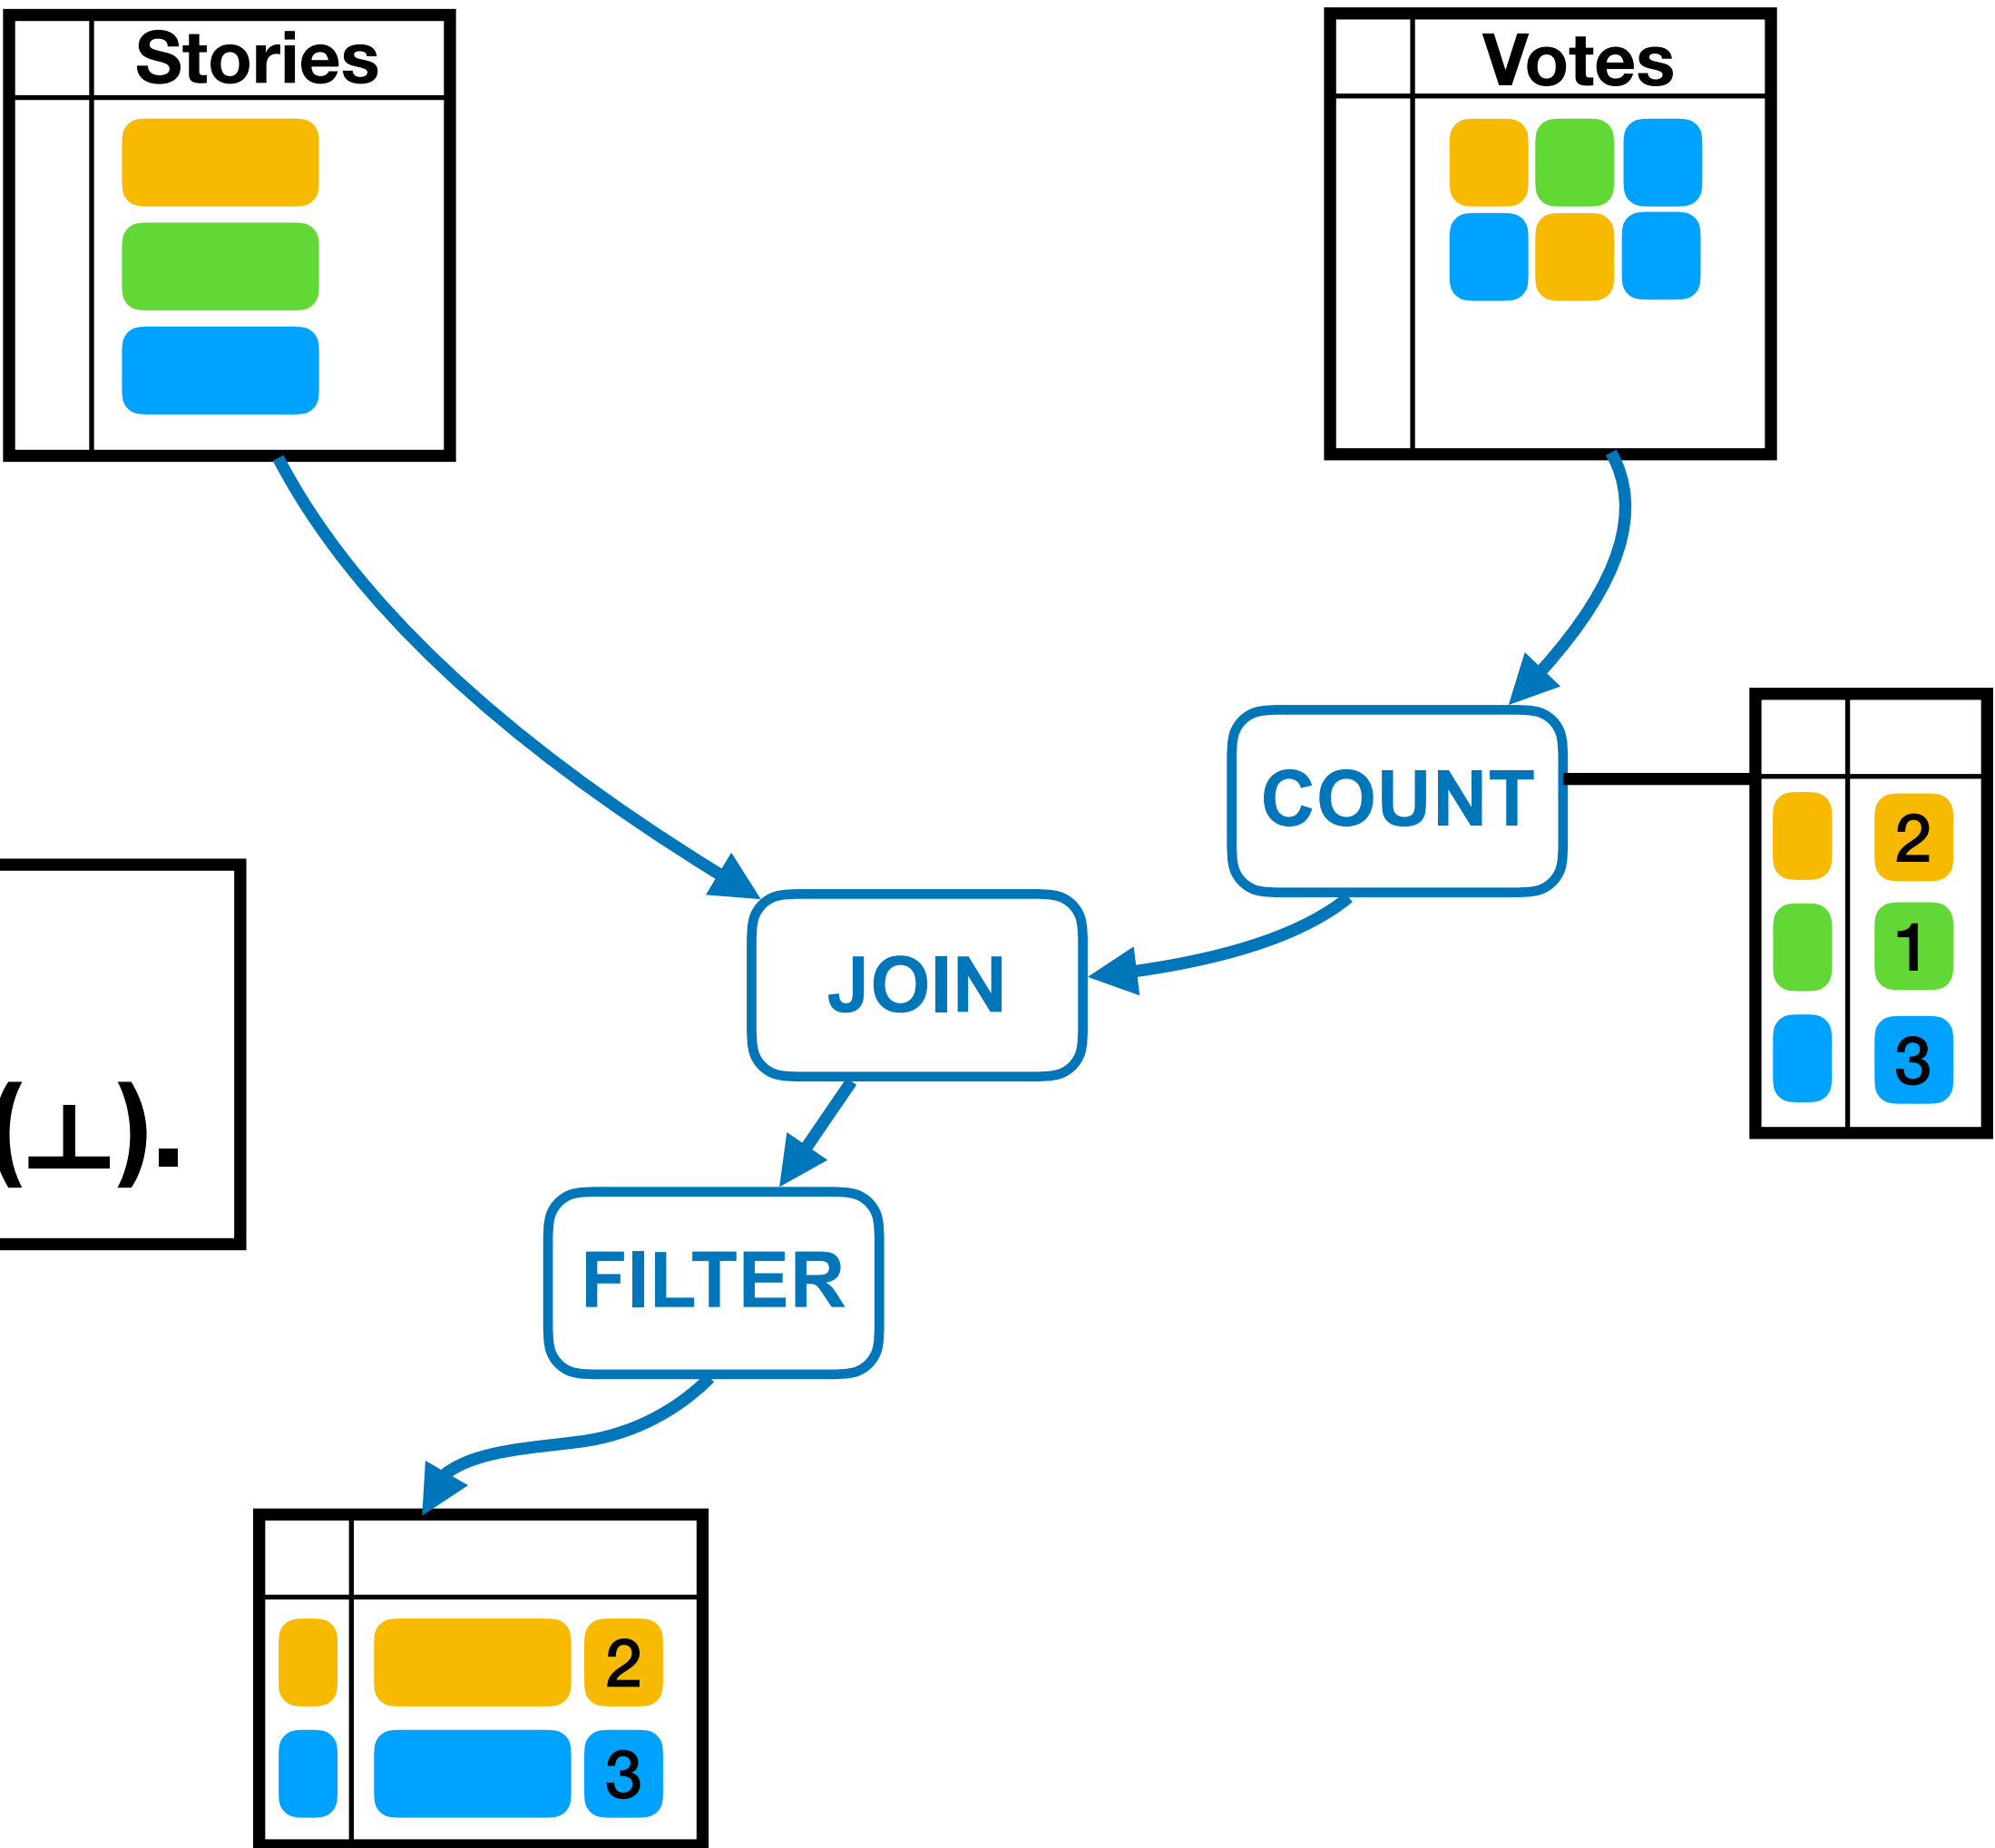
New model:

# Partially-stateful data-flow



# Partially-stateful data-flow

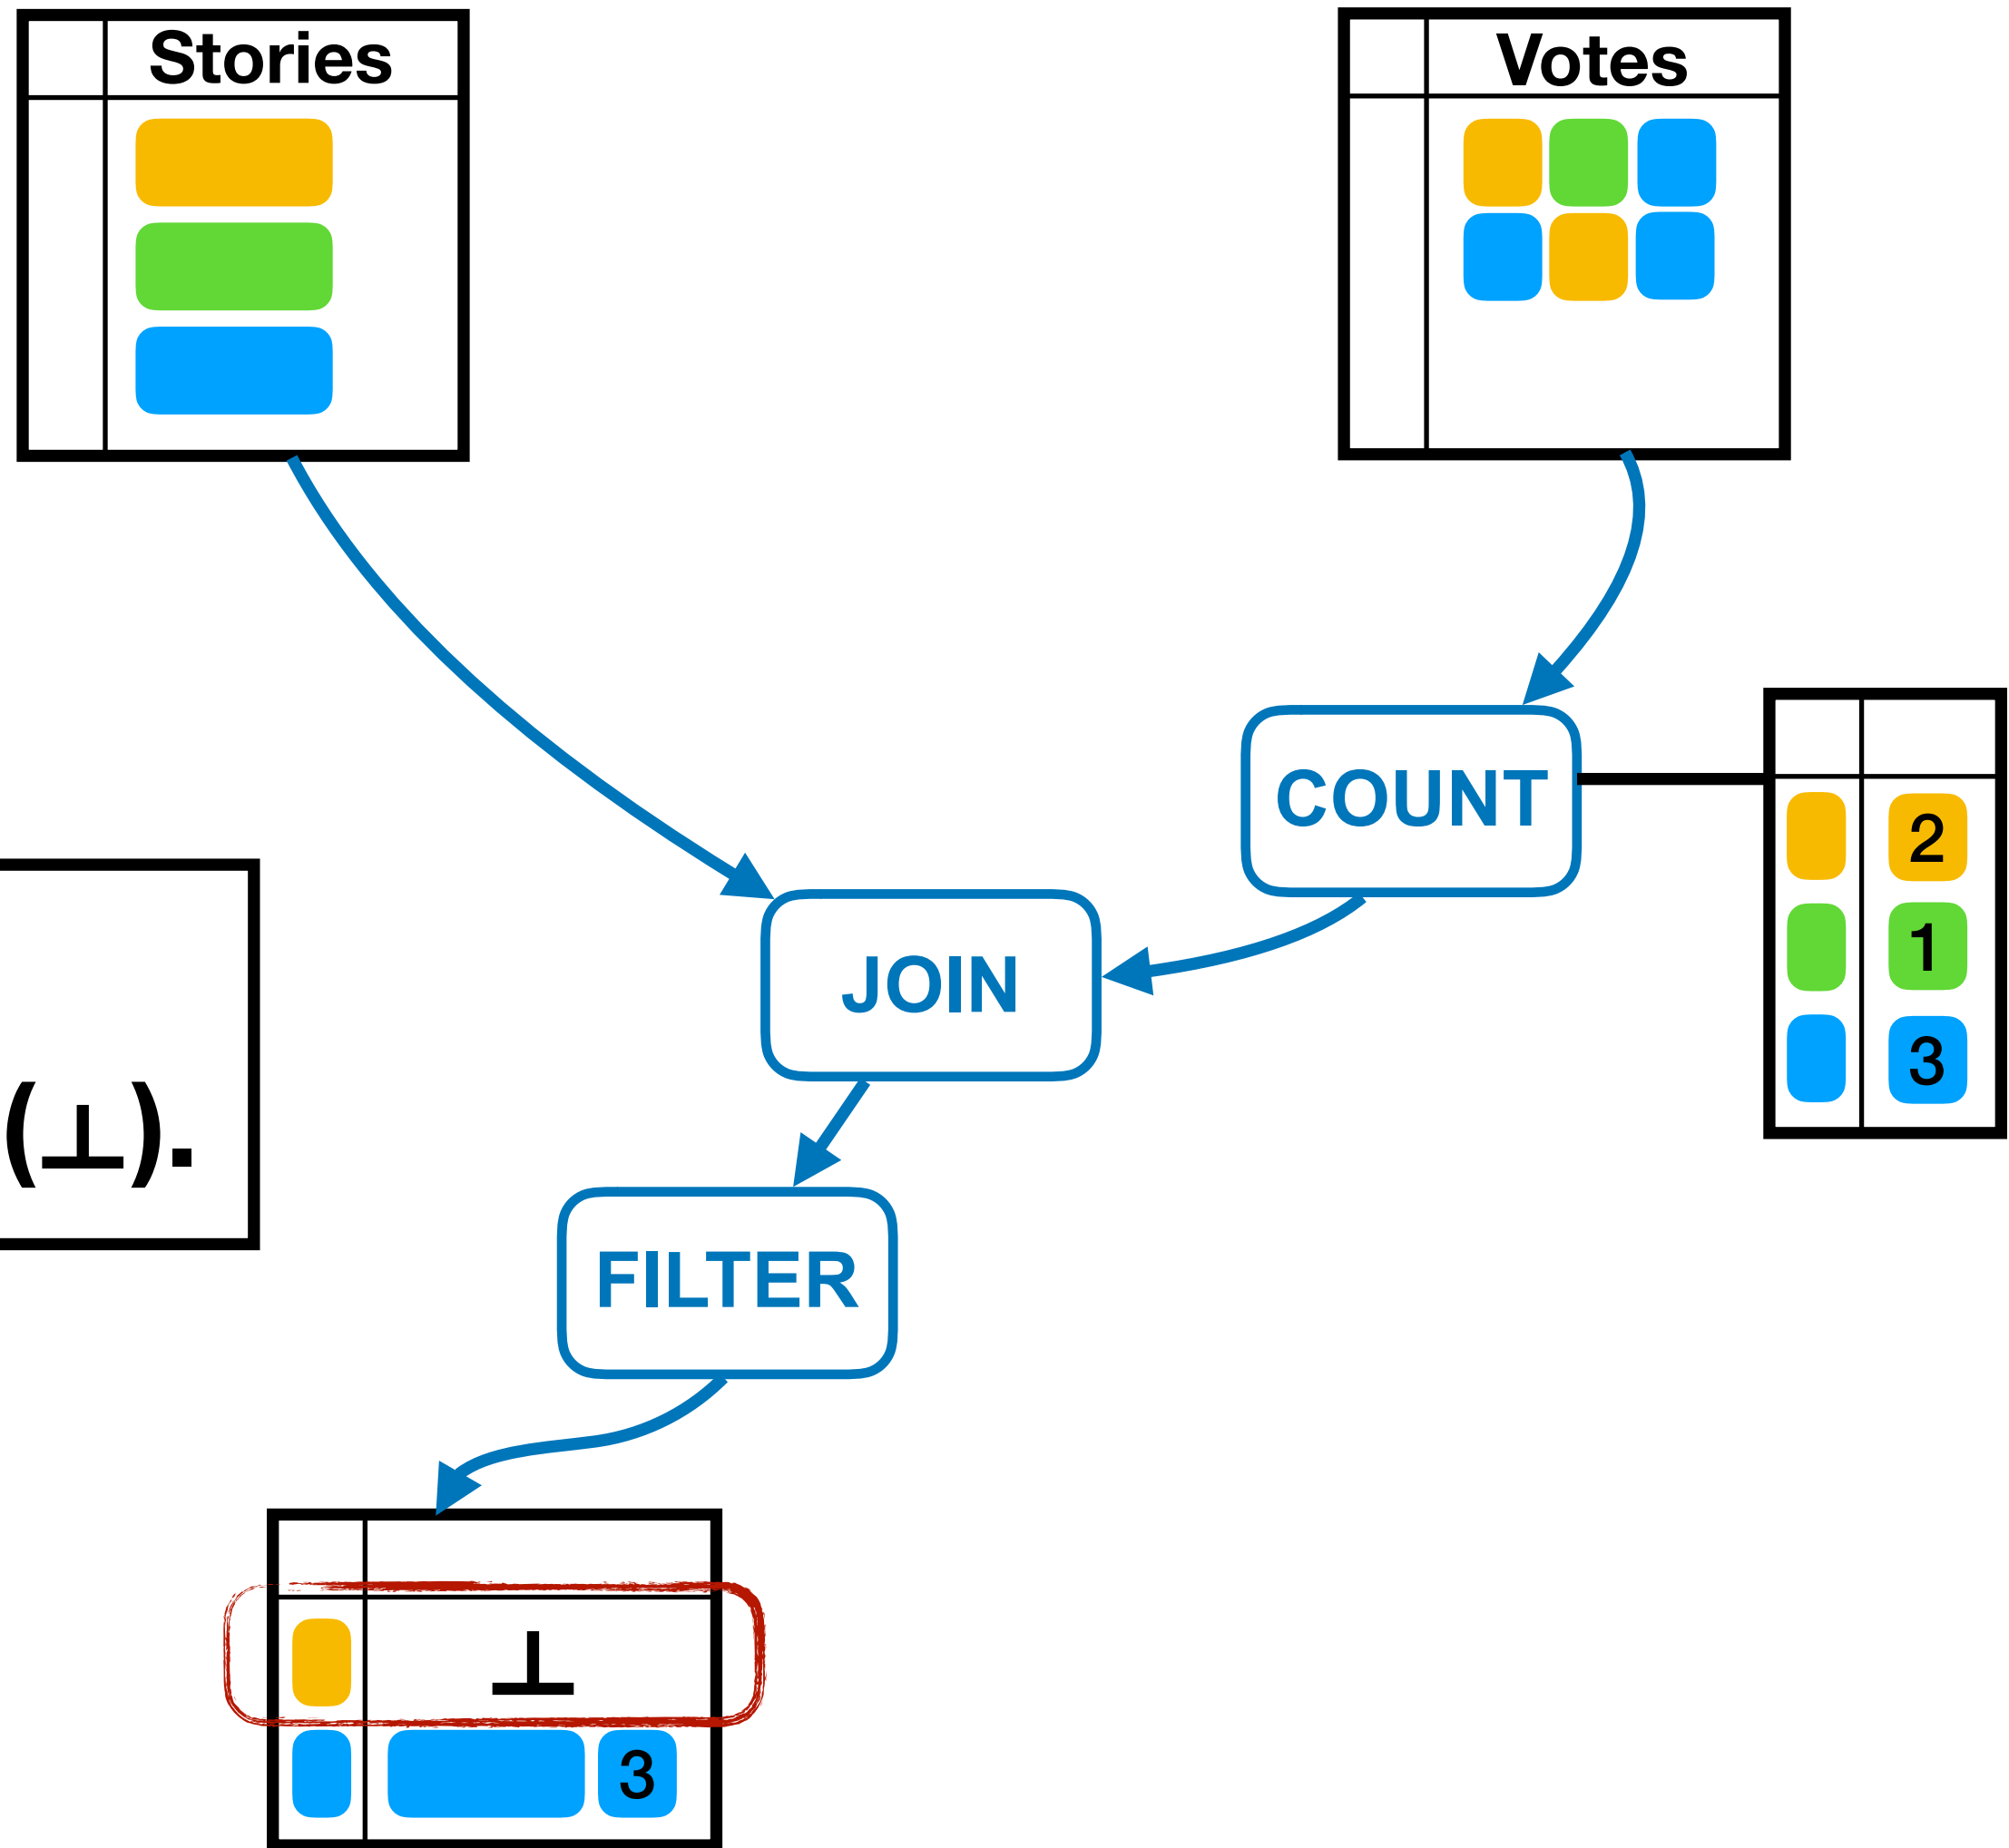
Data-flow state is *partial*: entries for some keys are absent ( $\perp$ ).



Frontend

# Partially-stateful data-flow


Data-flow state is *partial*: entries for some keys are absent ( $\perp$ ).



Frontend

# Partially-stateful data-flow

Data-flow state is *partial*: entries for some keys are absent ( $\perp$ ).

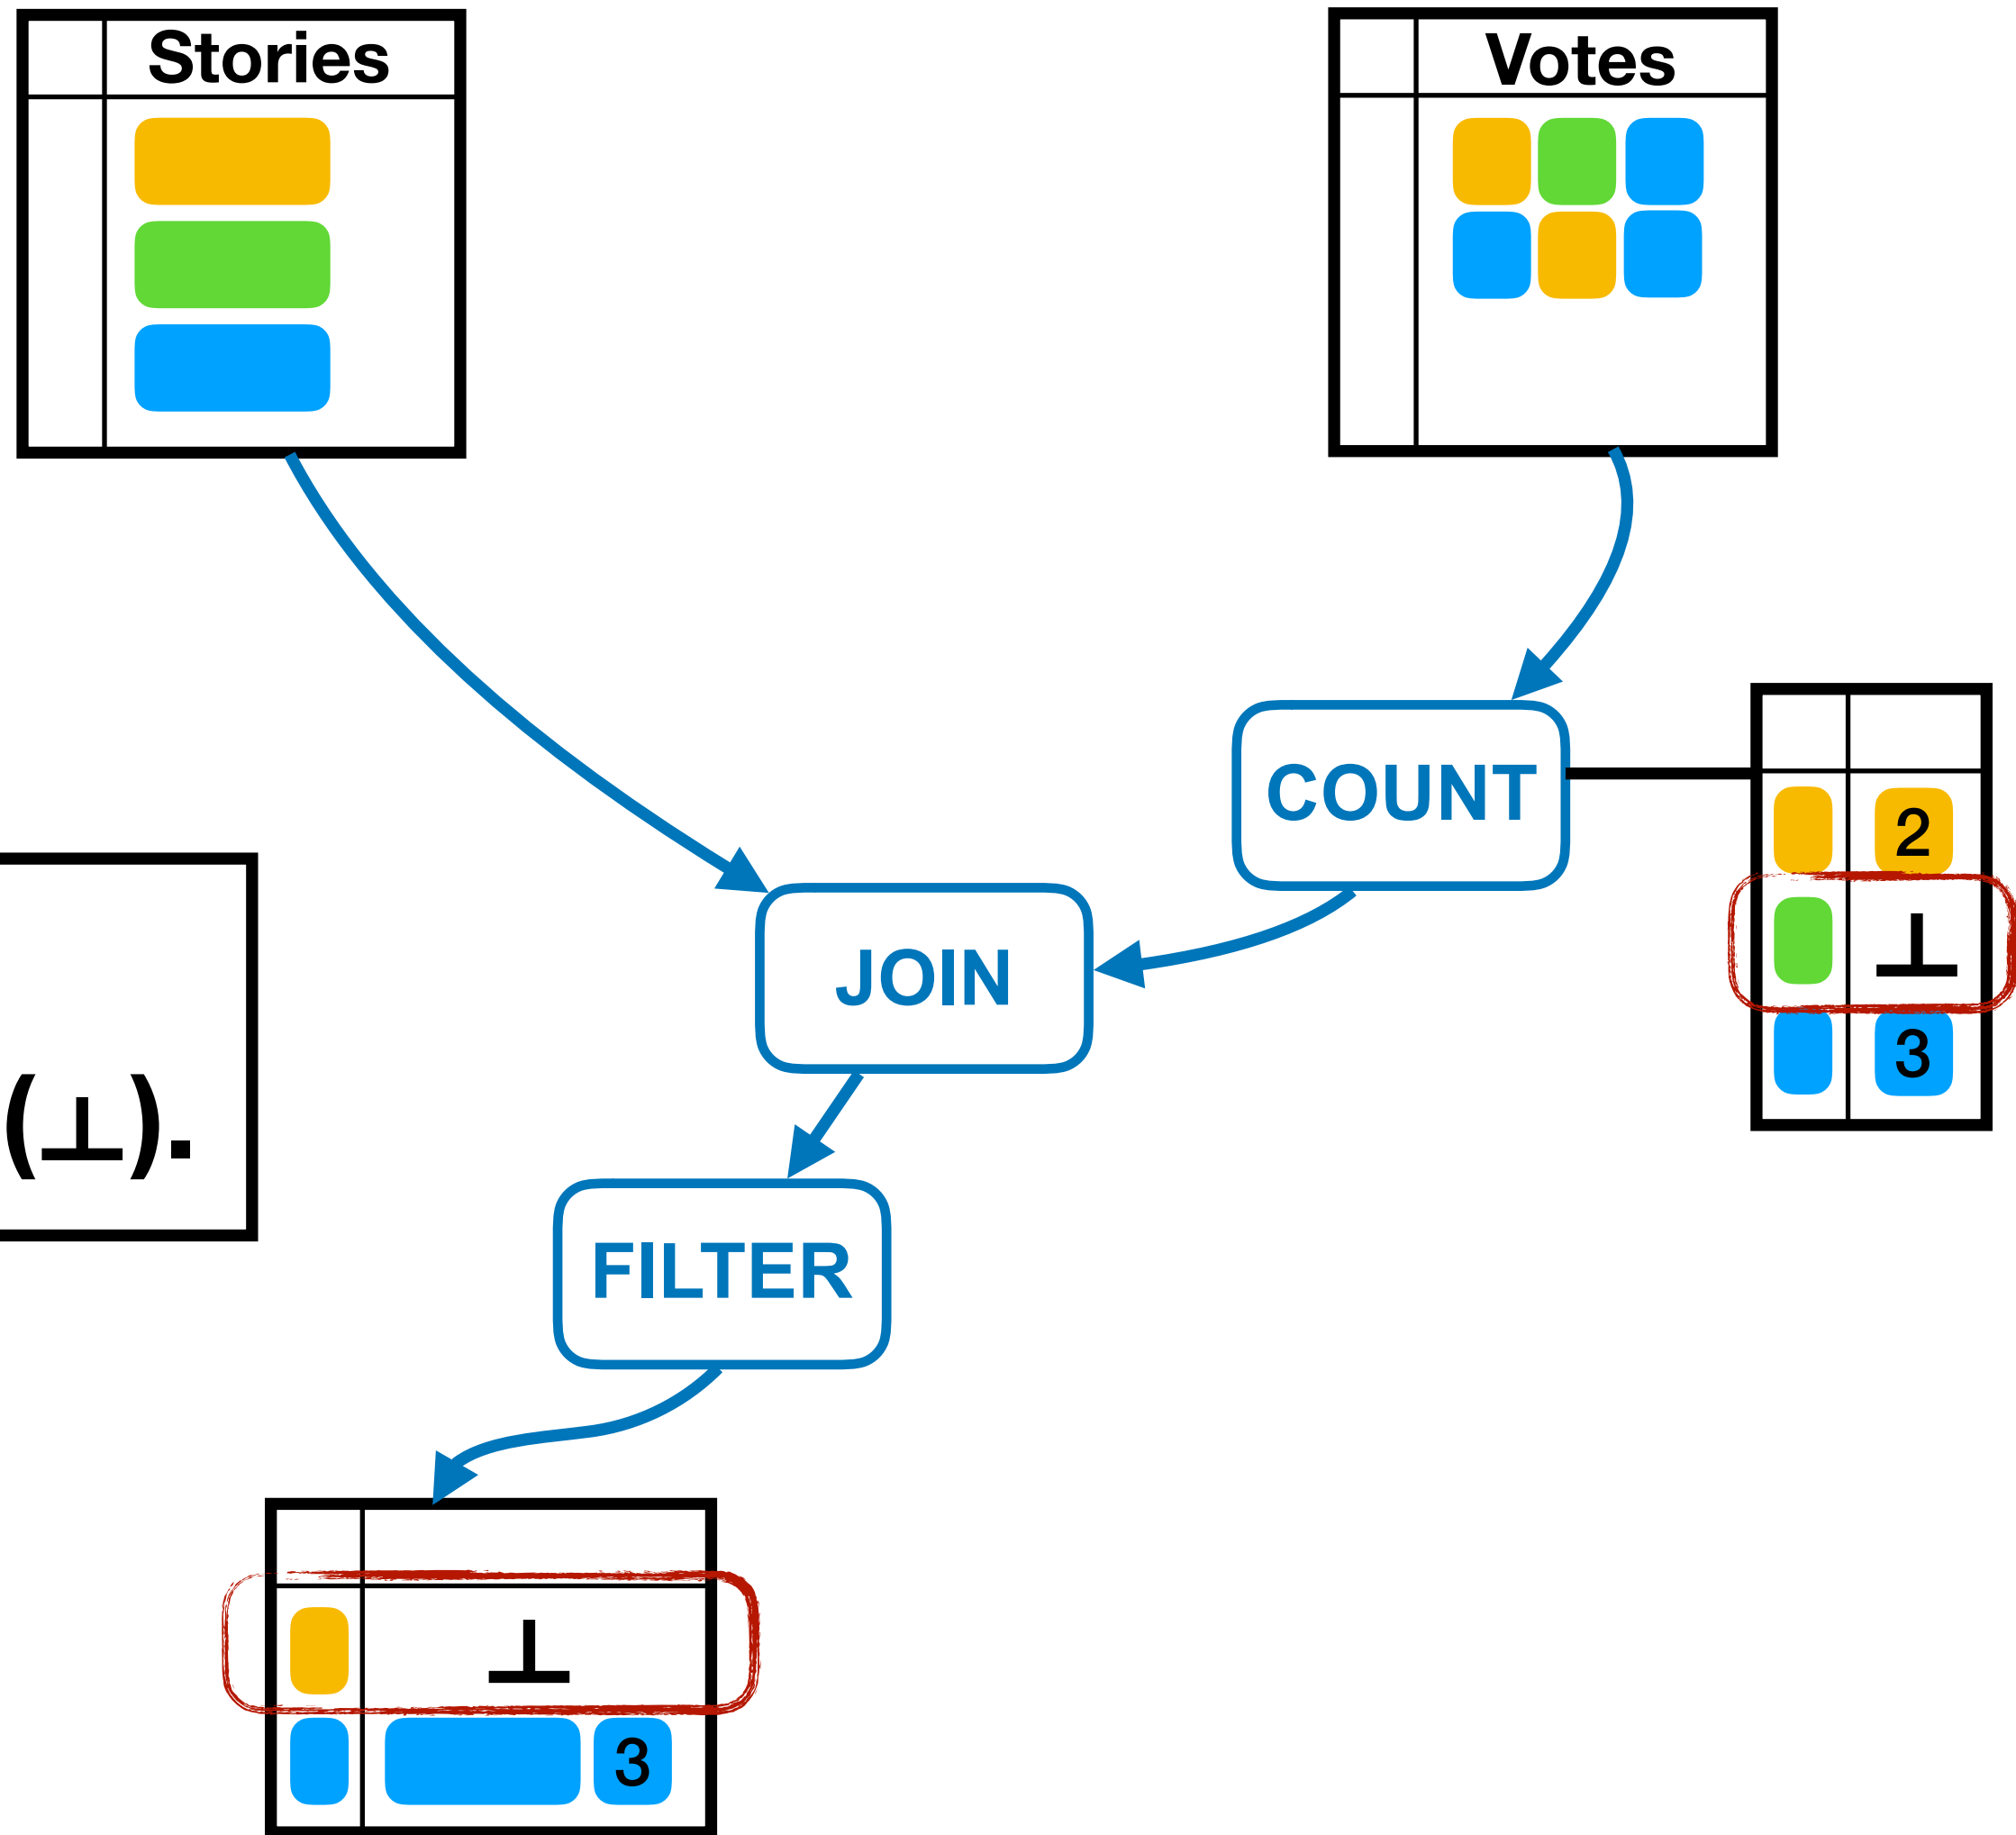


Frontend

# Partially-stateful data-flow

Data-flow state is *partial*: entries for some keys are absent ( $\perp$ ).

Lower memory footprint.

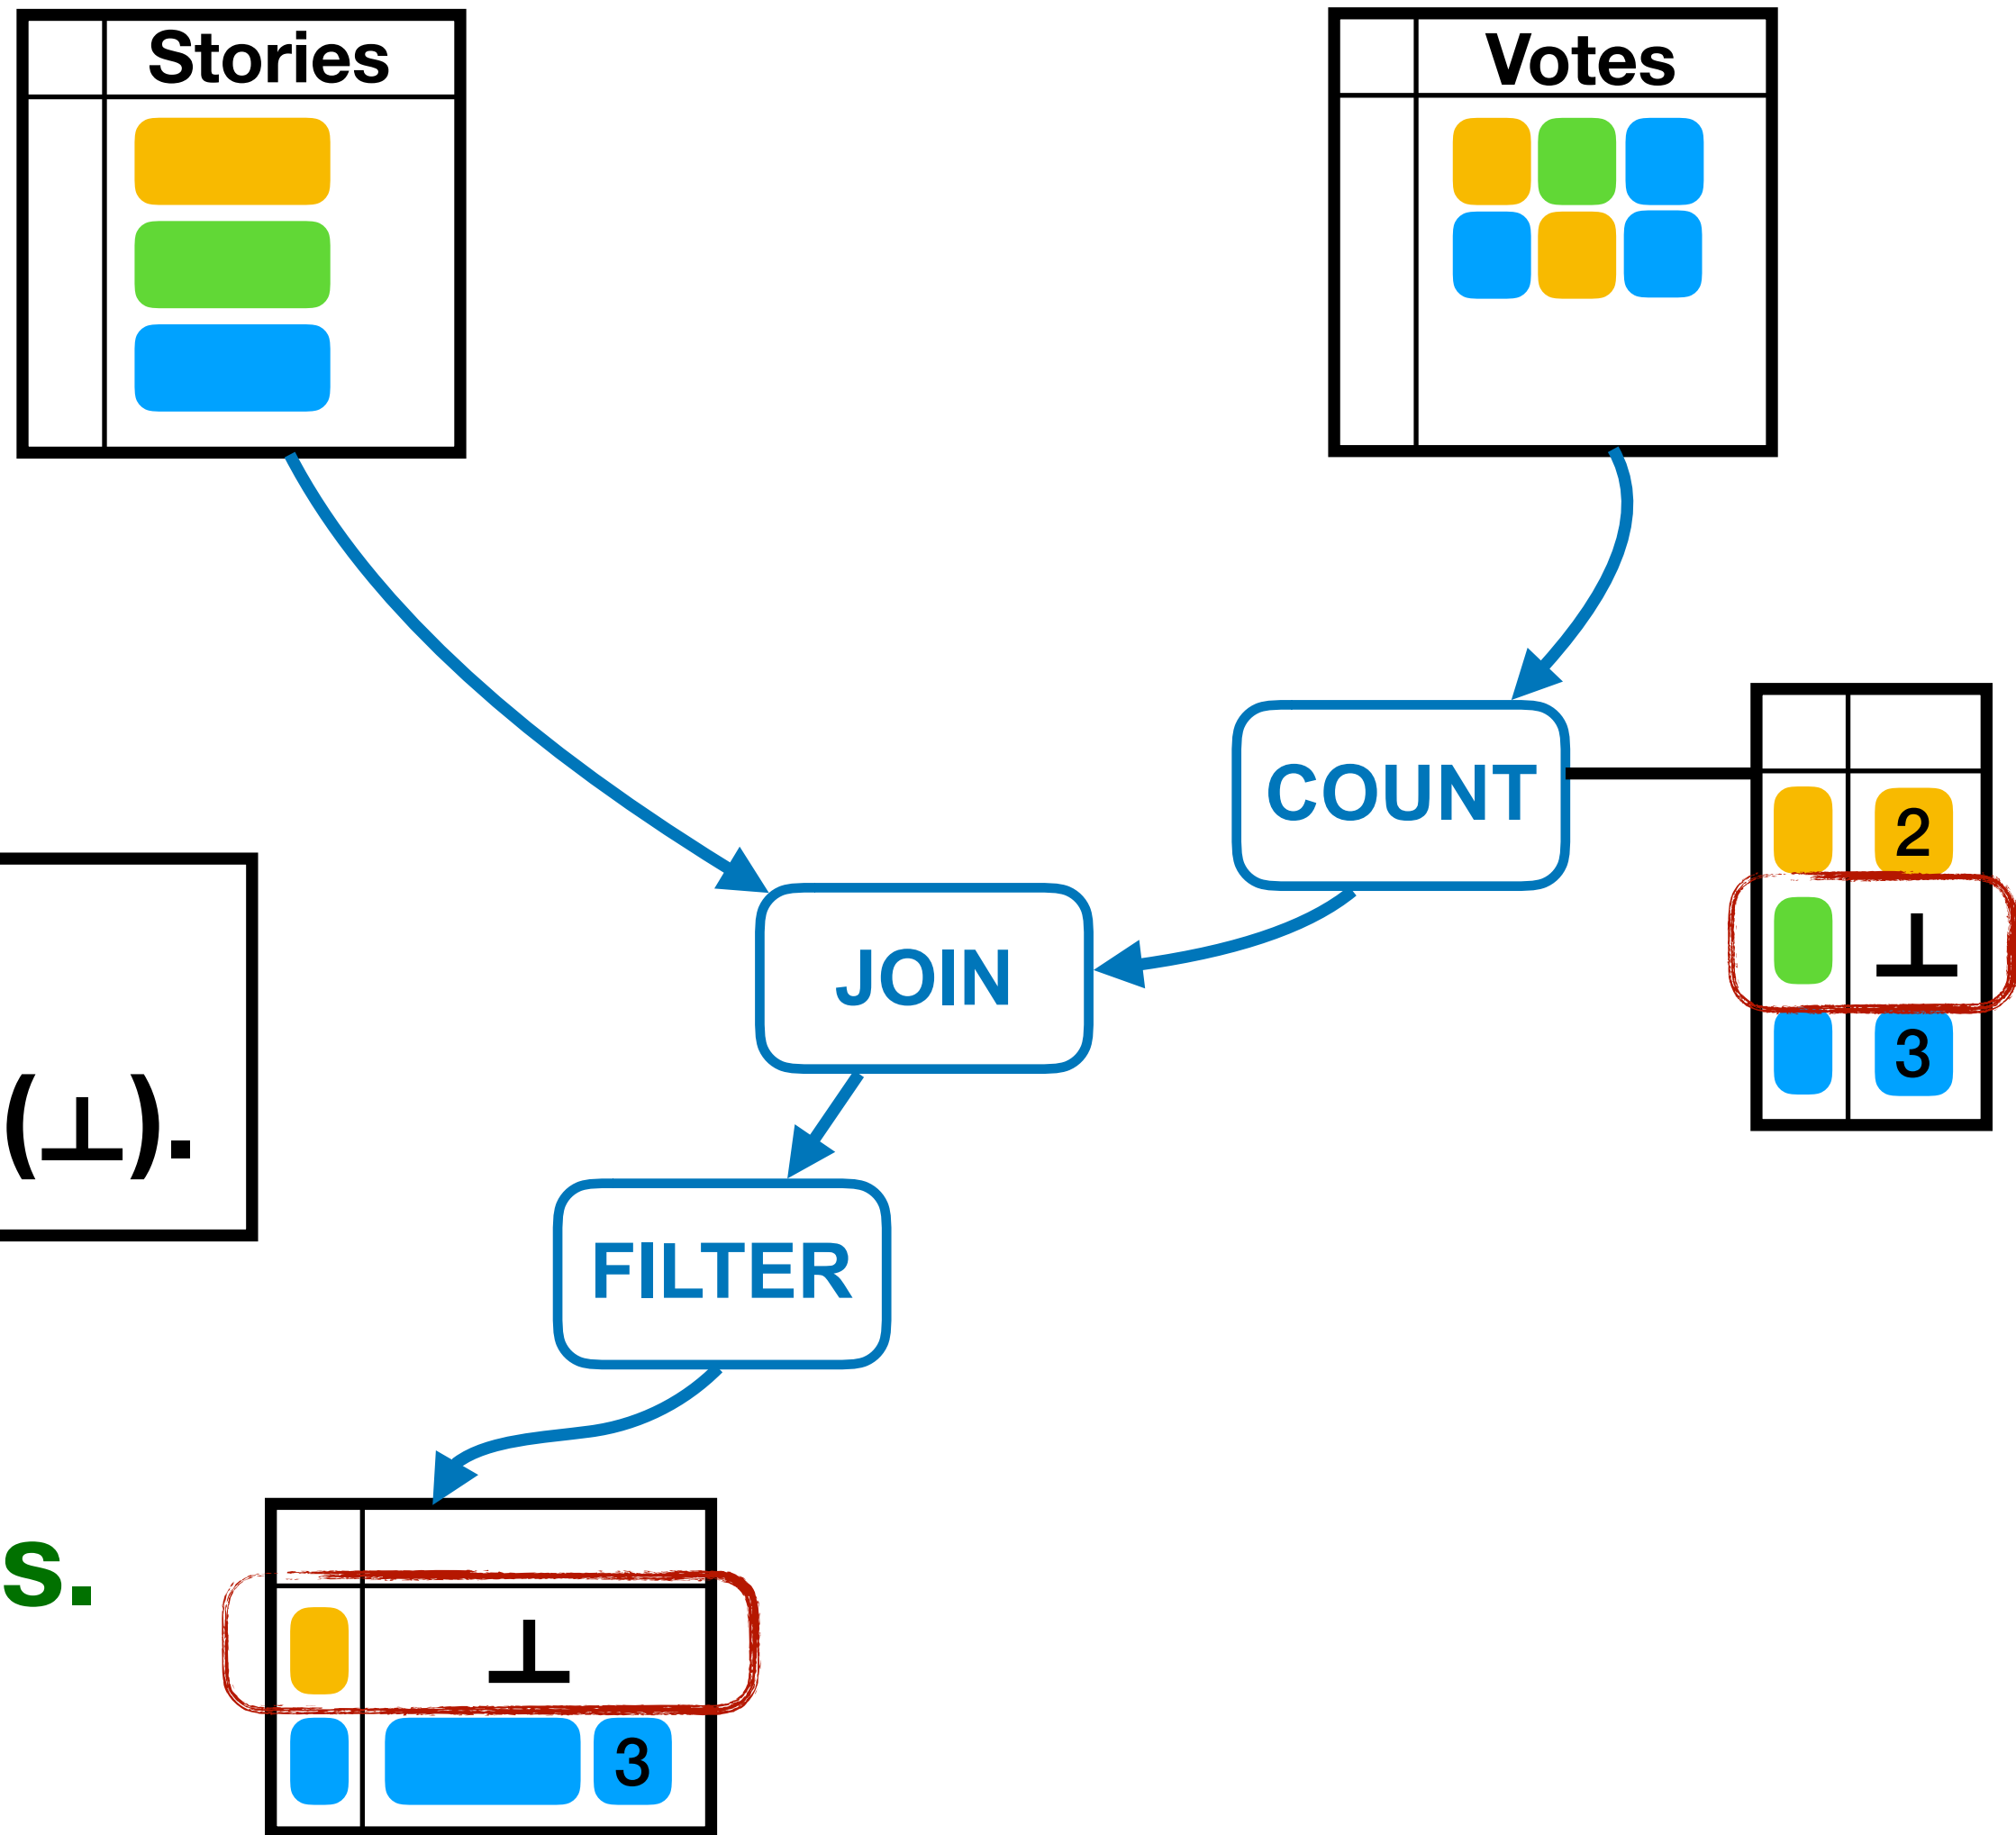


Frontend

# Partially-stateful data-flow

Data-flow state is *partial*: entries for some keys are absent ( $\perp$ ).

Lower memory footprint.  
No need to update absent entries.

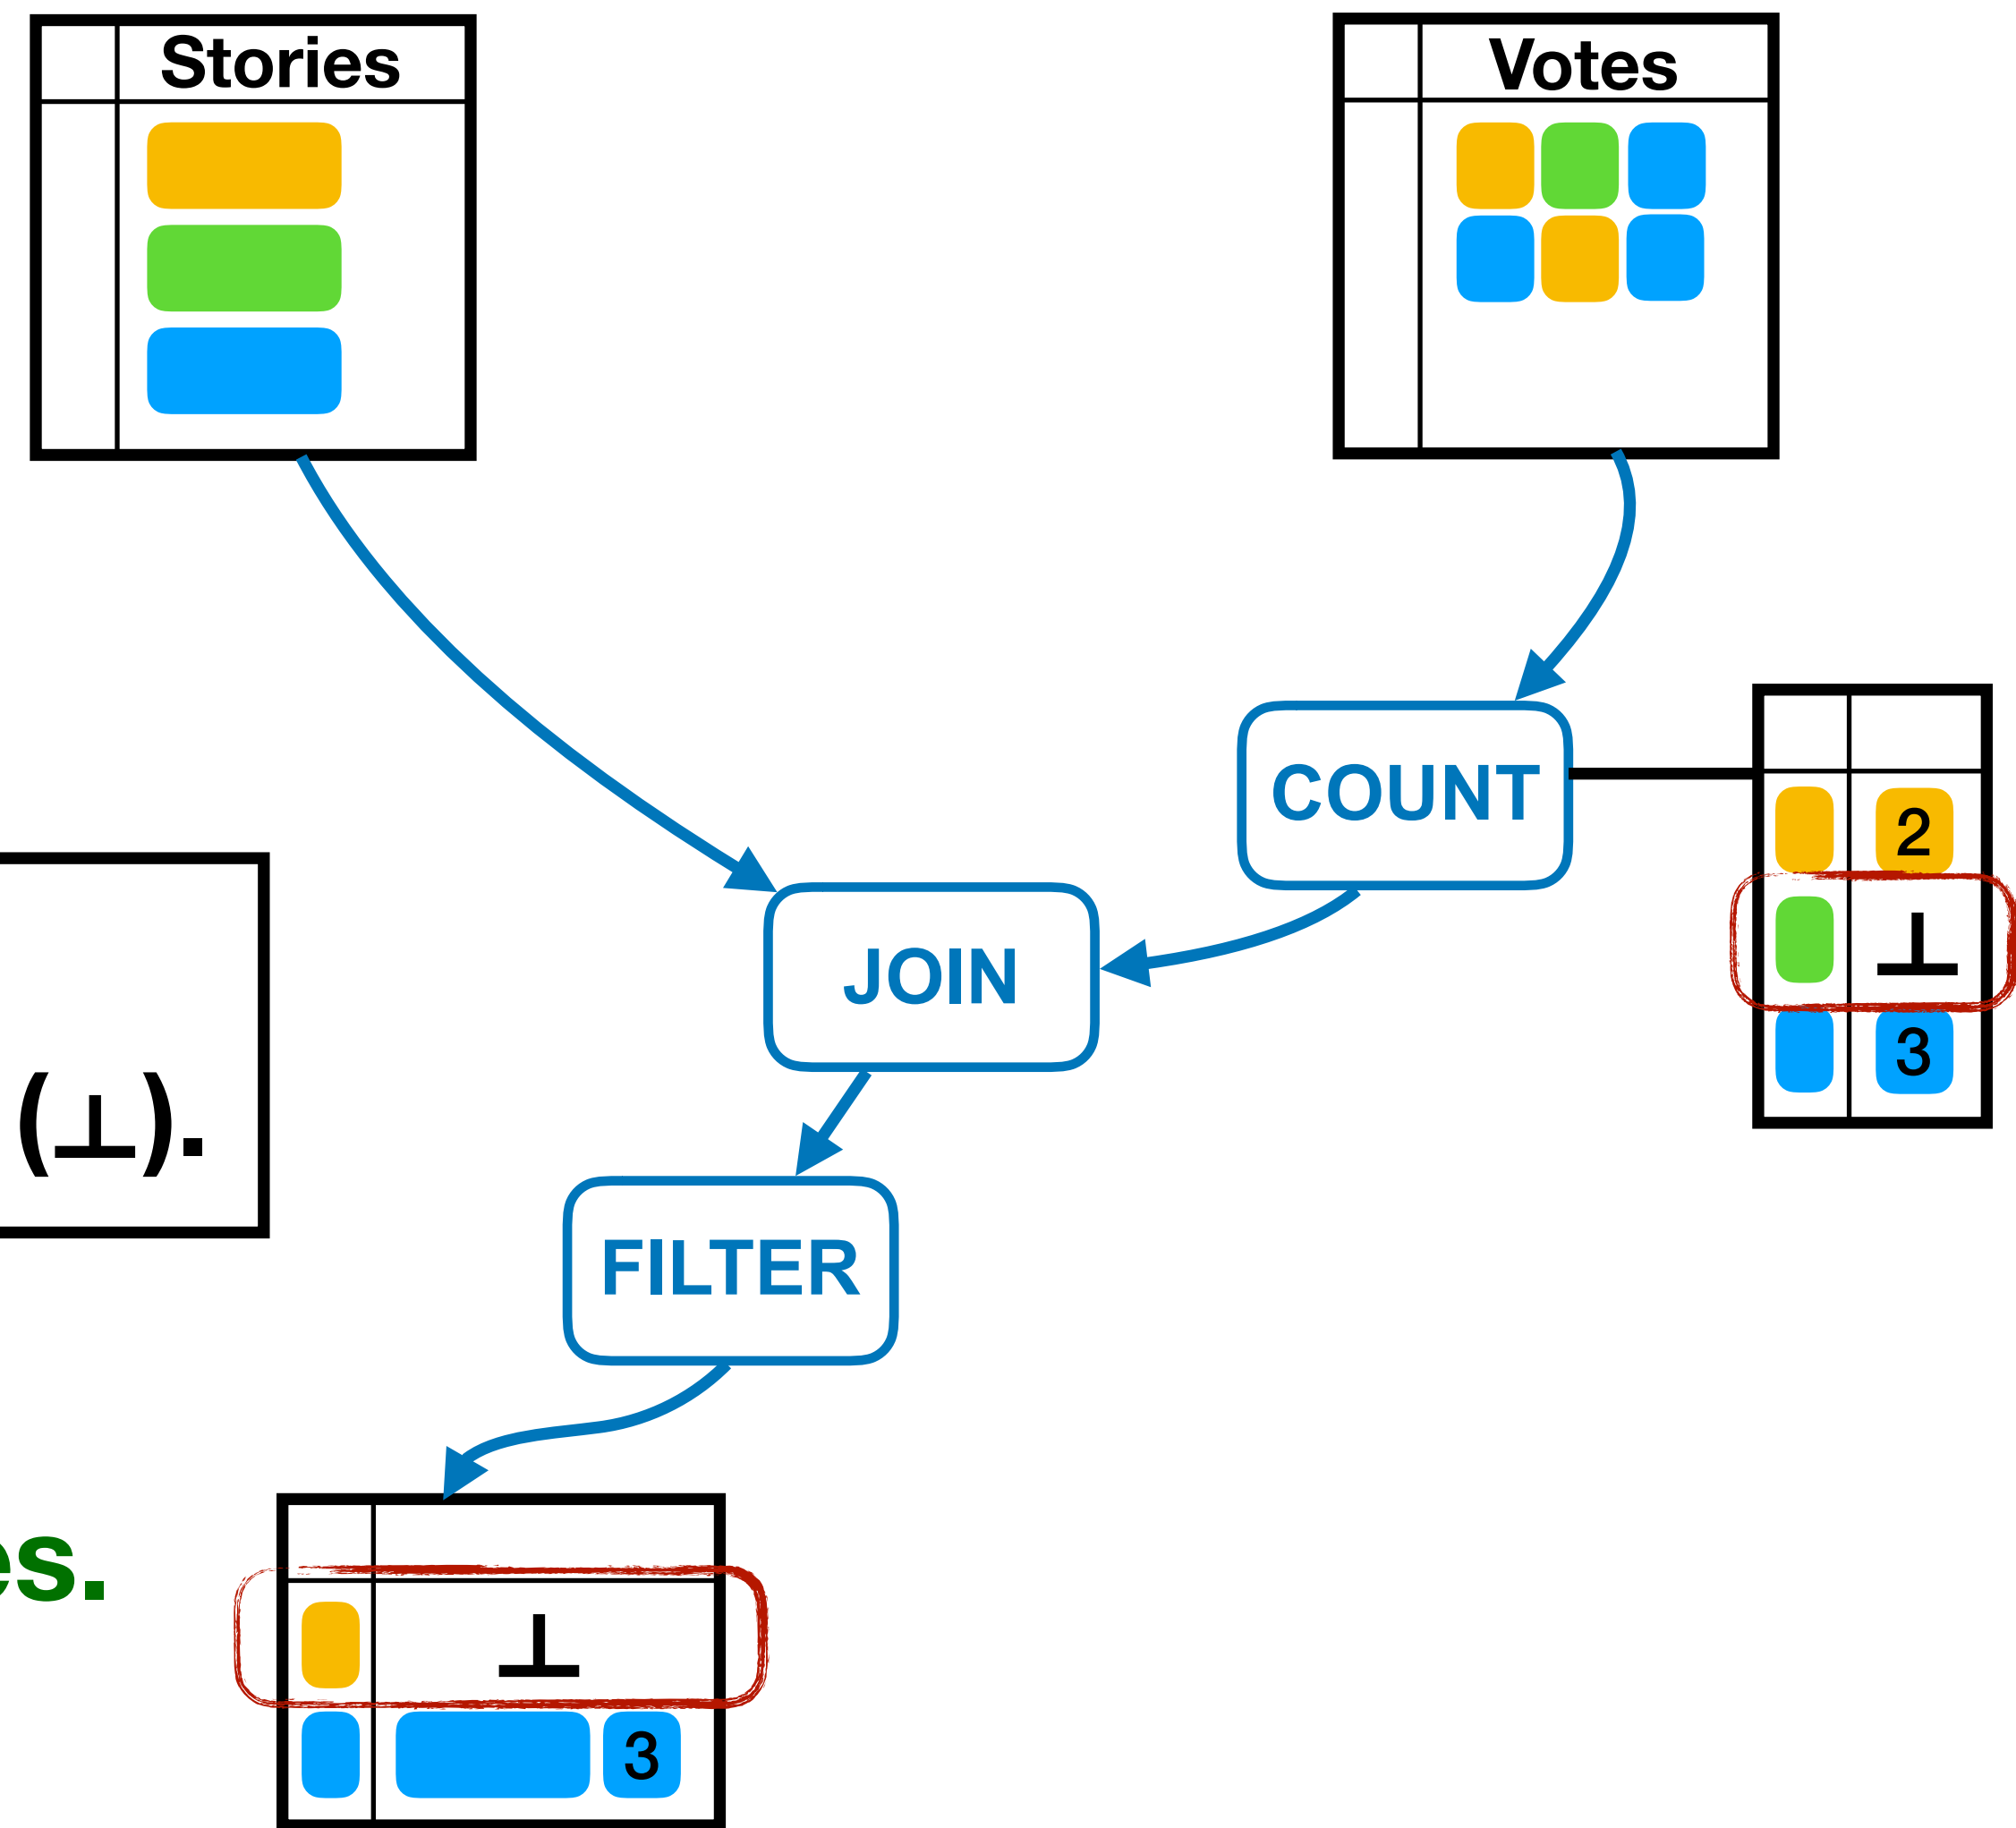


Frontend

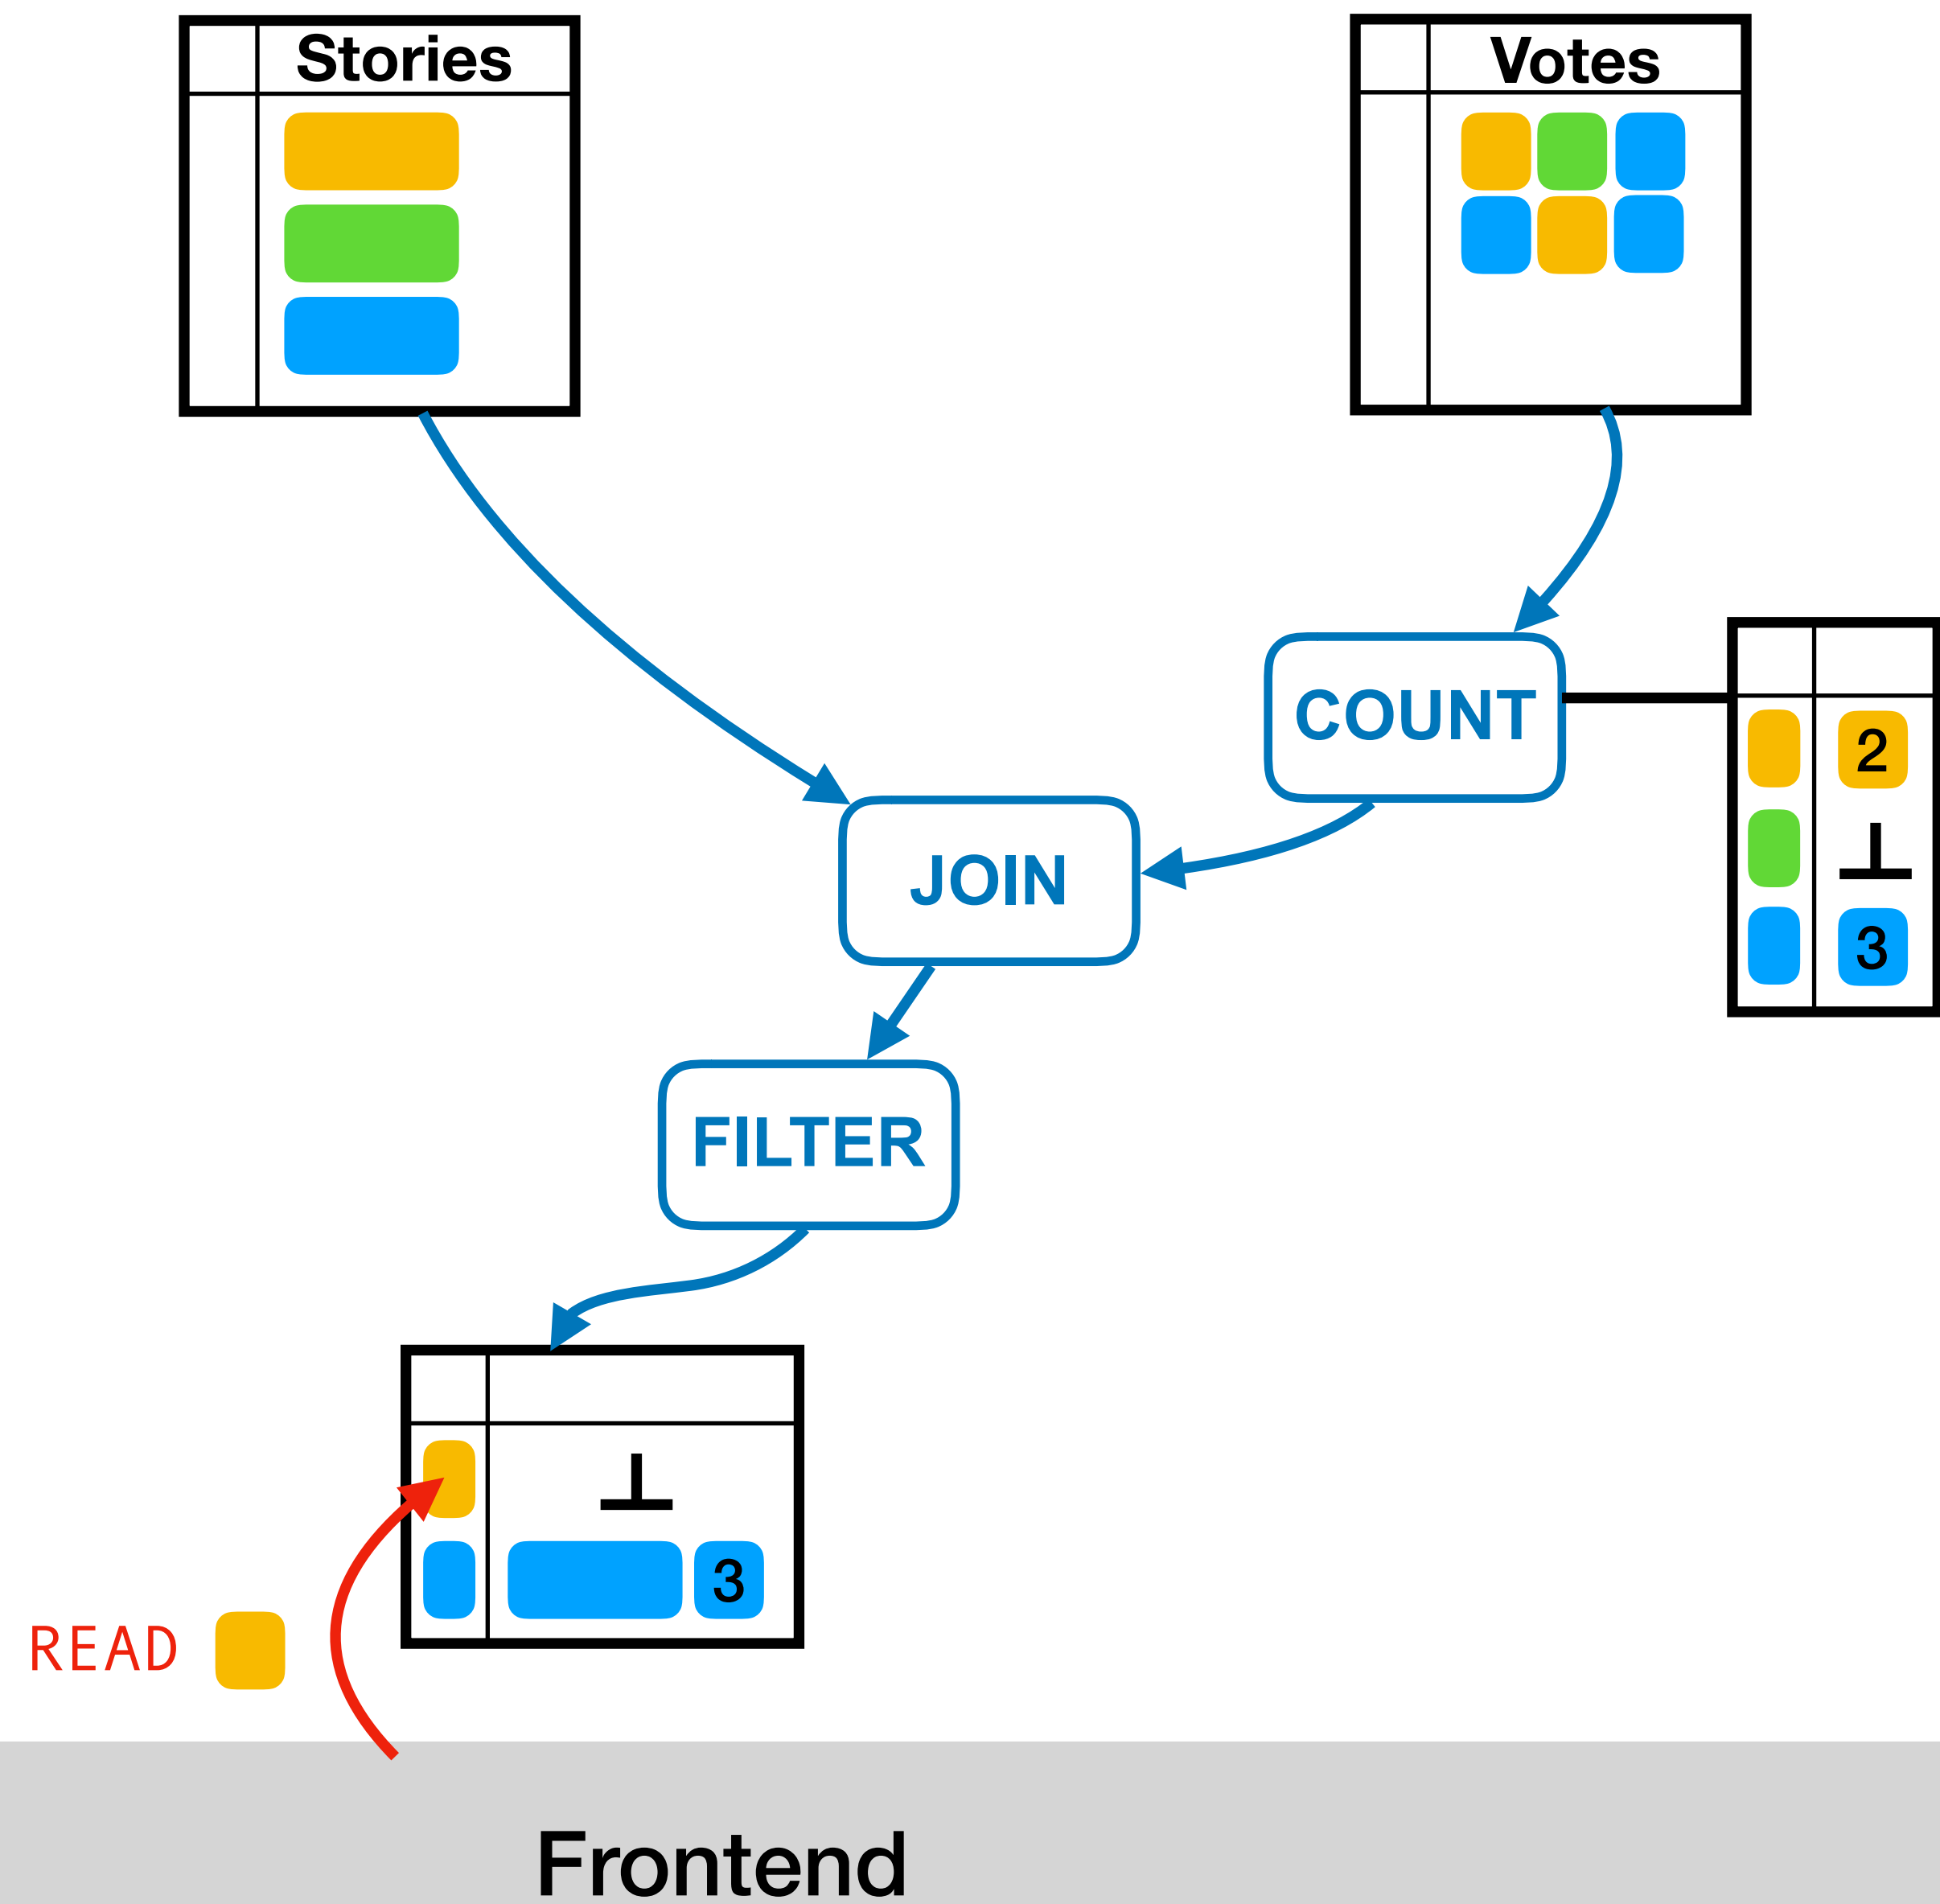
# Partially-stateful data-flow

Data-flow state is *partial*: entries for some keys are absent ( $\perp$ ).

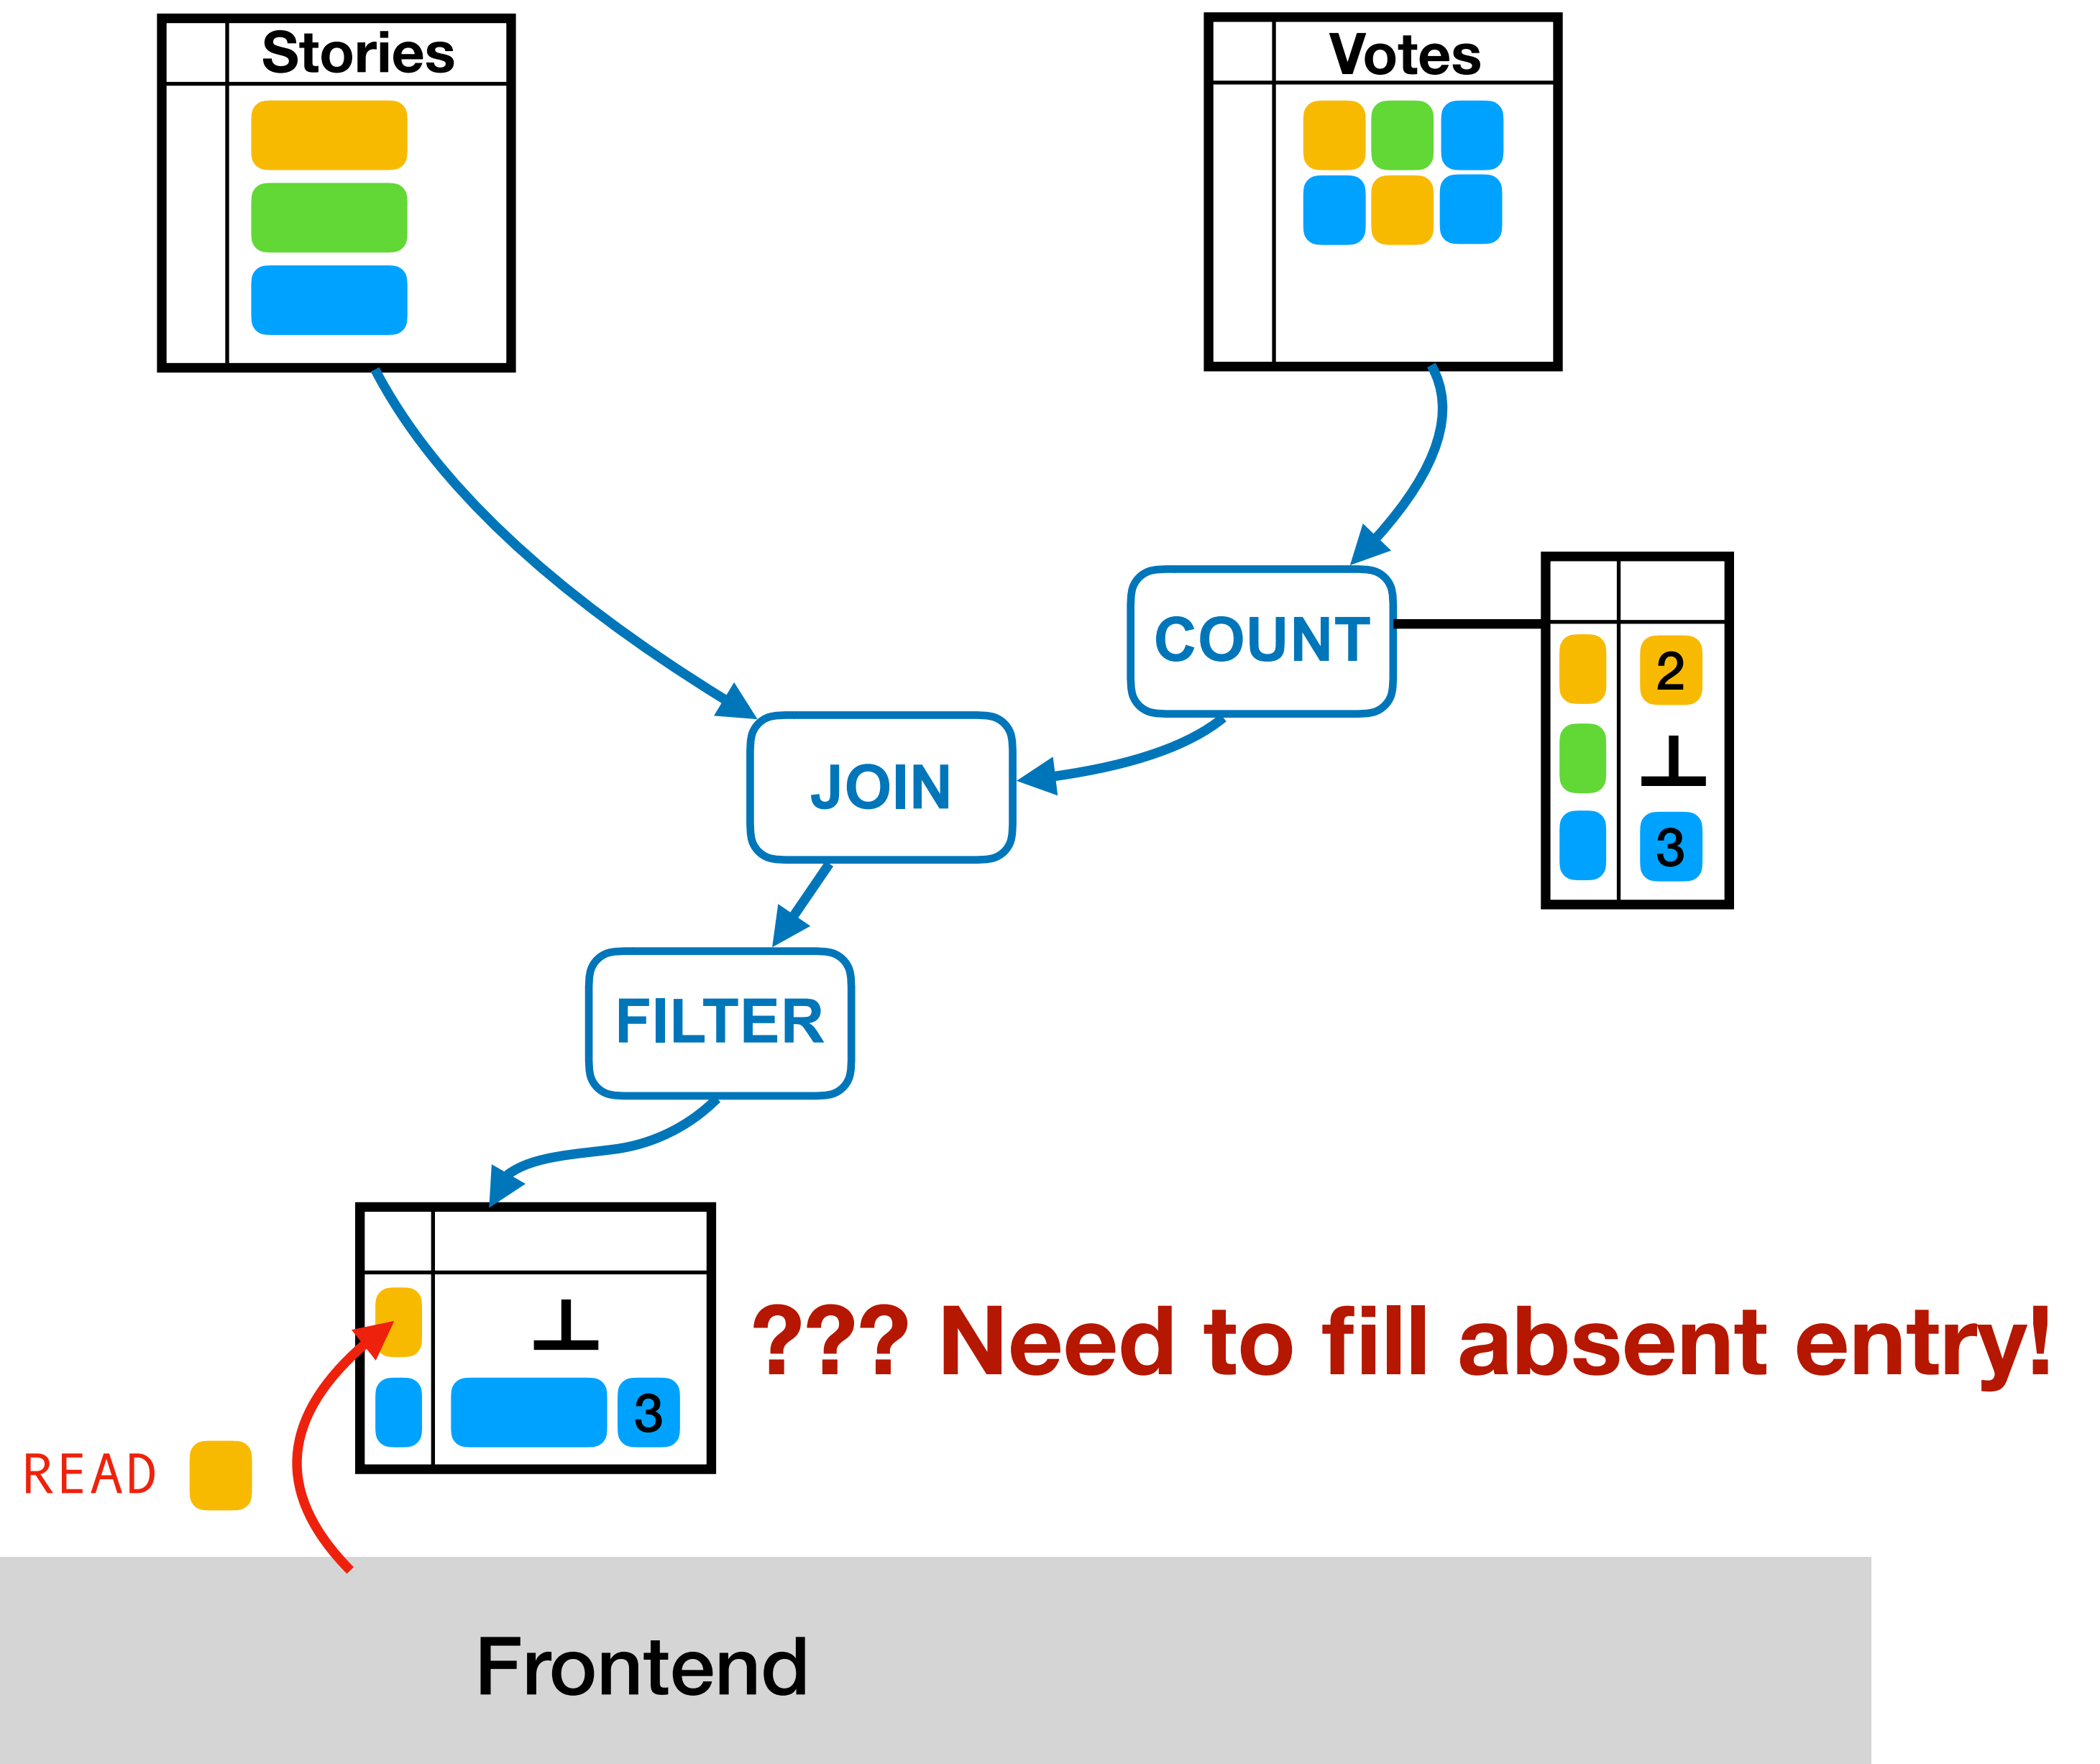
Lower memory footprint.  
No need to update absent entries.  
Enables live data-flow changes.



# Partially-stateful data-flow: upqueries



# Partially-stateful data-flow: upqueries

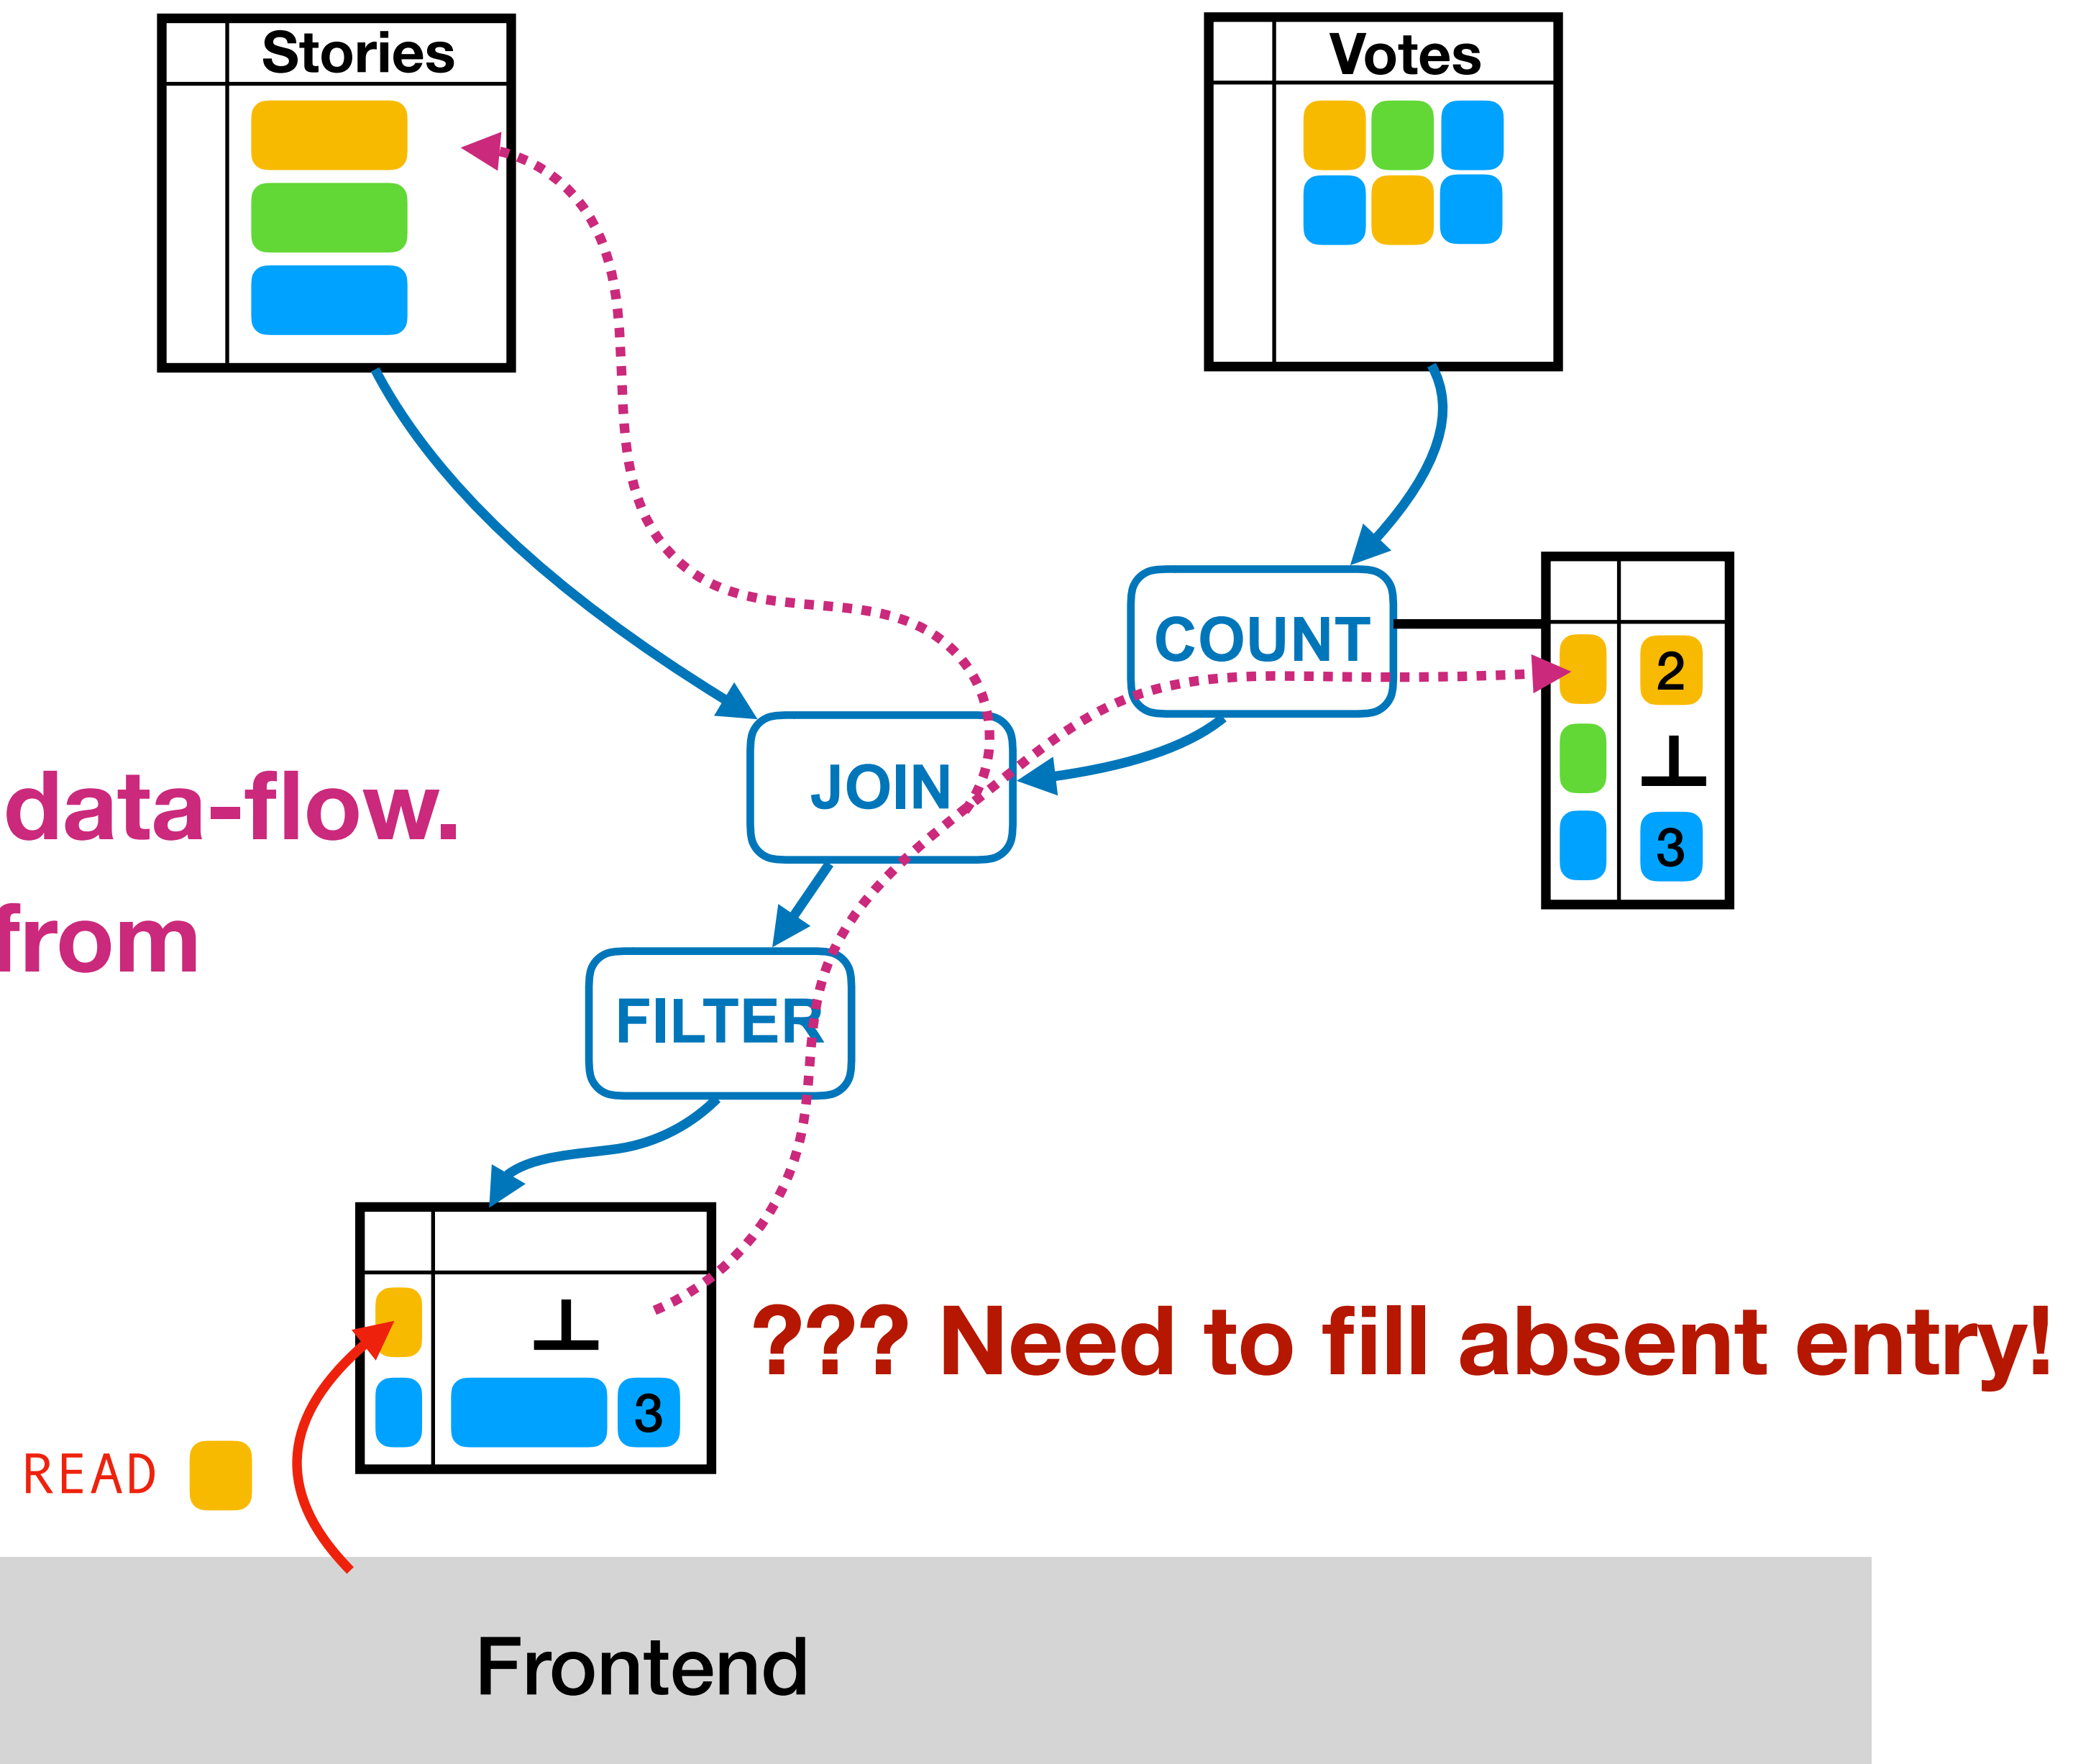




# Partially-stateful data-flow: upqueries

Solution: *upquery* through data-flow.

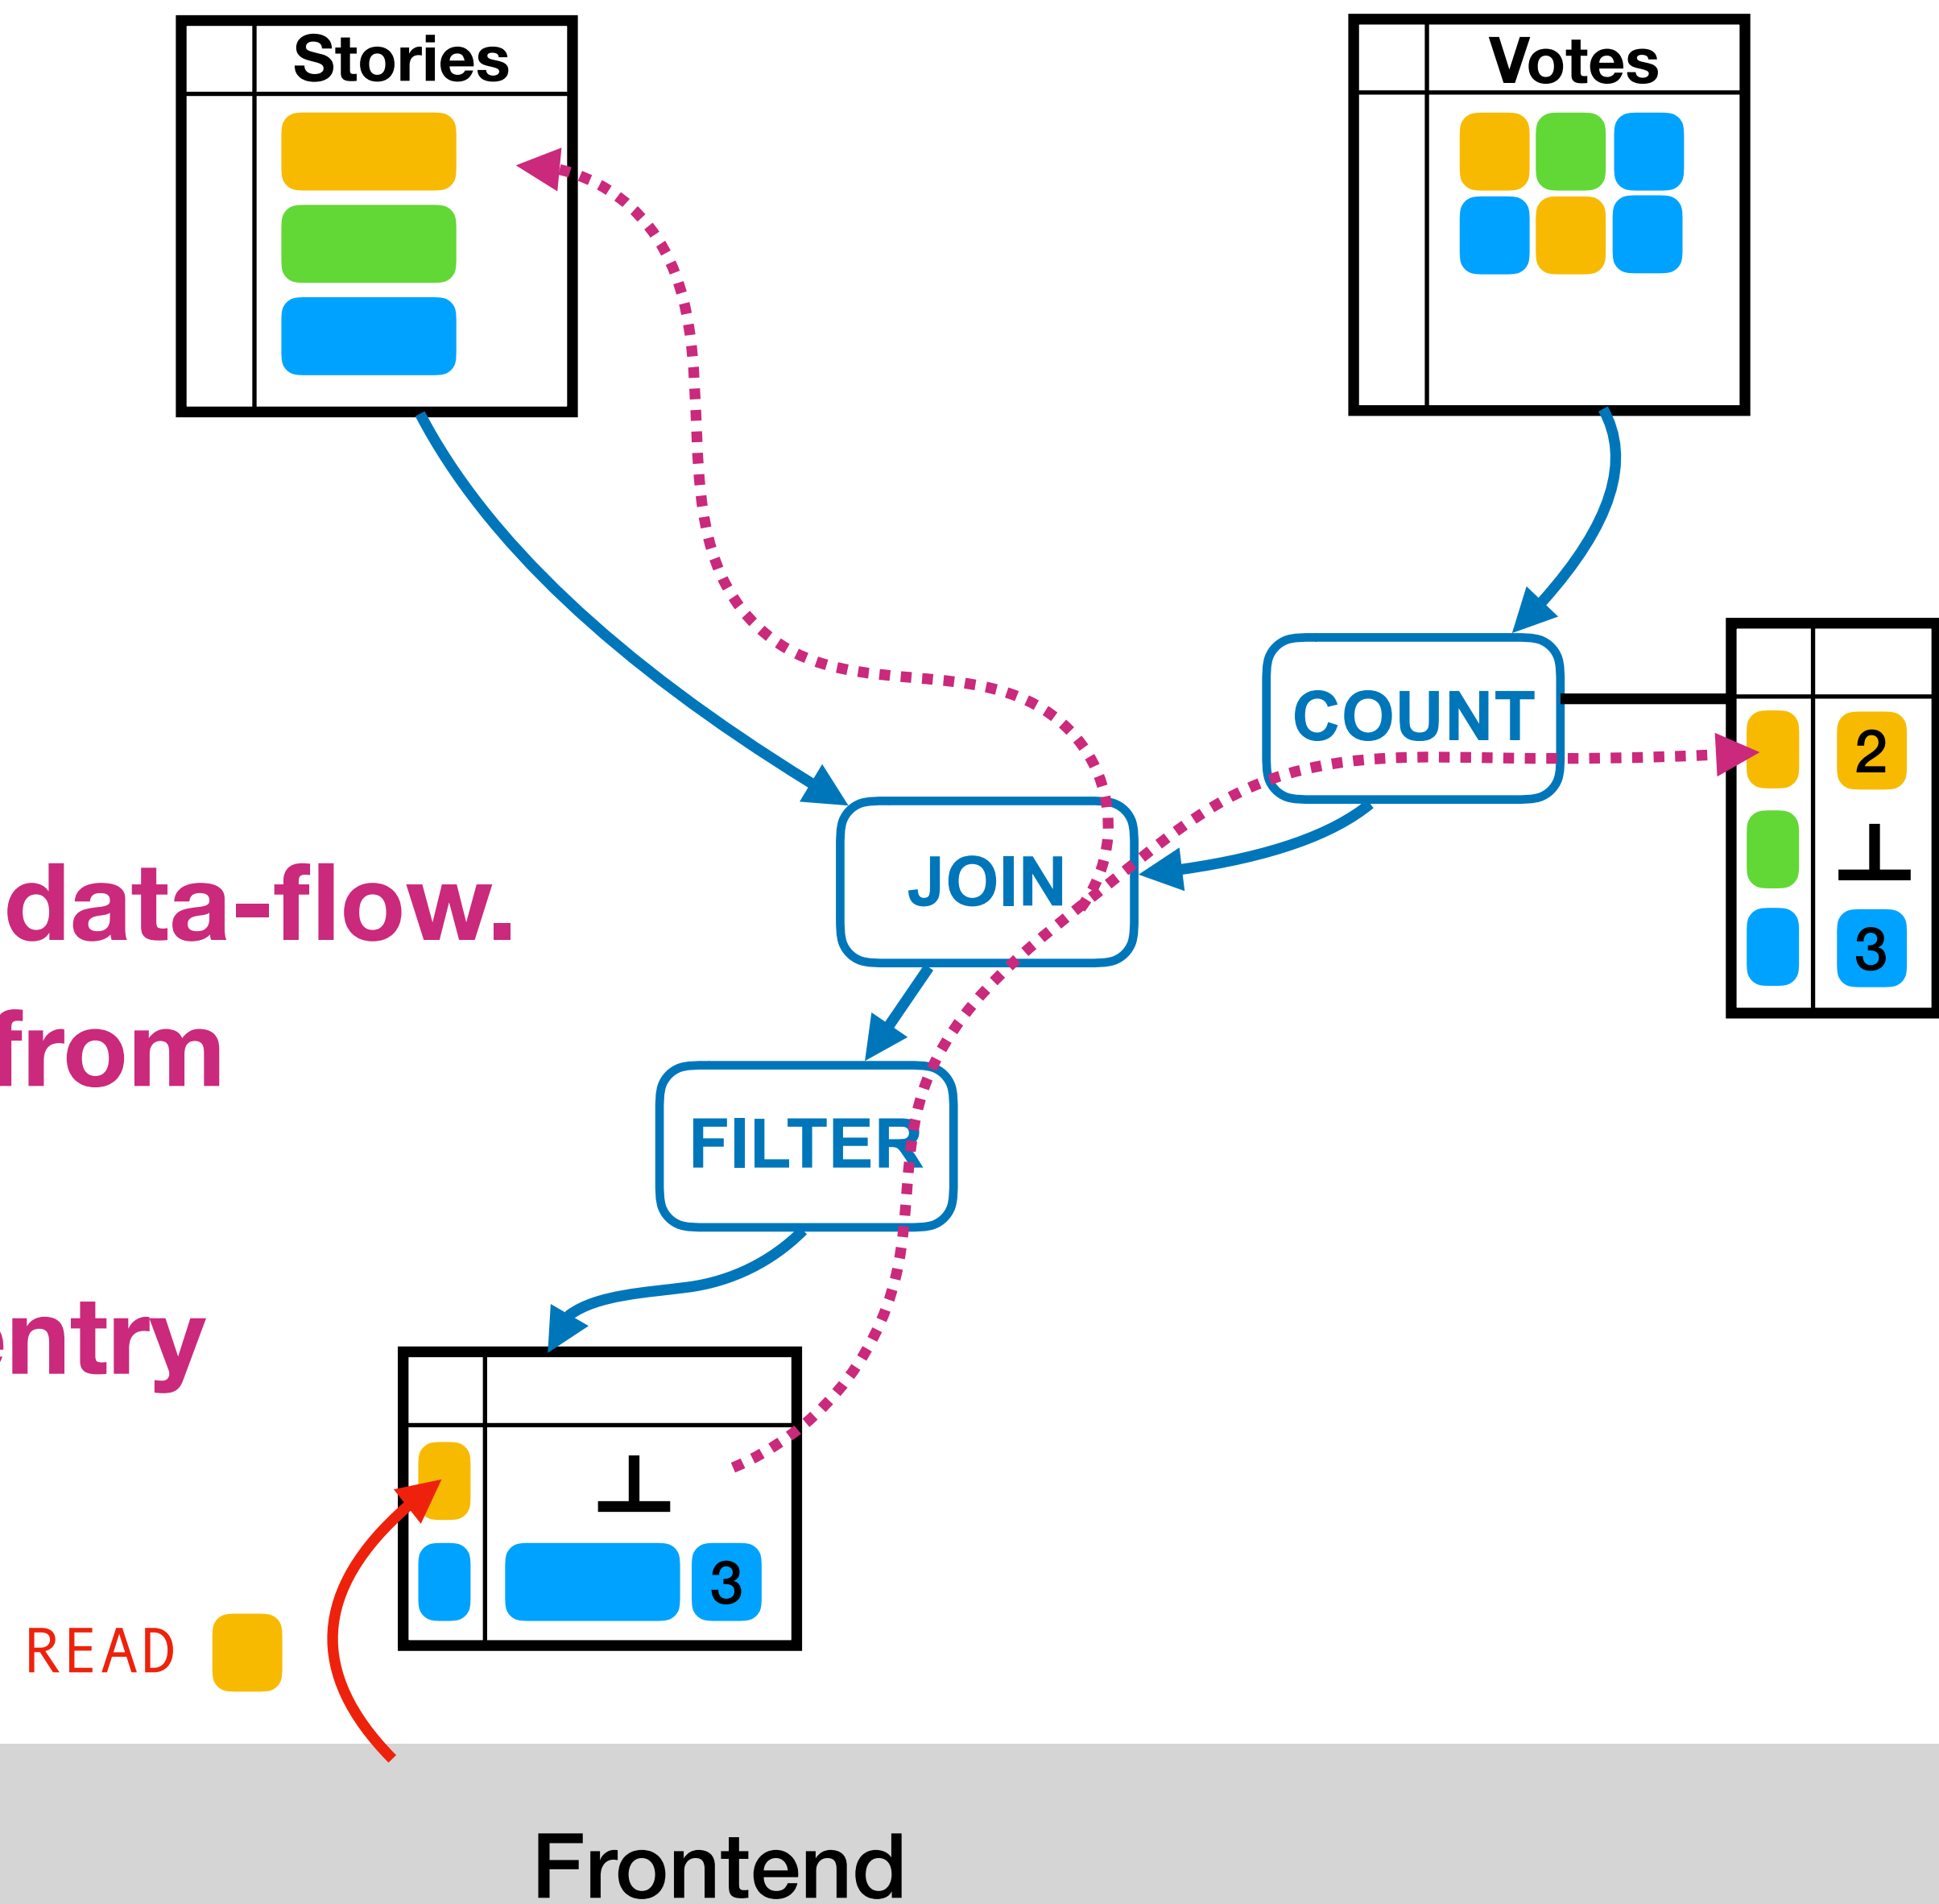
- Compute missing entry from upstream state



# Partially-stateful data-flow: upqueries

Solution: *upquery* through data-flow.

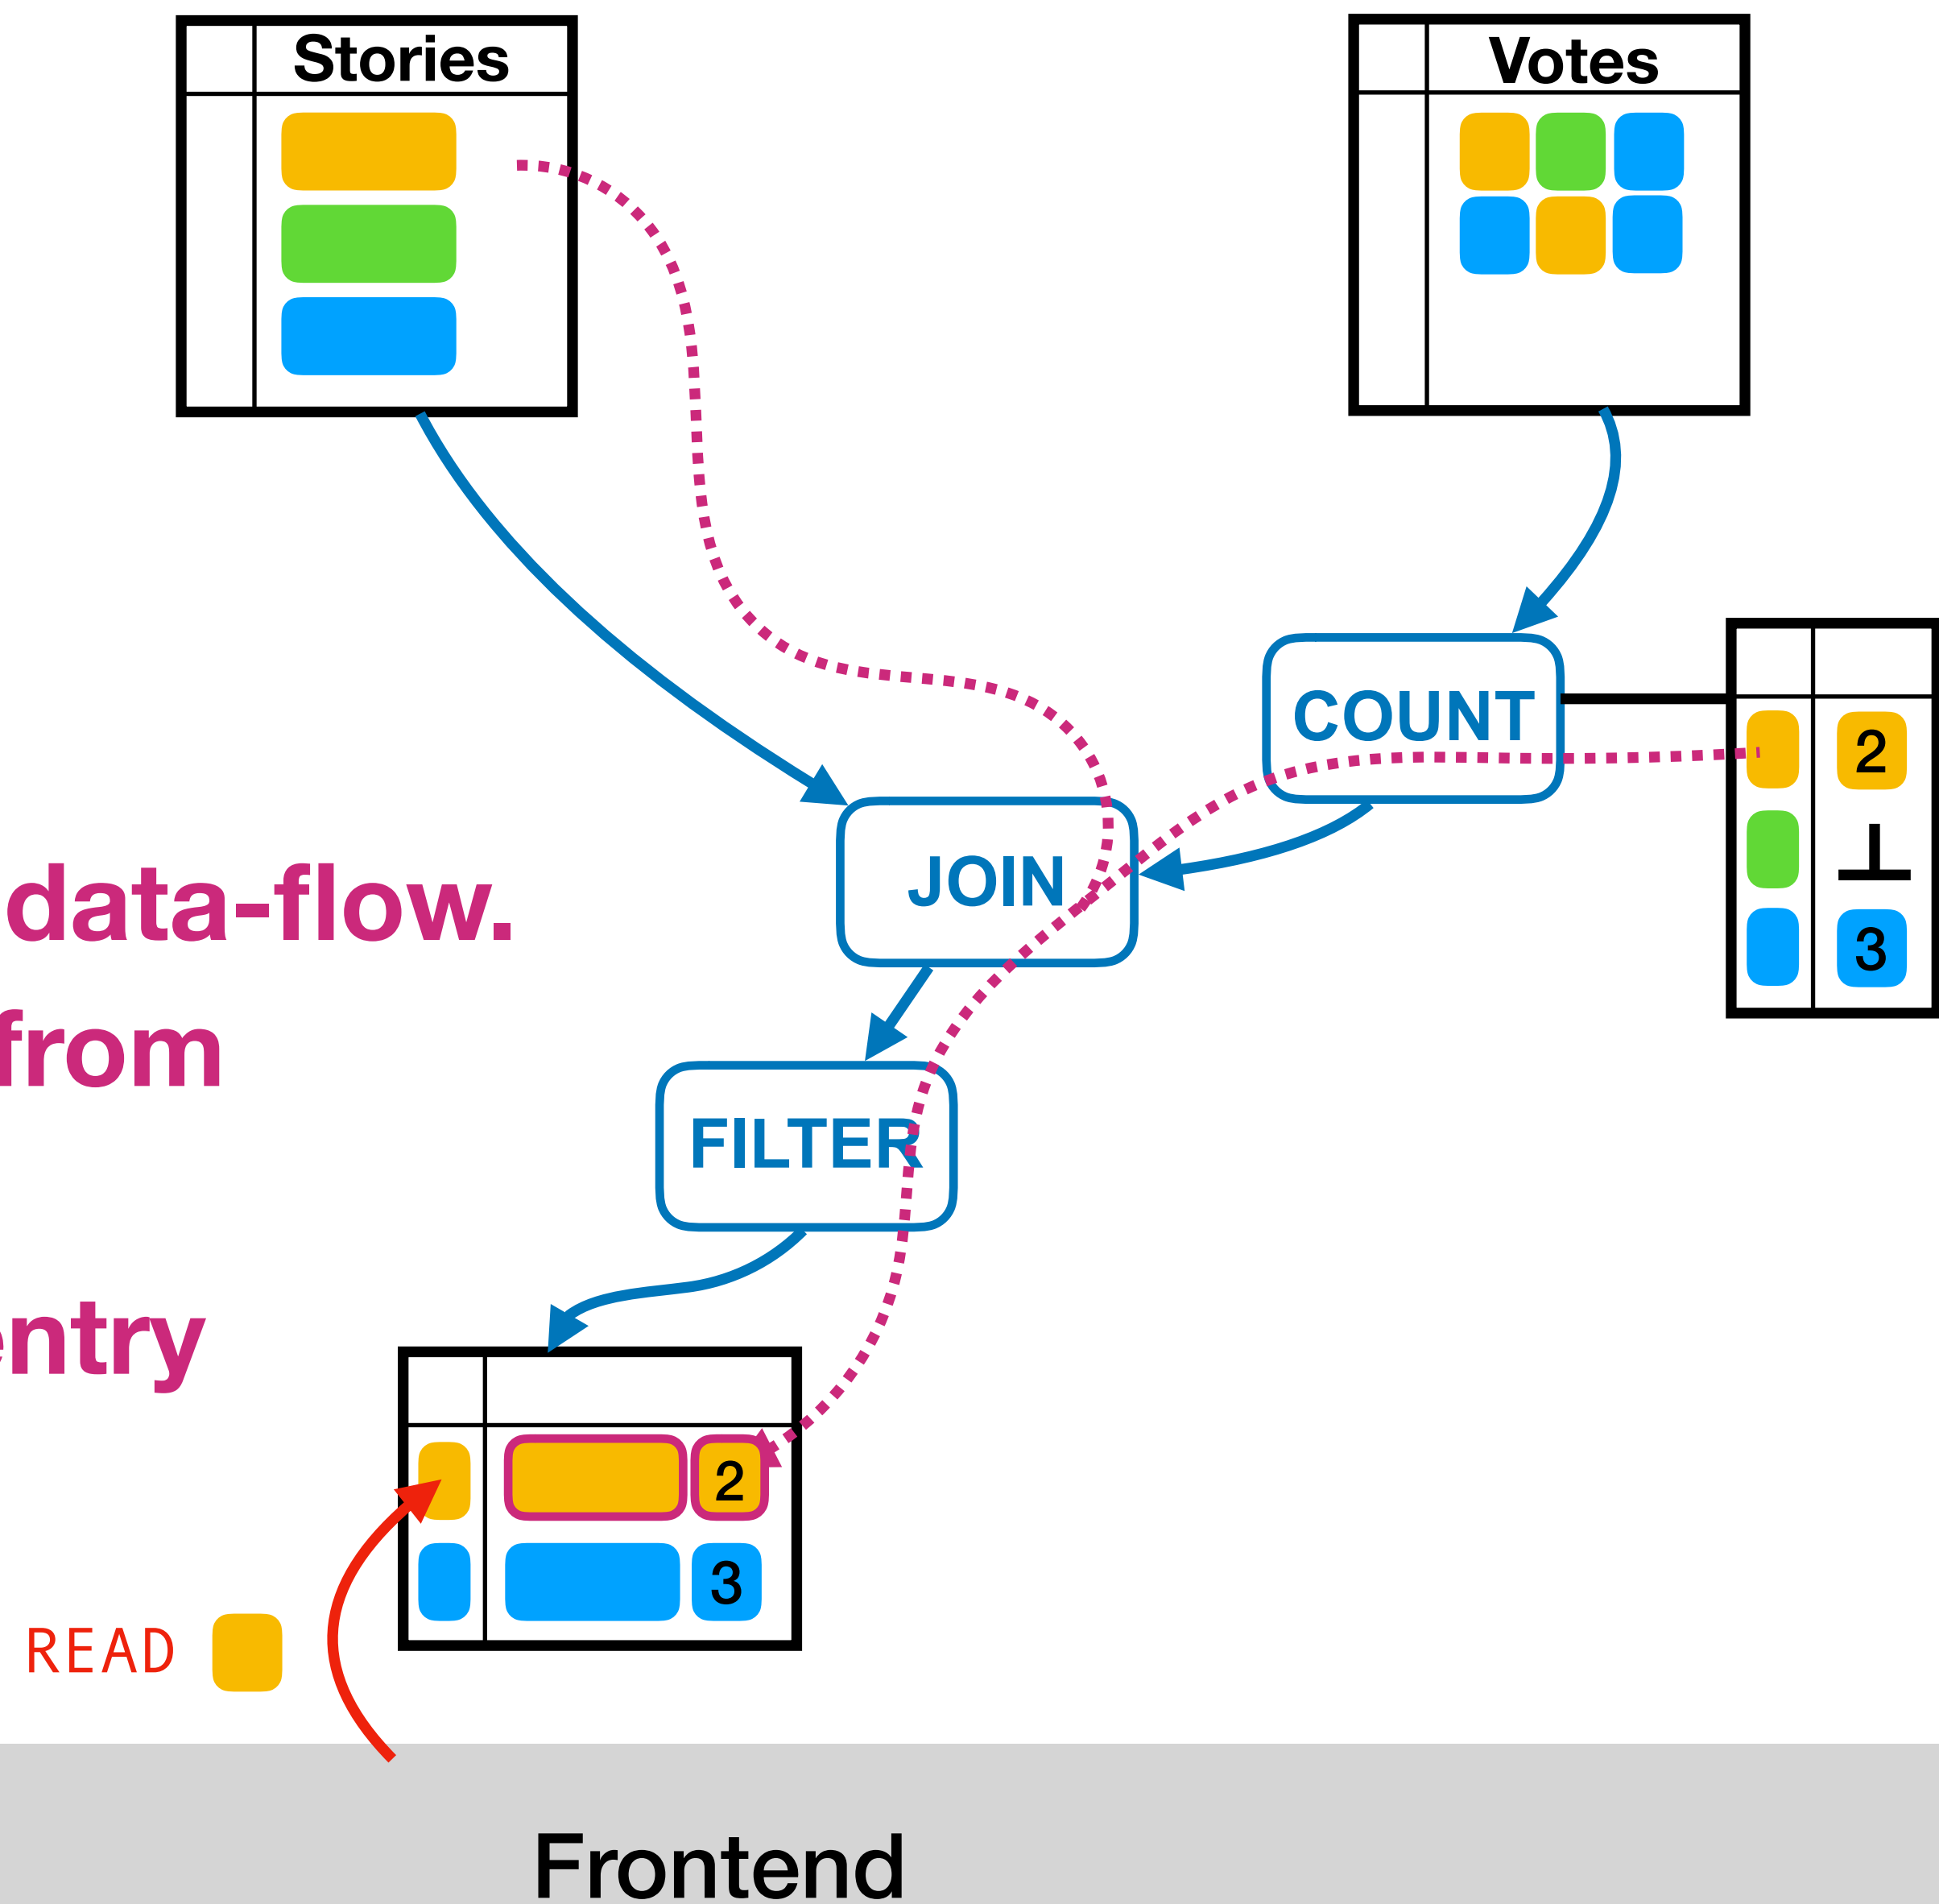
- Compute missing entry from upstream state
- Response fills missing entry



# Partially-stateful data-flow: upqueries

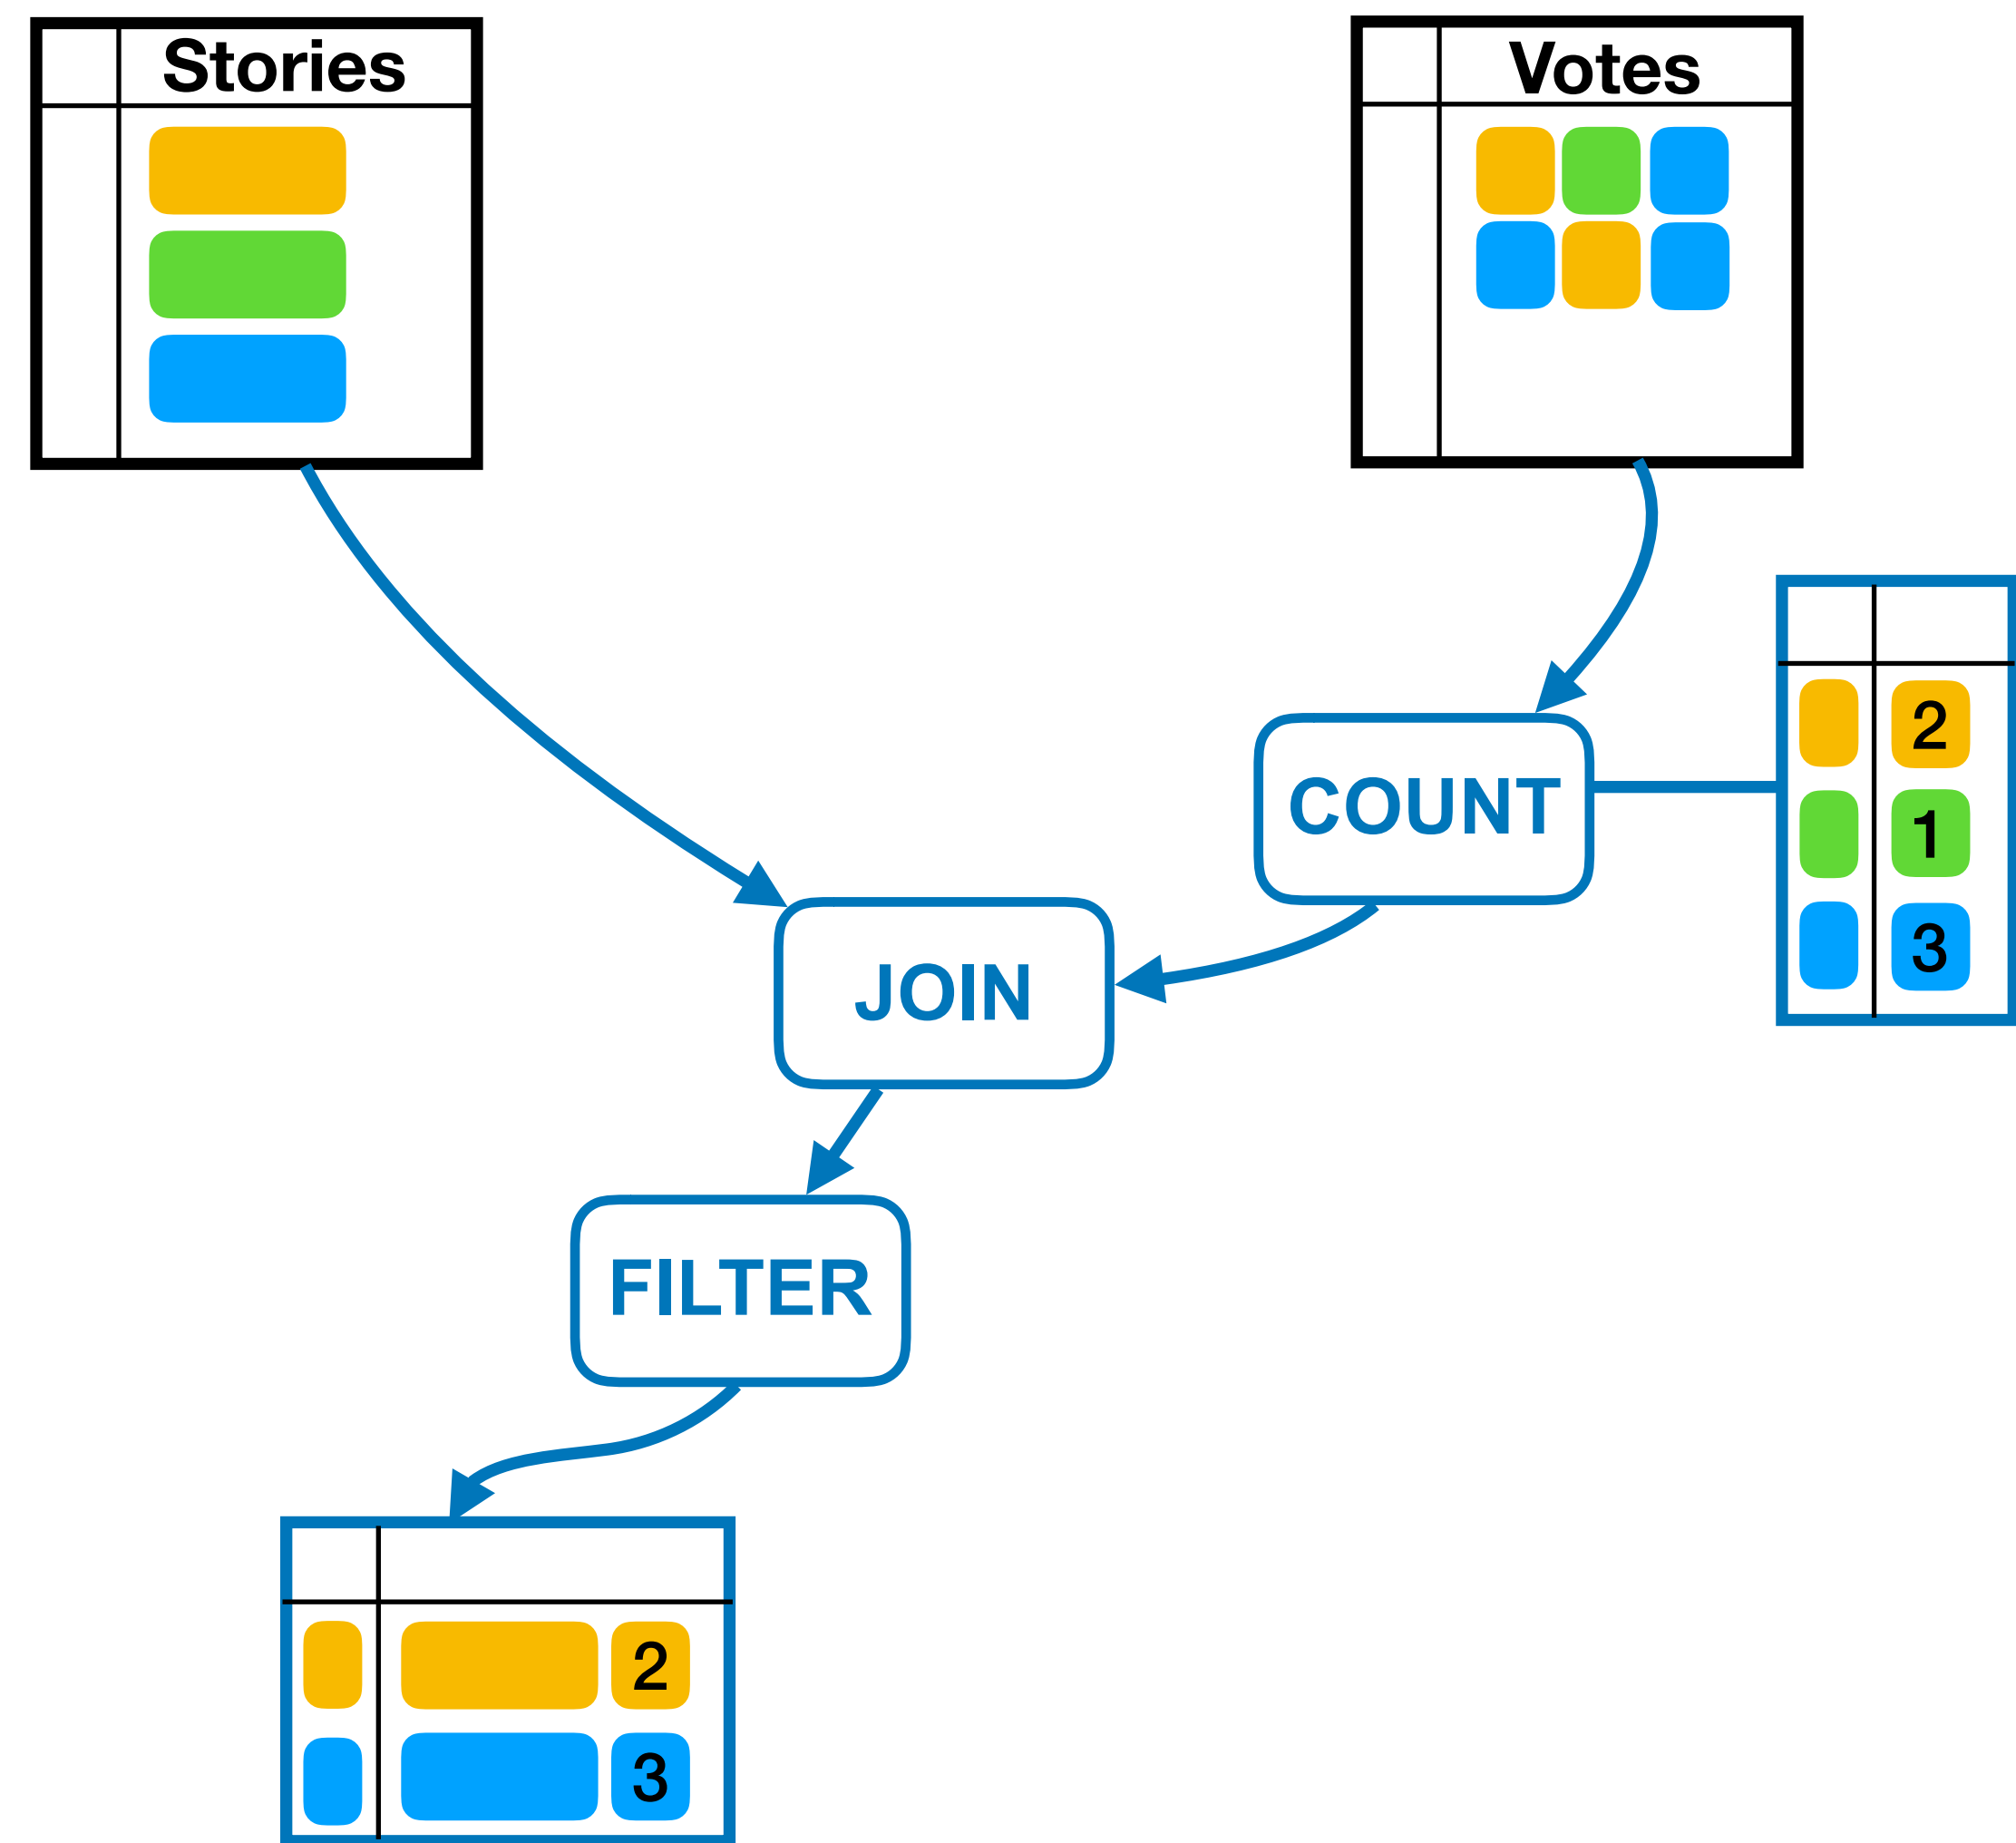
Solution: *upquery* through data-flow.

- Compute missing entry from upstream state
- Response fills missing entry



# Partial state enables live data-flow changes

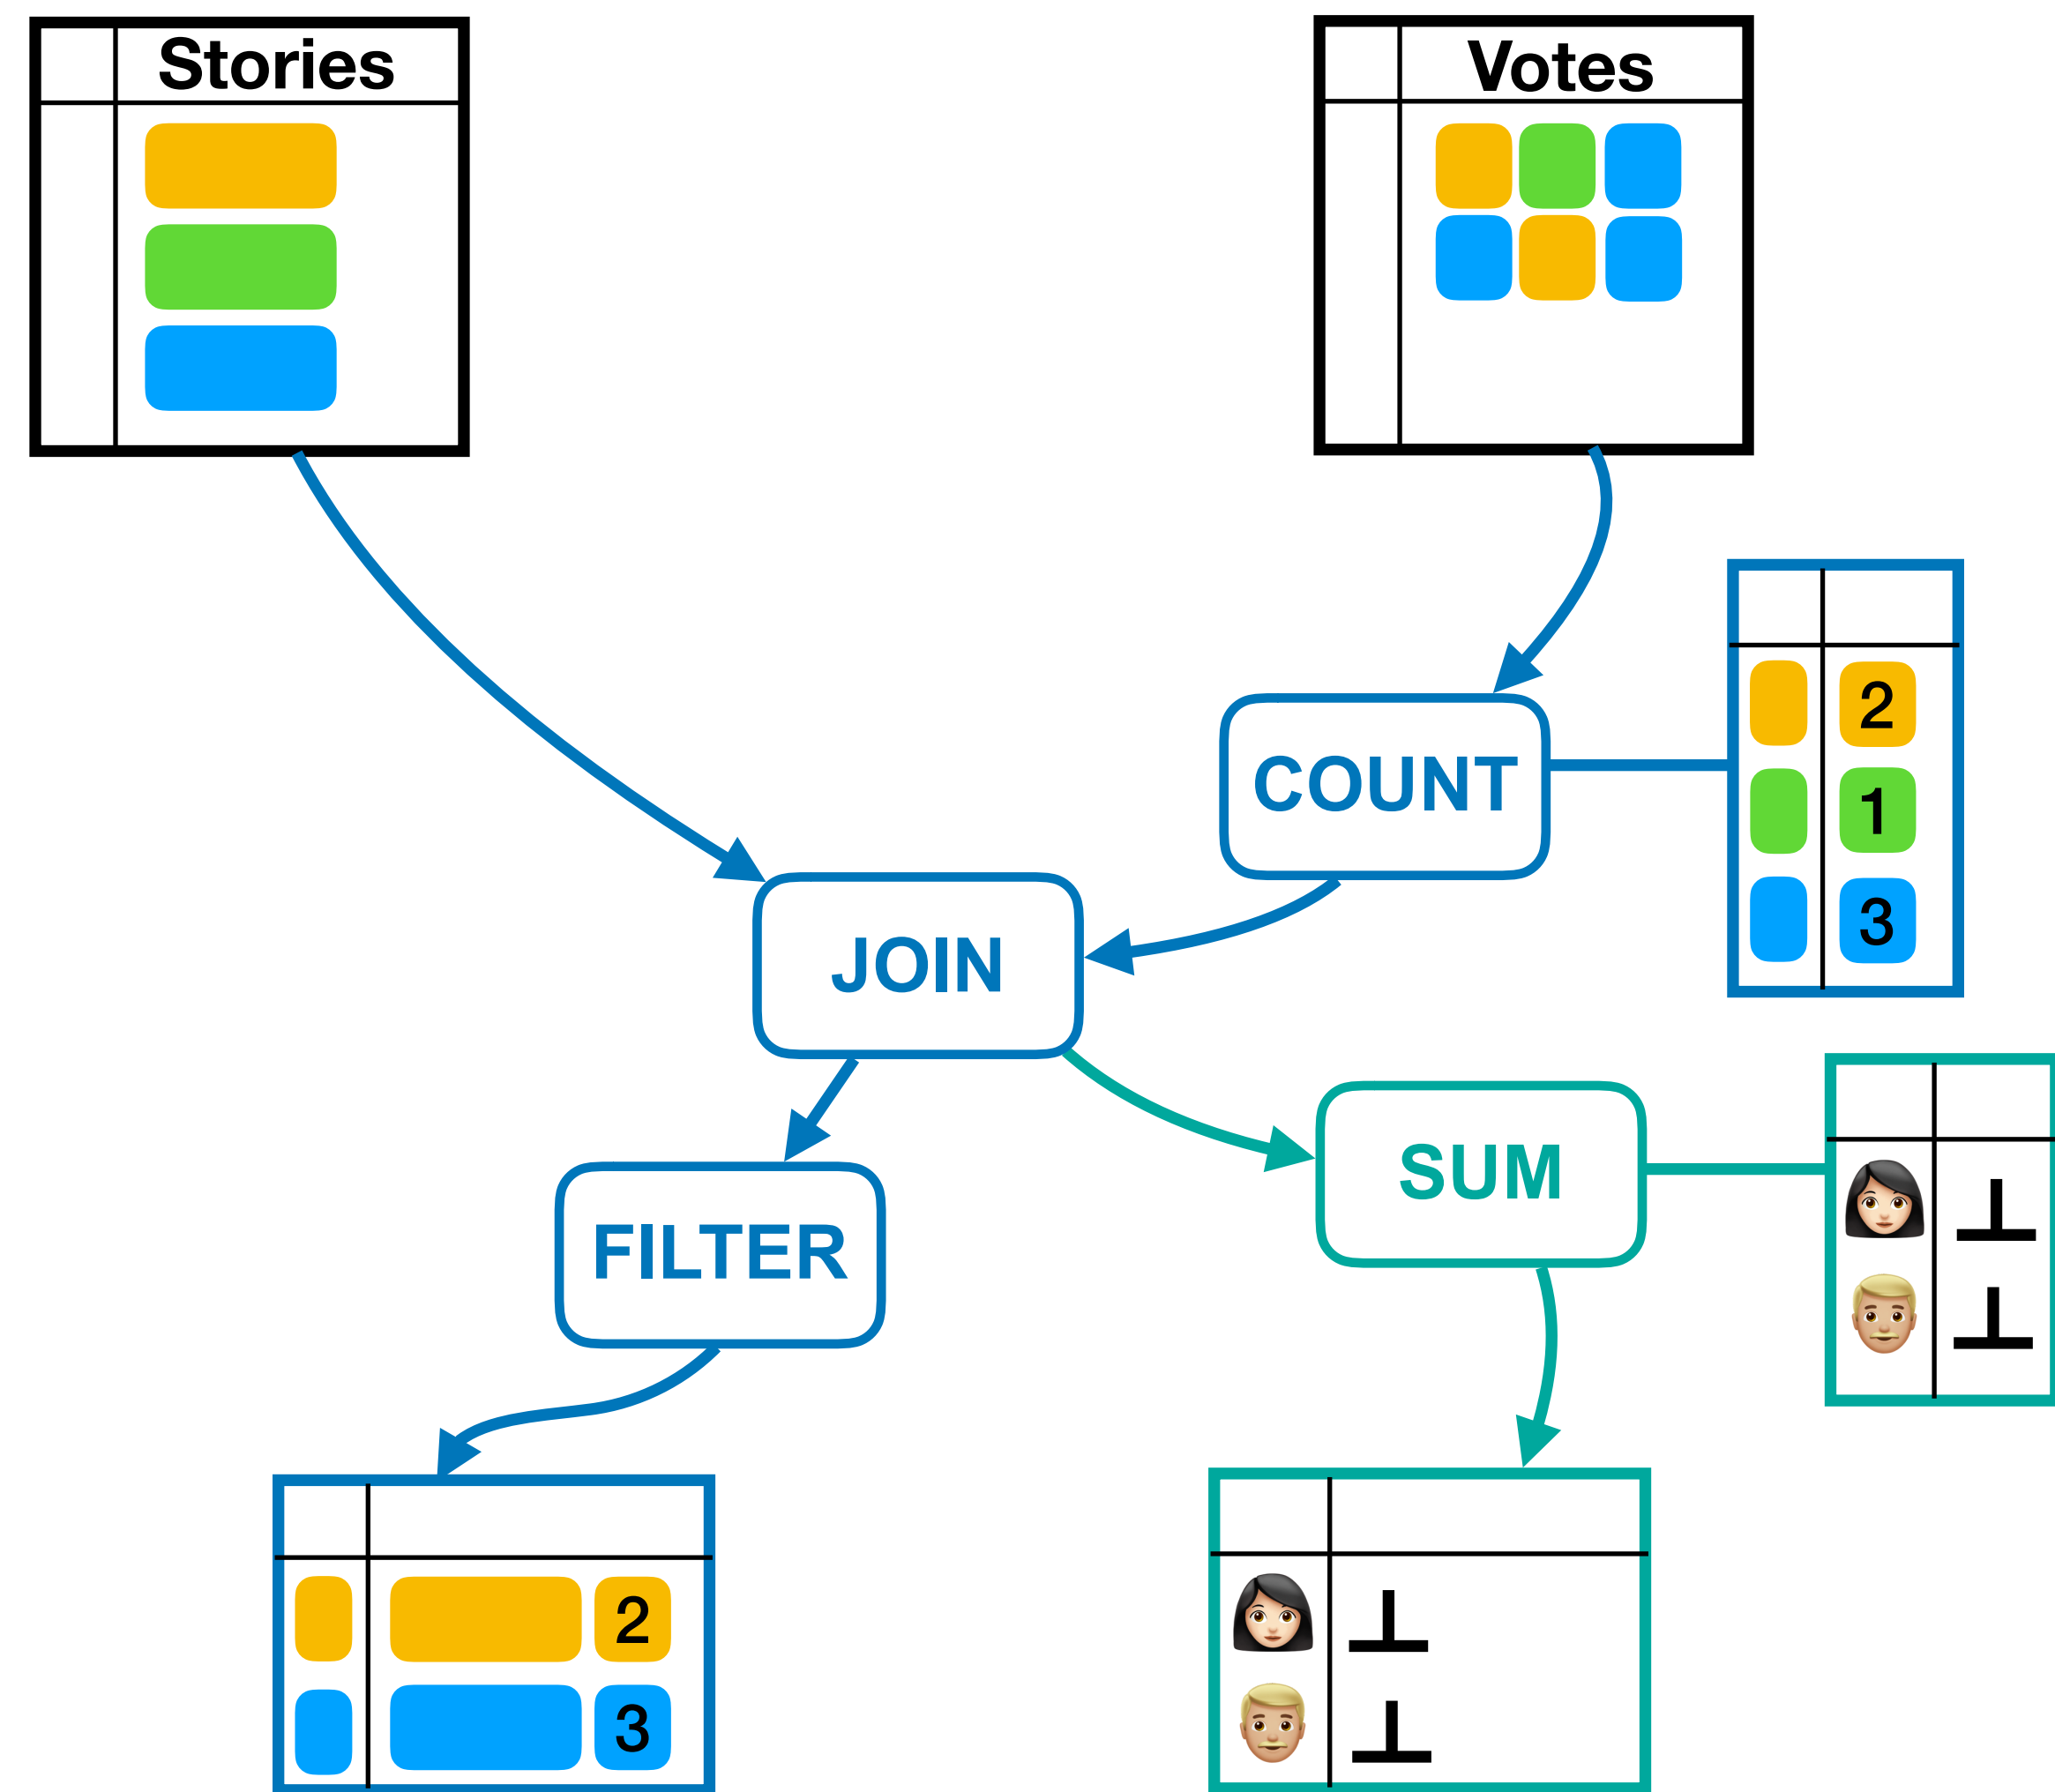
Start new views and operator state **empty**, fill via **upqueries**.



Frontend

# Partial state enables live data-flow changes

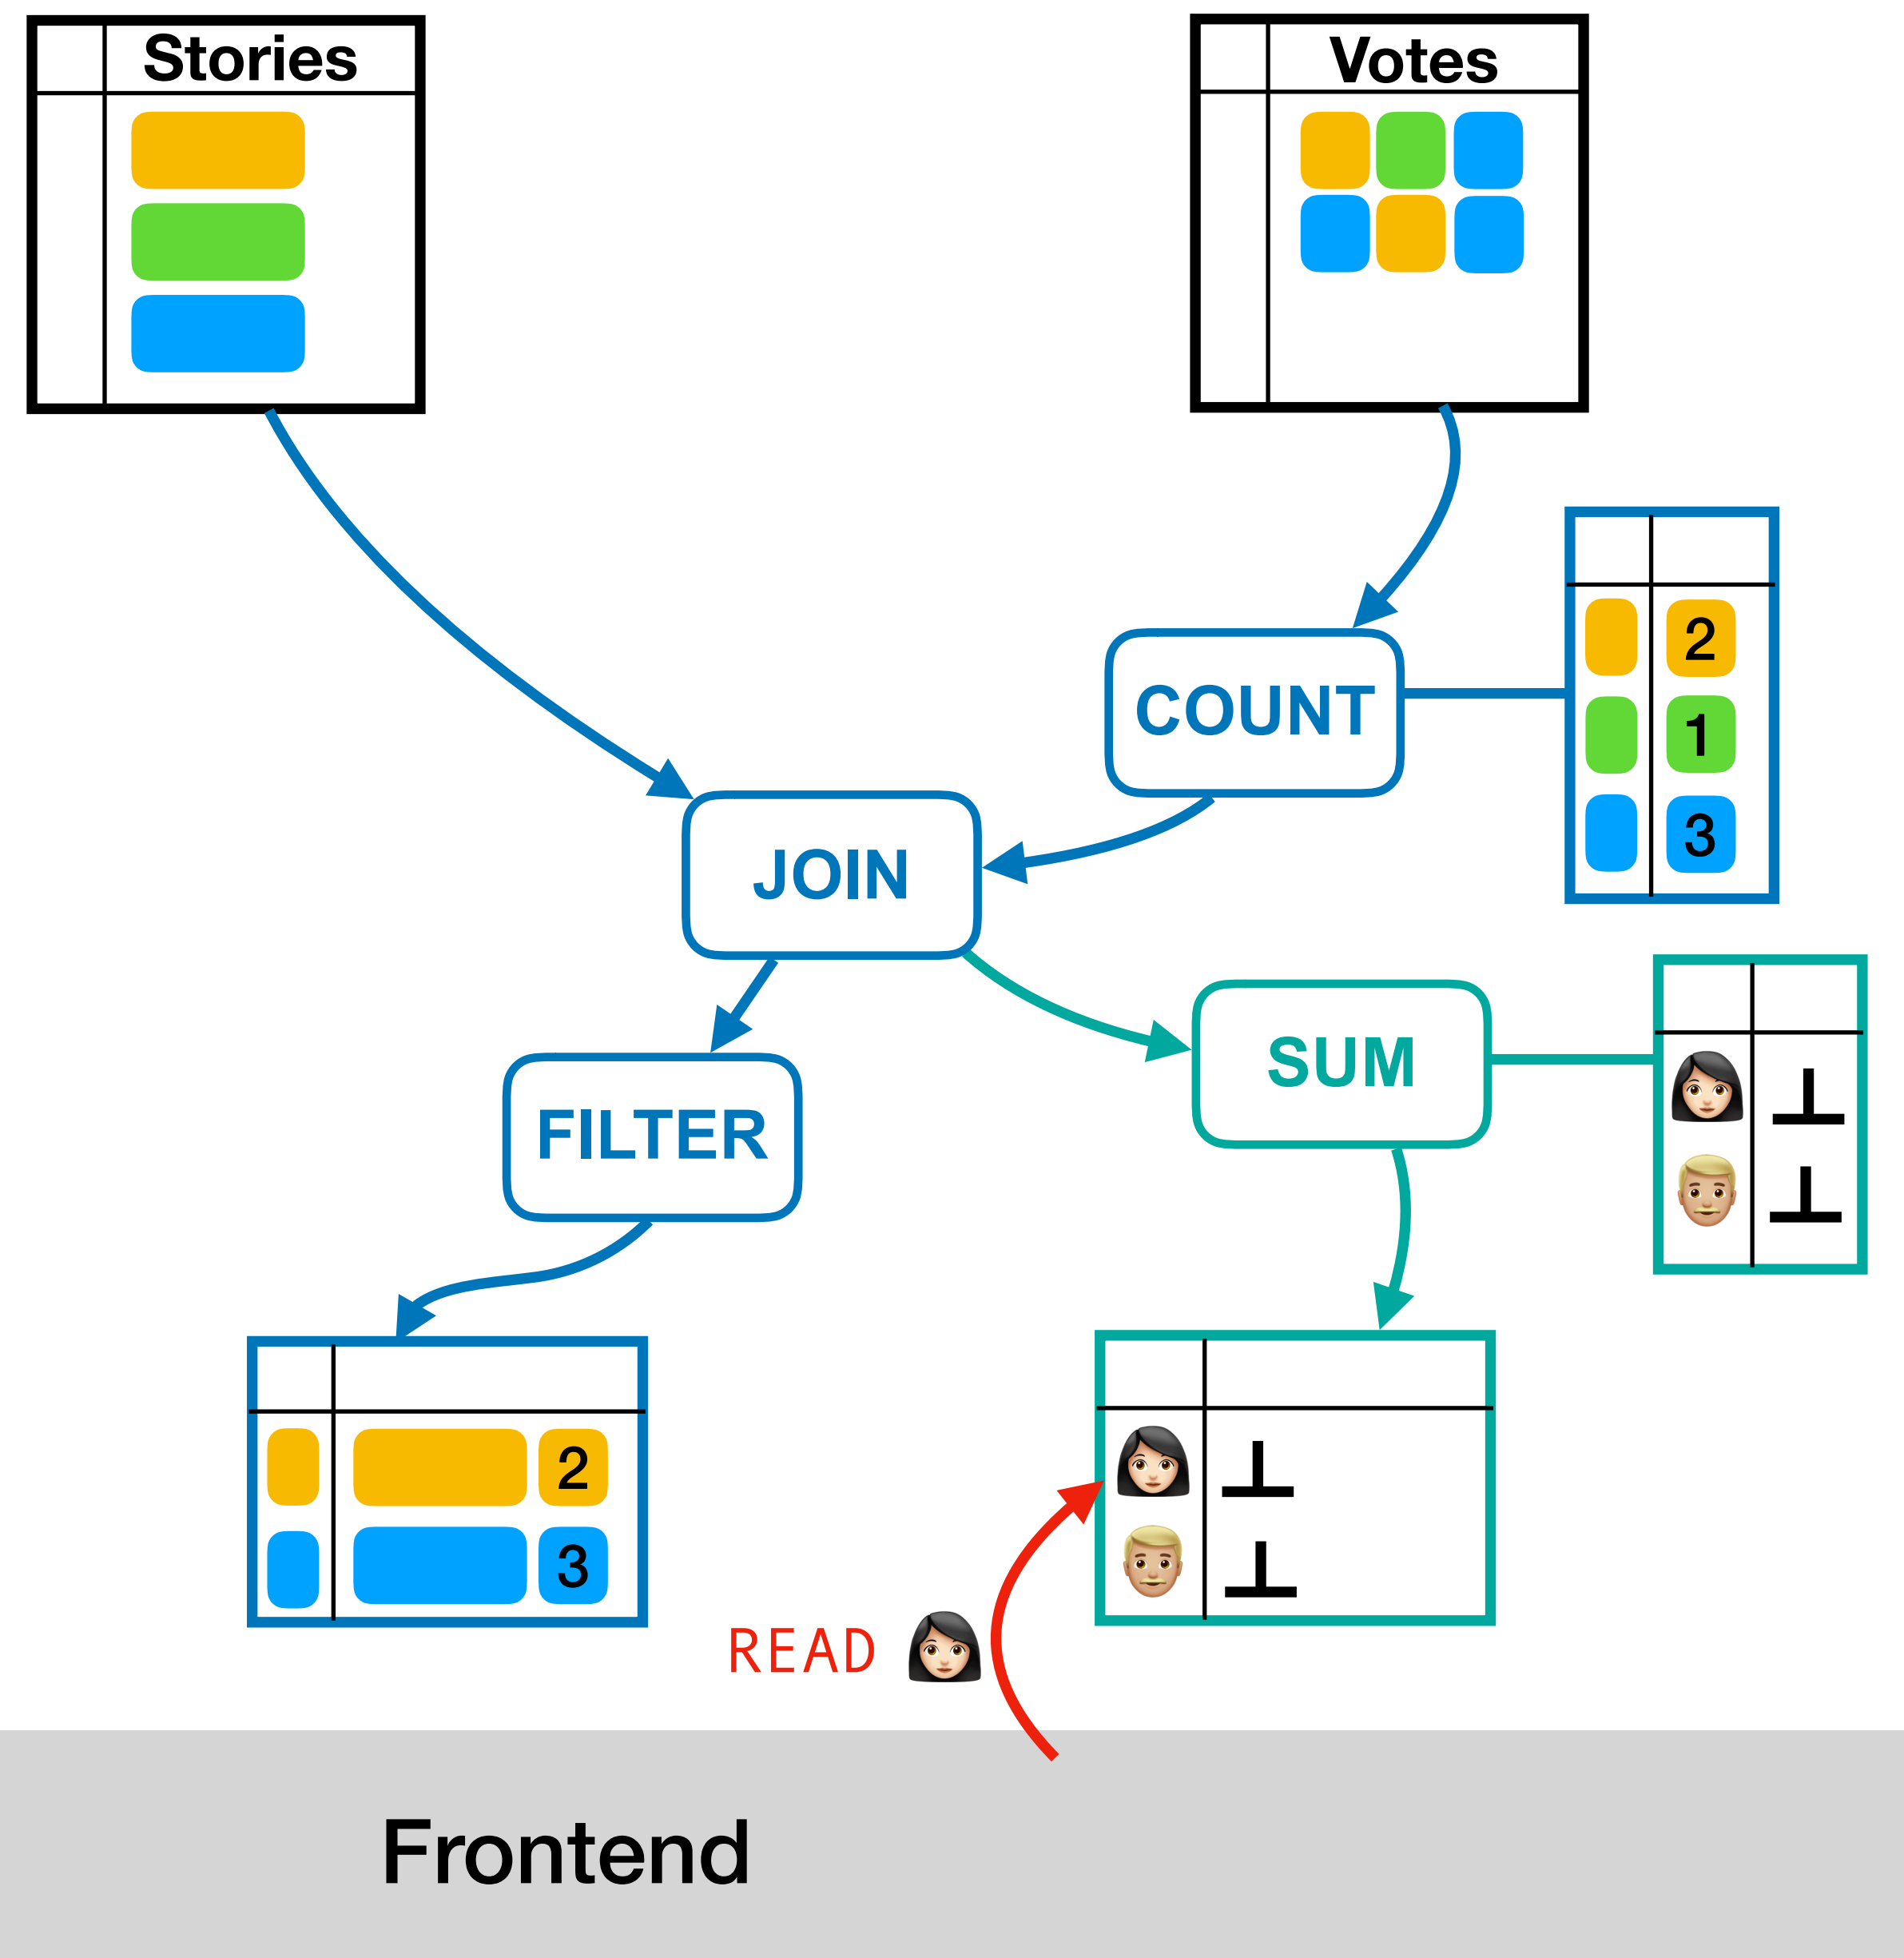
Start new views and operator state **empty**, fill via **upqueries**.



Frontend

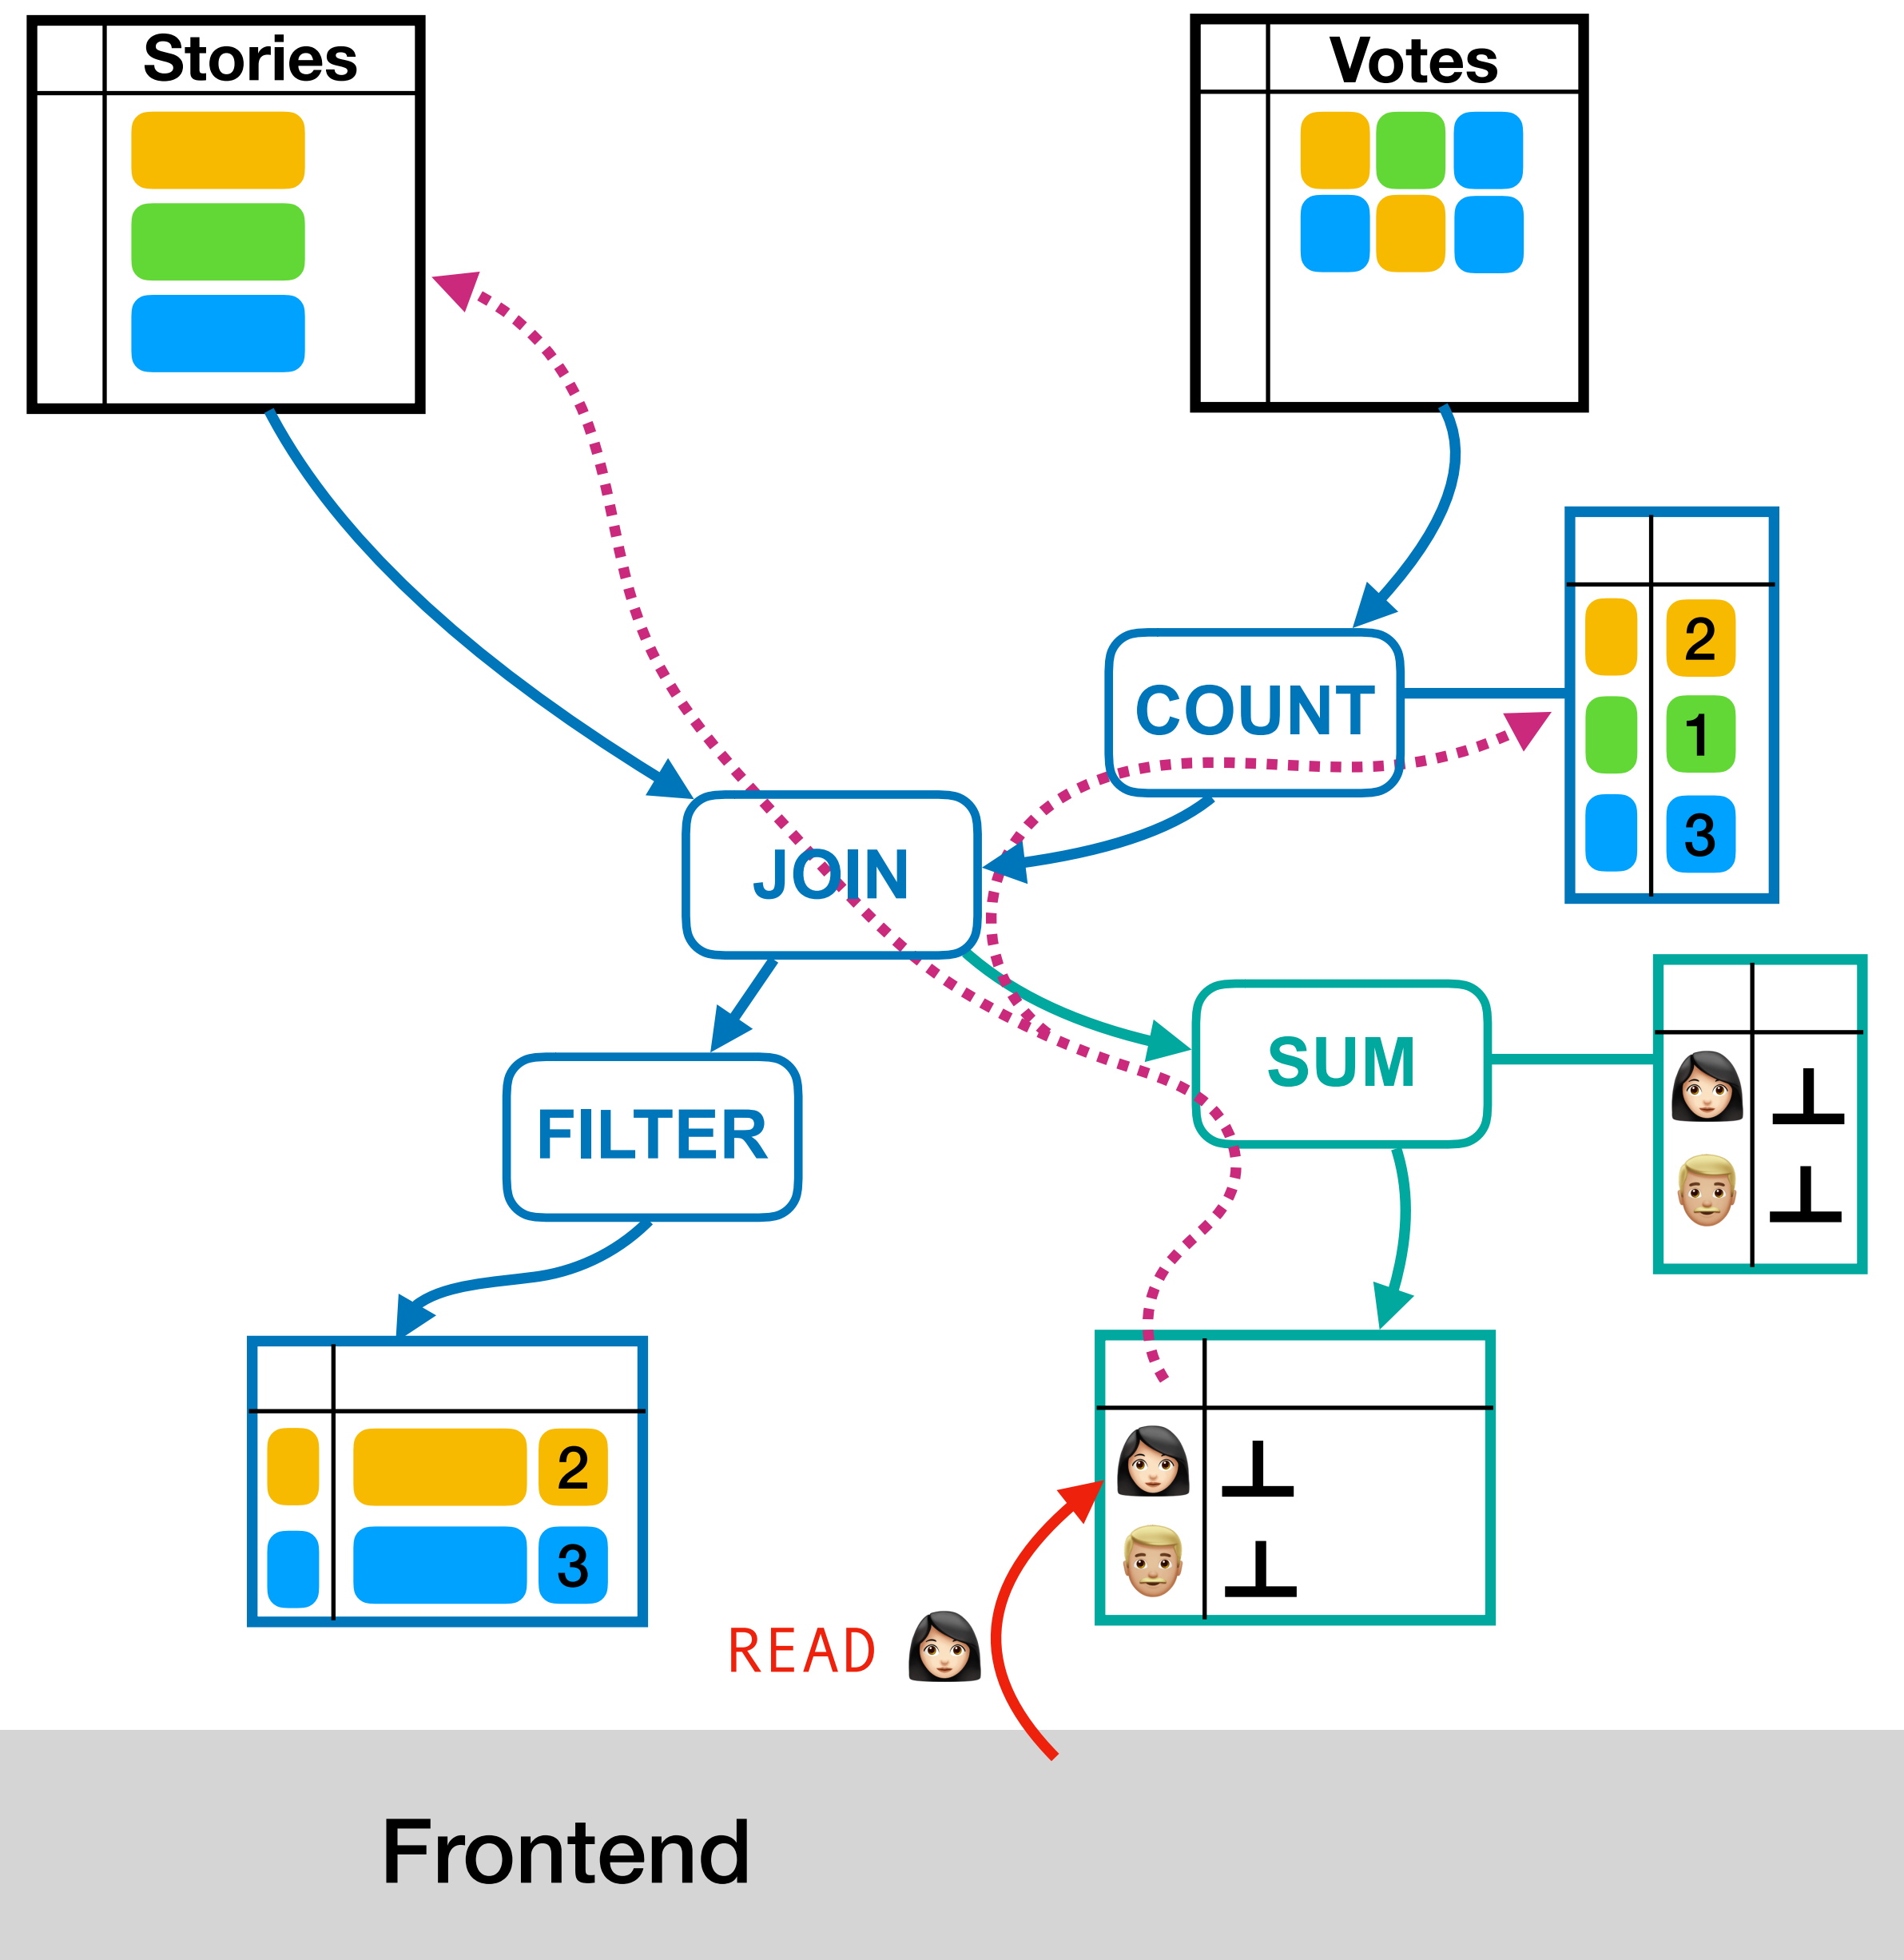
# Partial state enables live data-flow changes

Start new views and operator state **empty**, fill via **upqueries**.



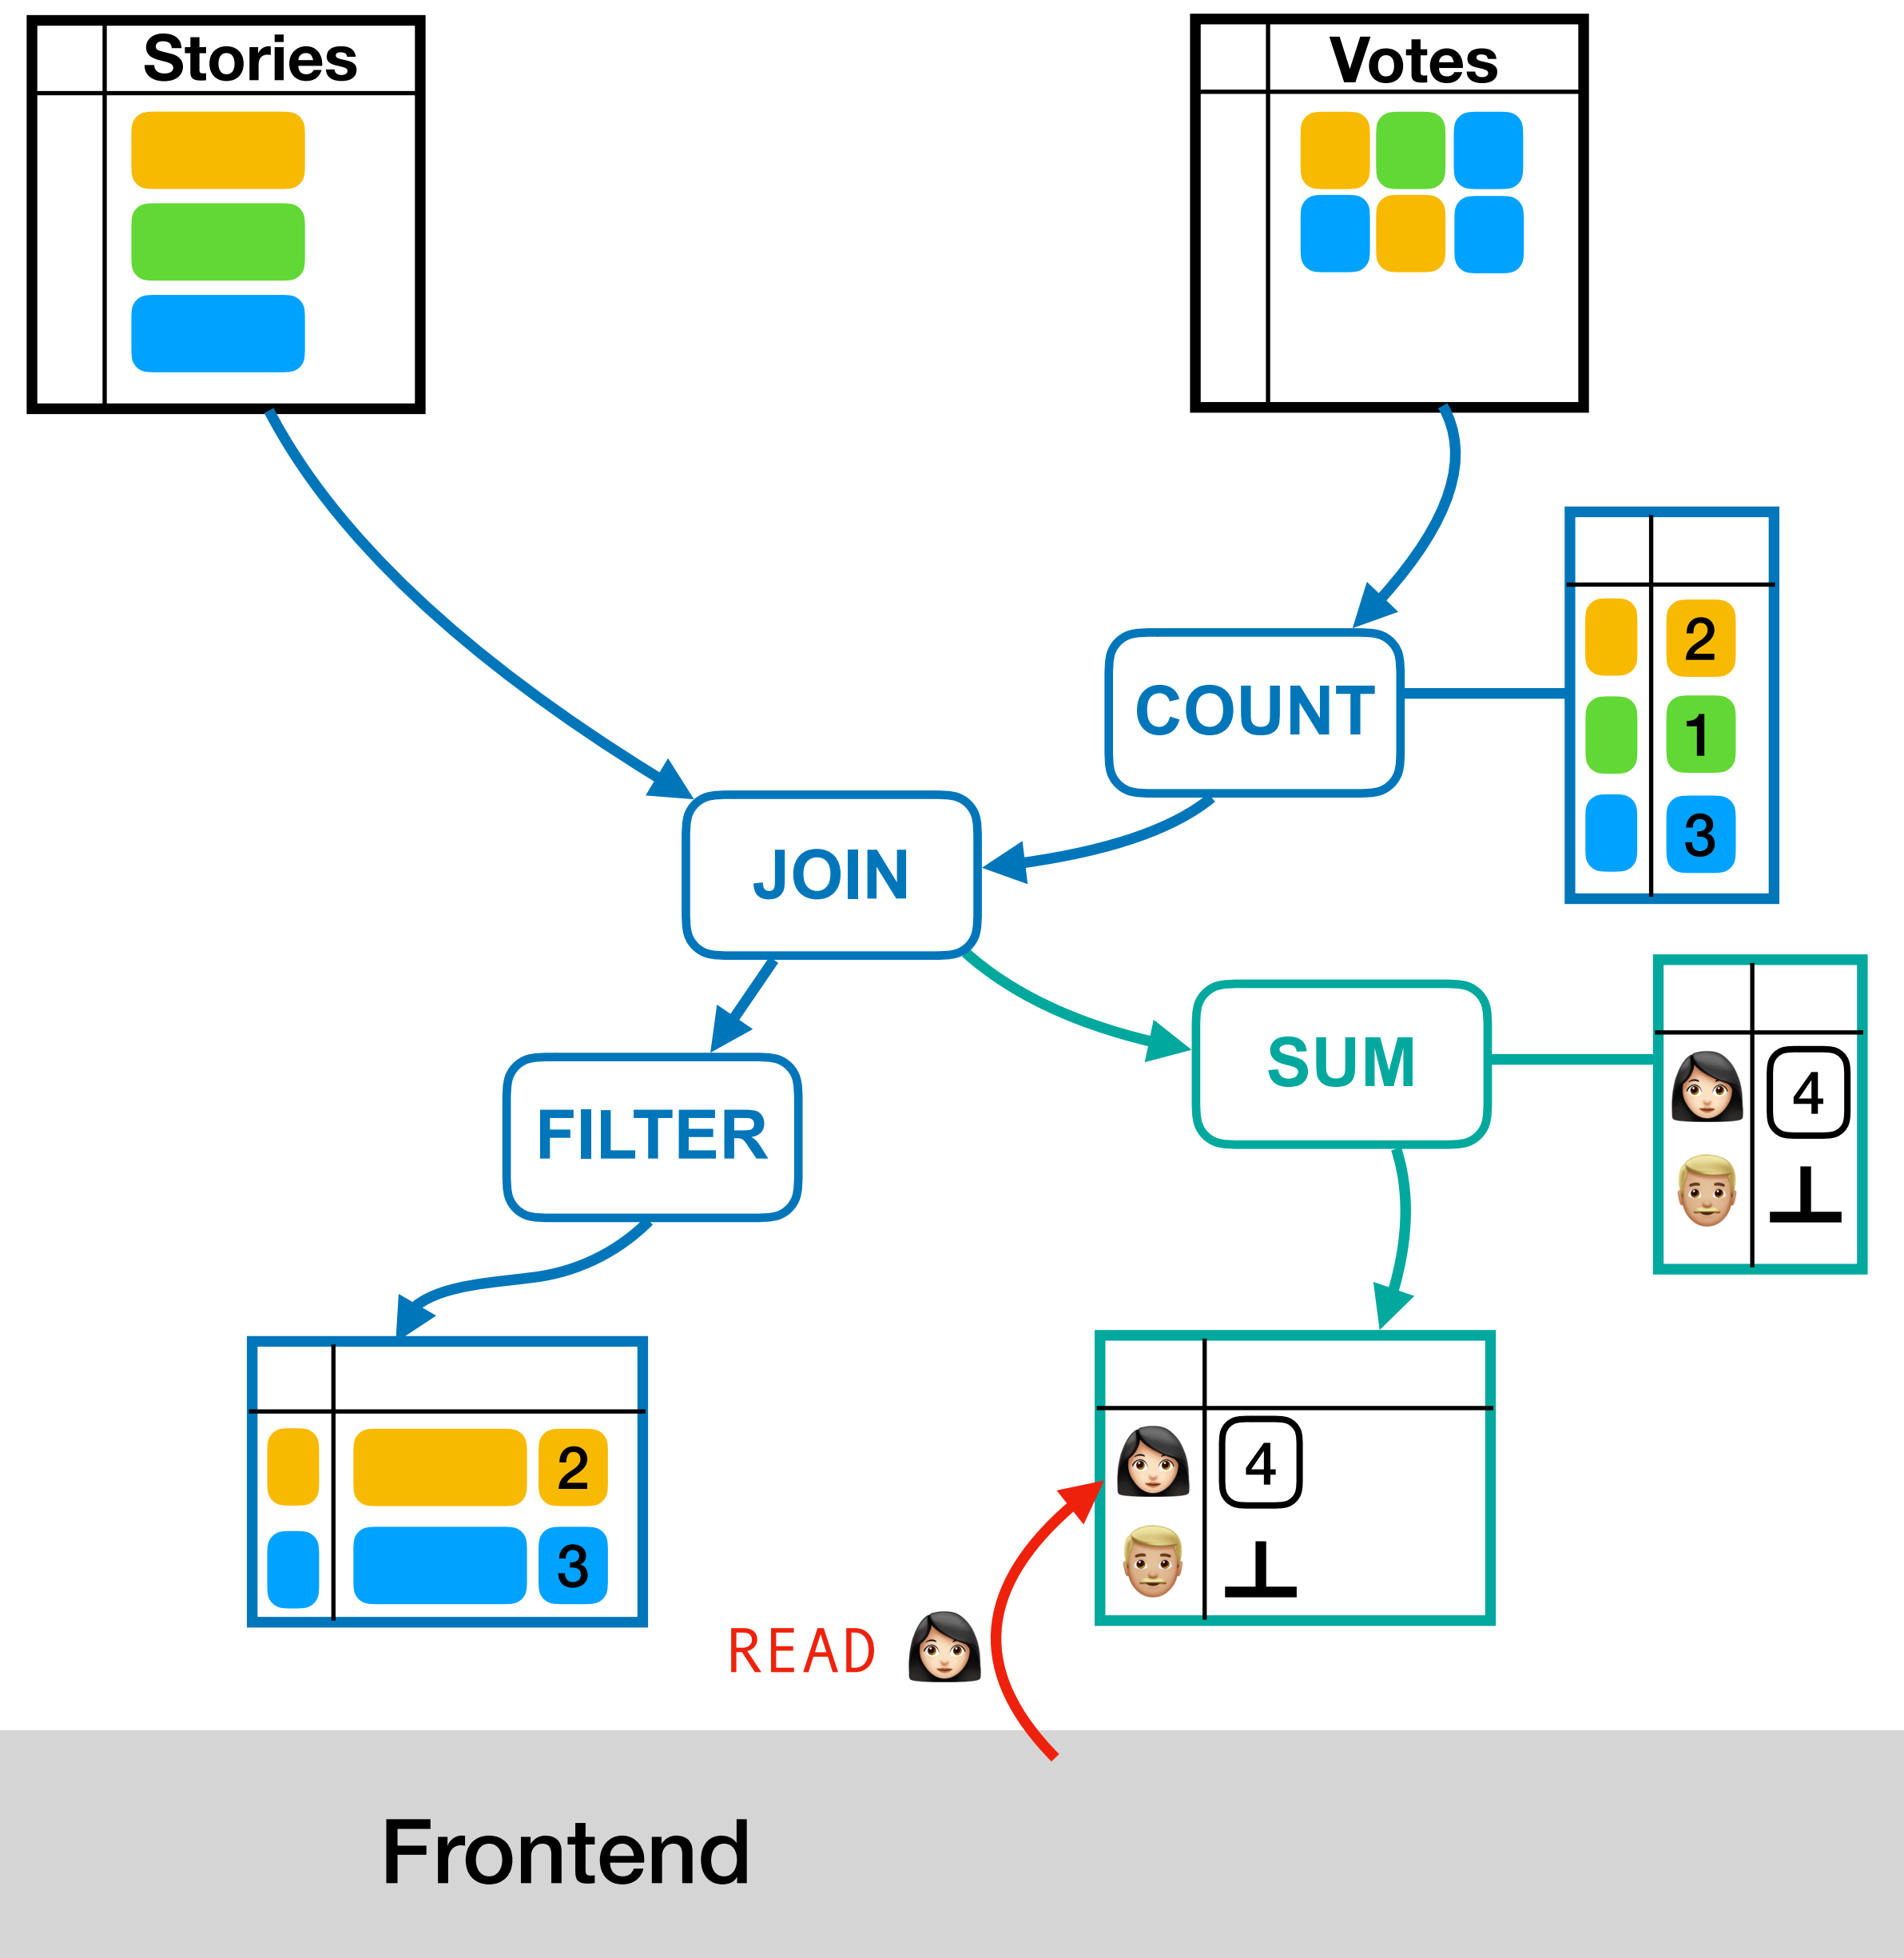
# Partial state enables live data-flow changes

Start new views and operator state **empty**, fill via **upqueries**.



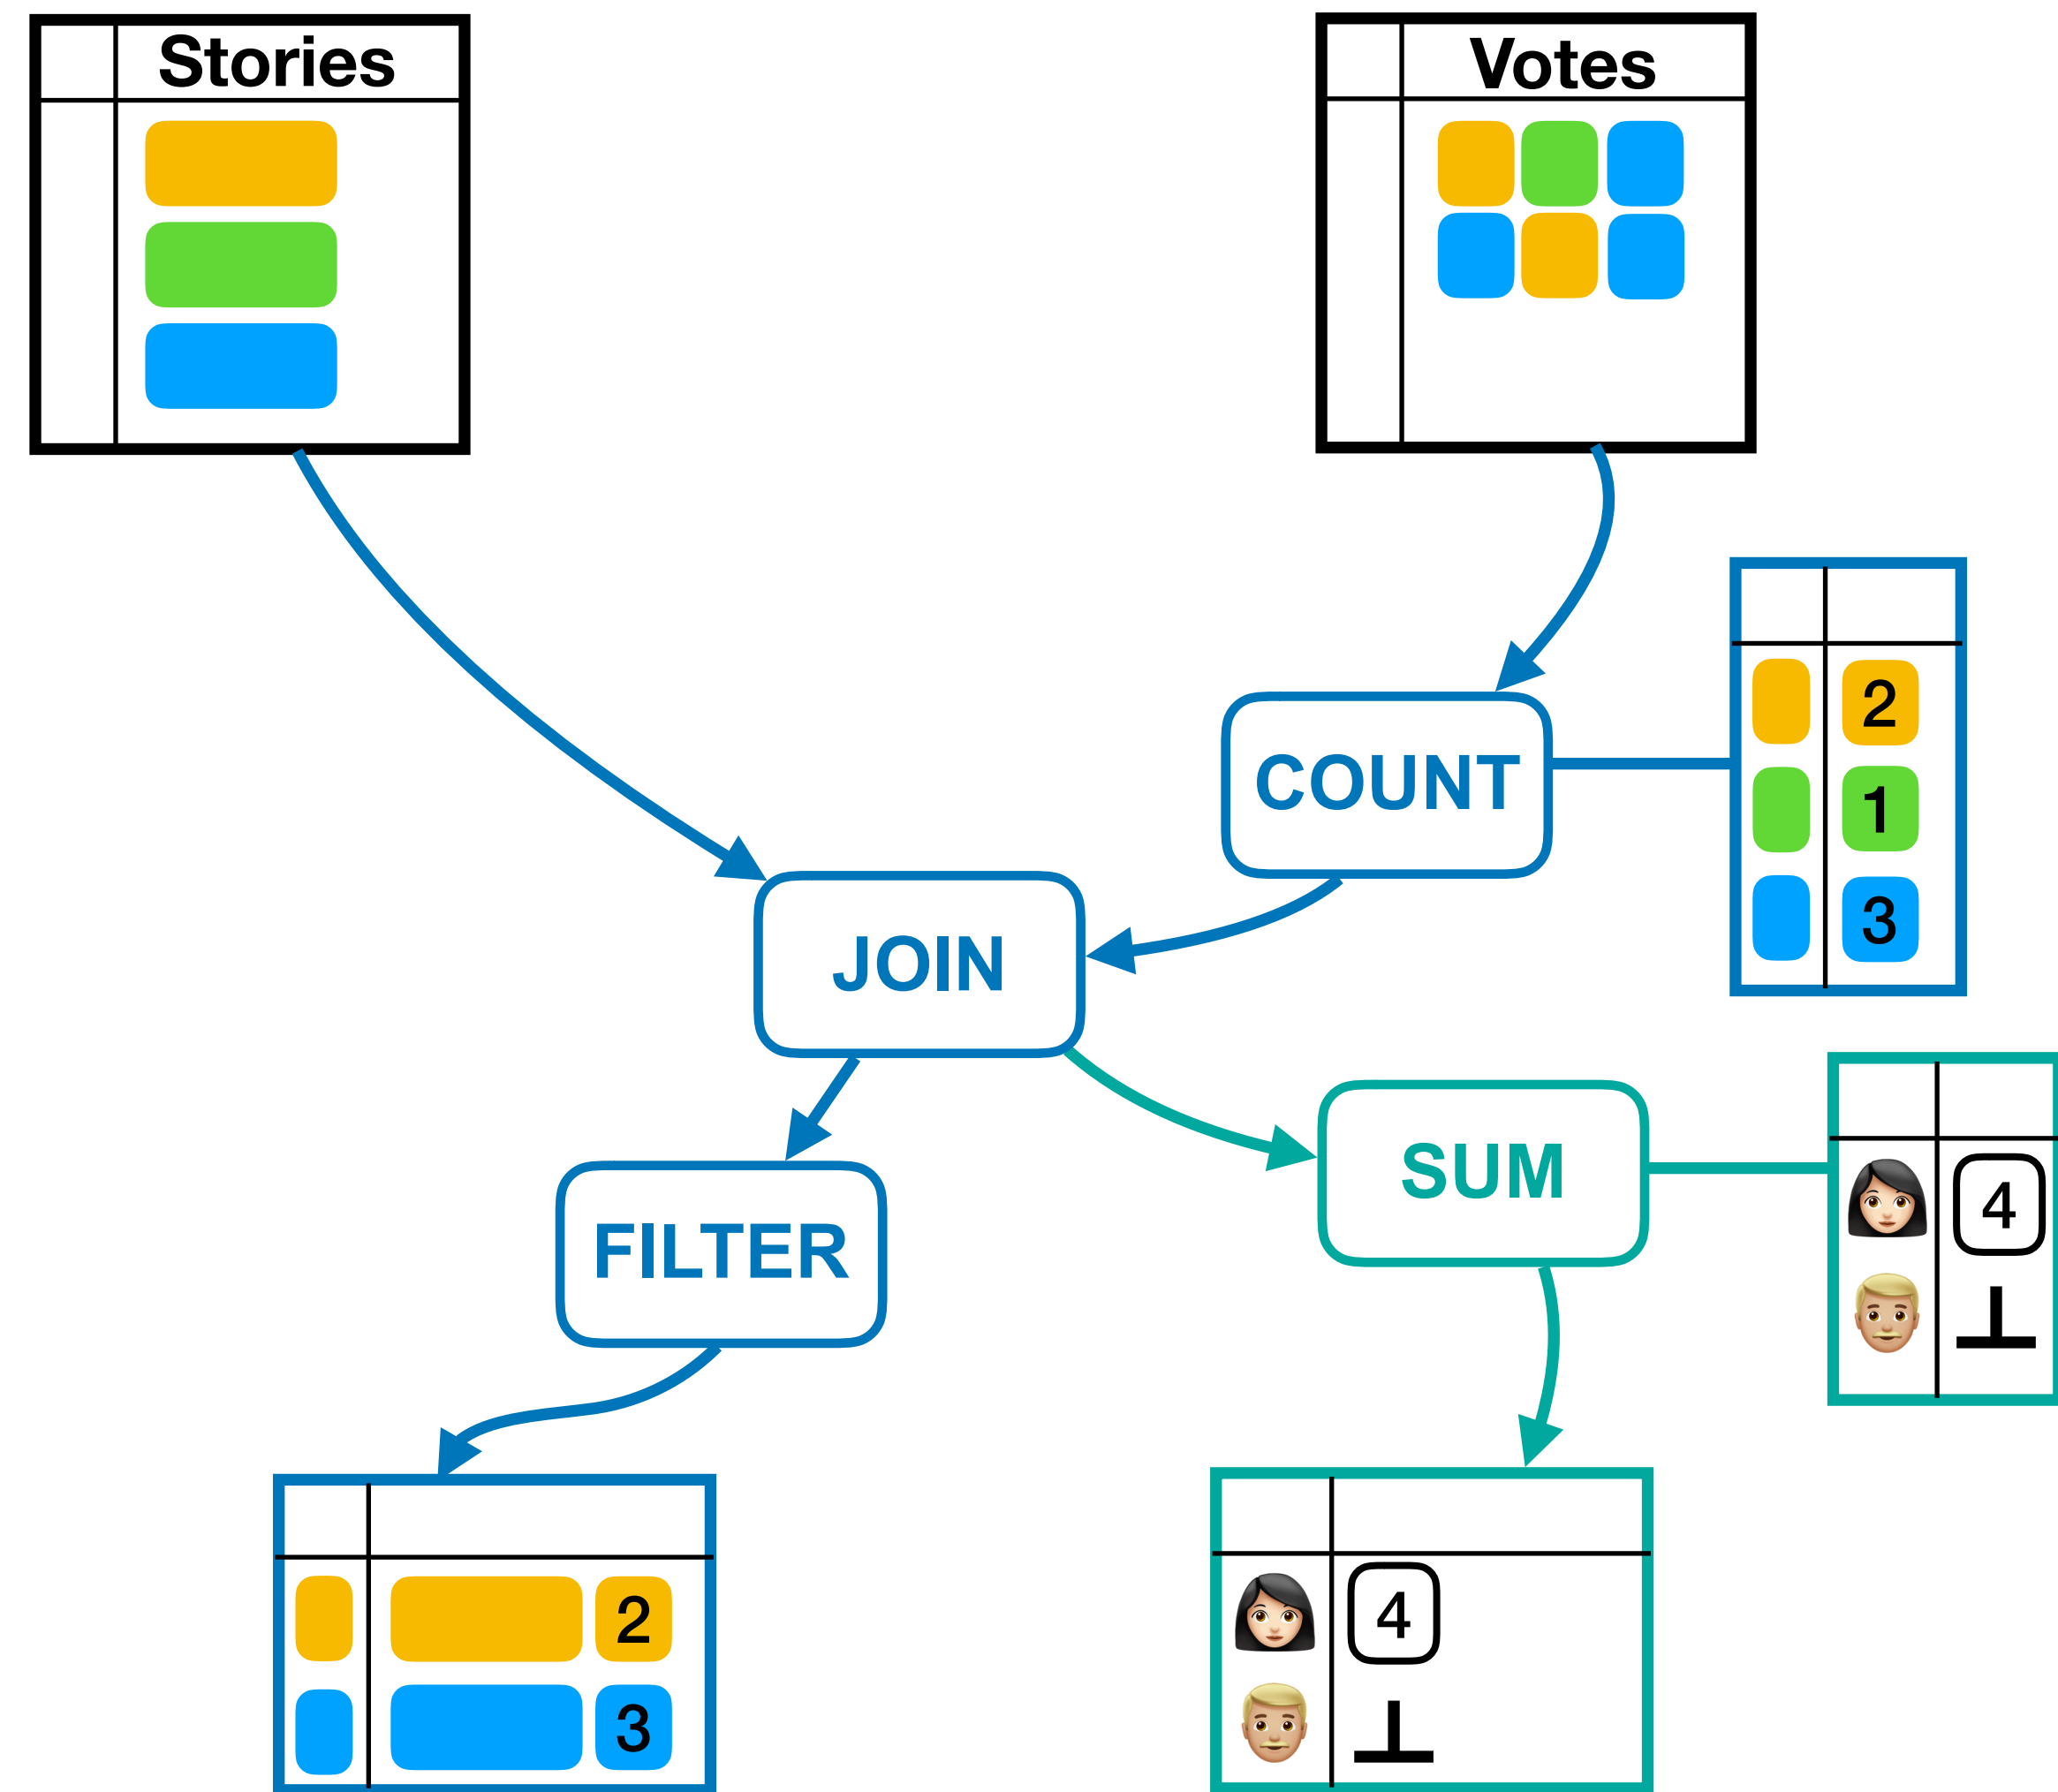
# Partial state enables live data-flow changes

Start new views and operator state **empty**, fill via **upqueries**.





# High performance requires concurrency

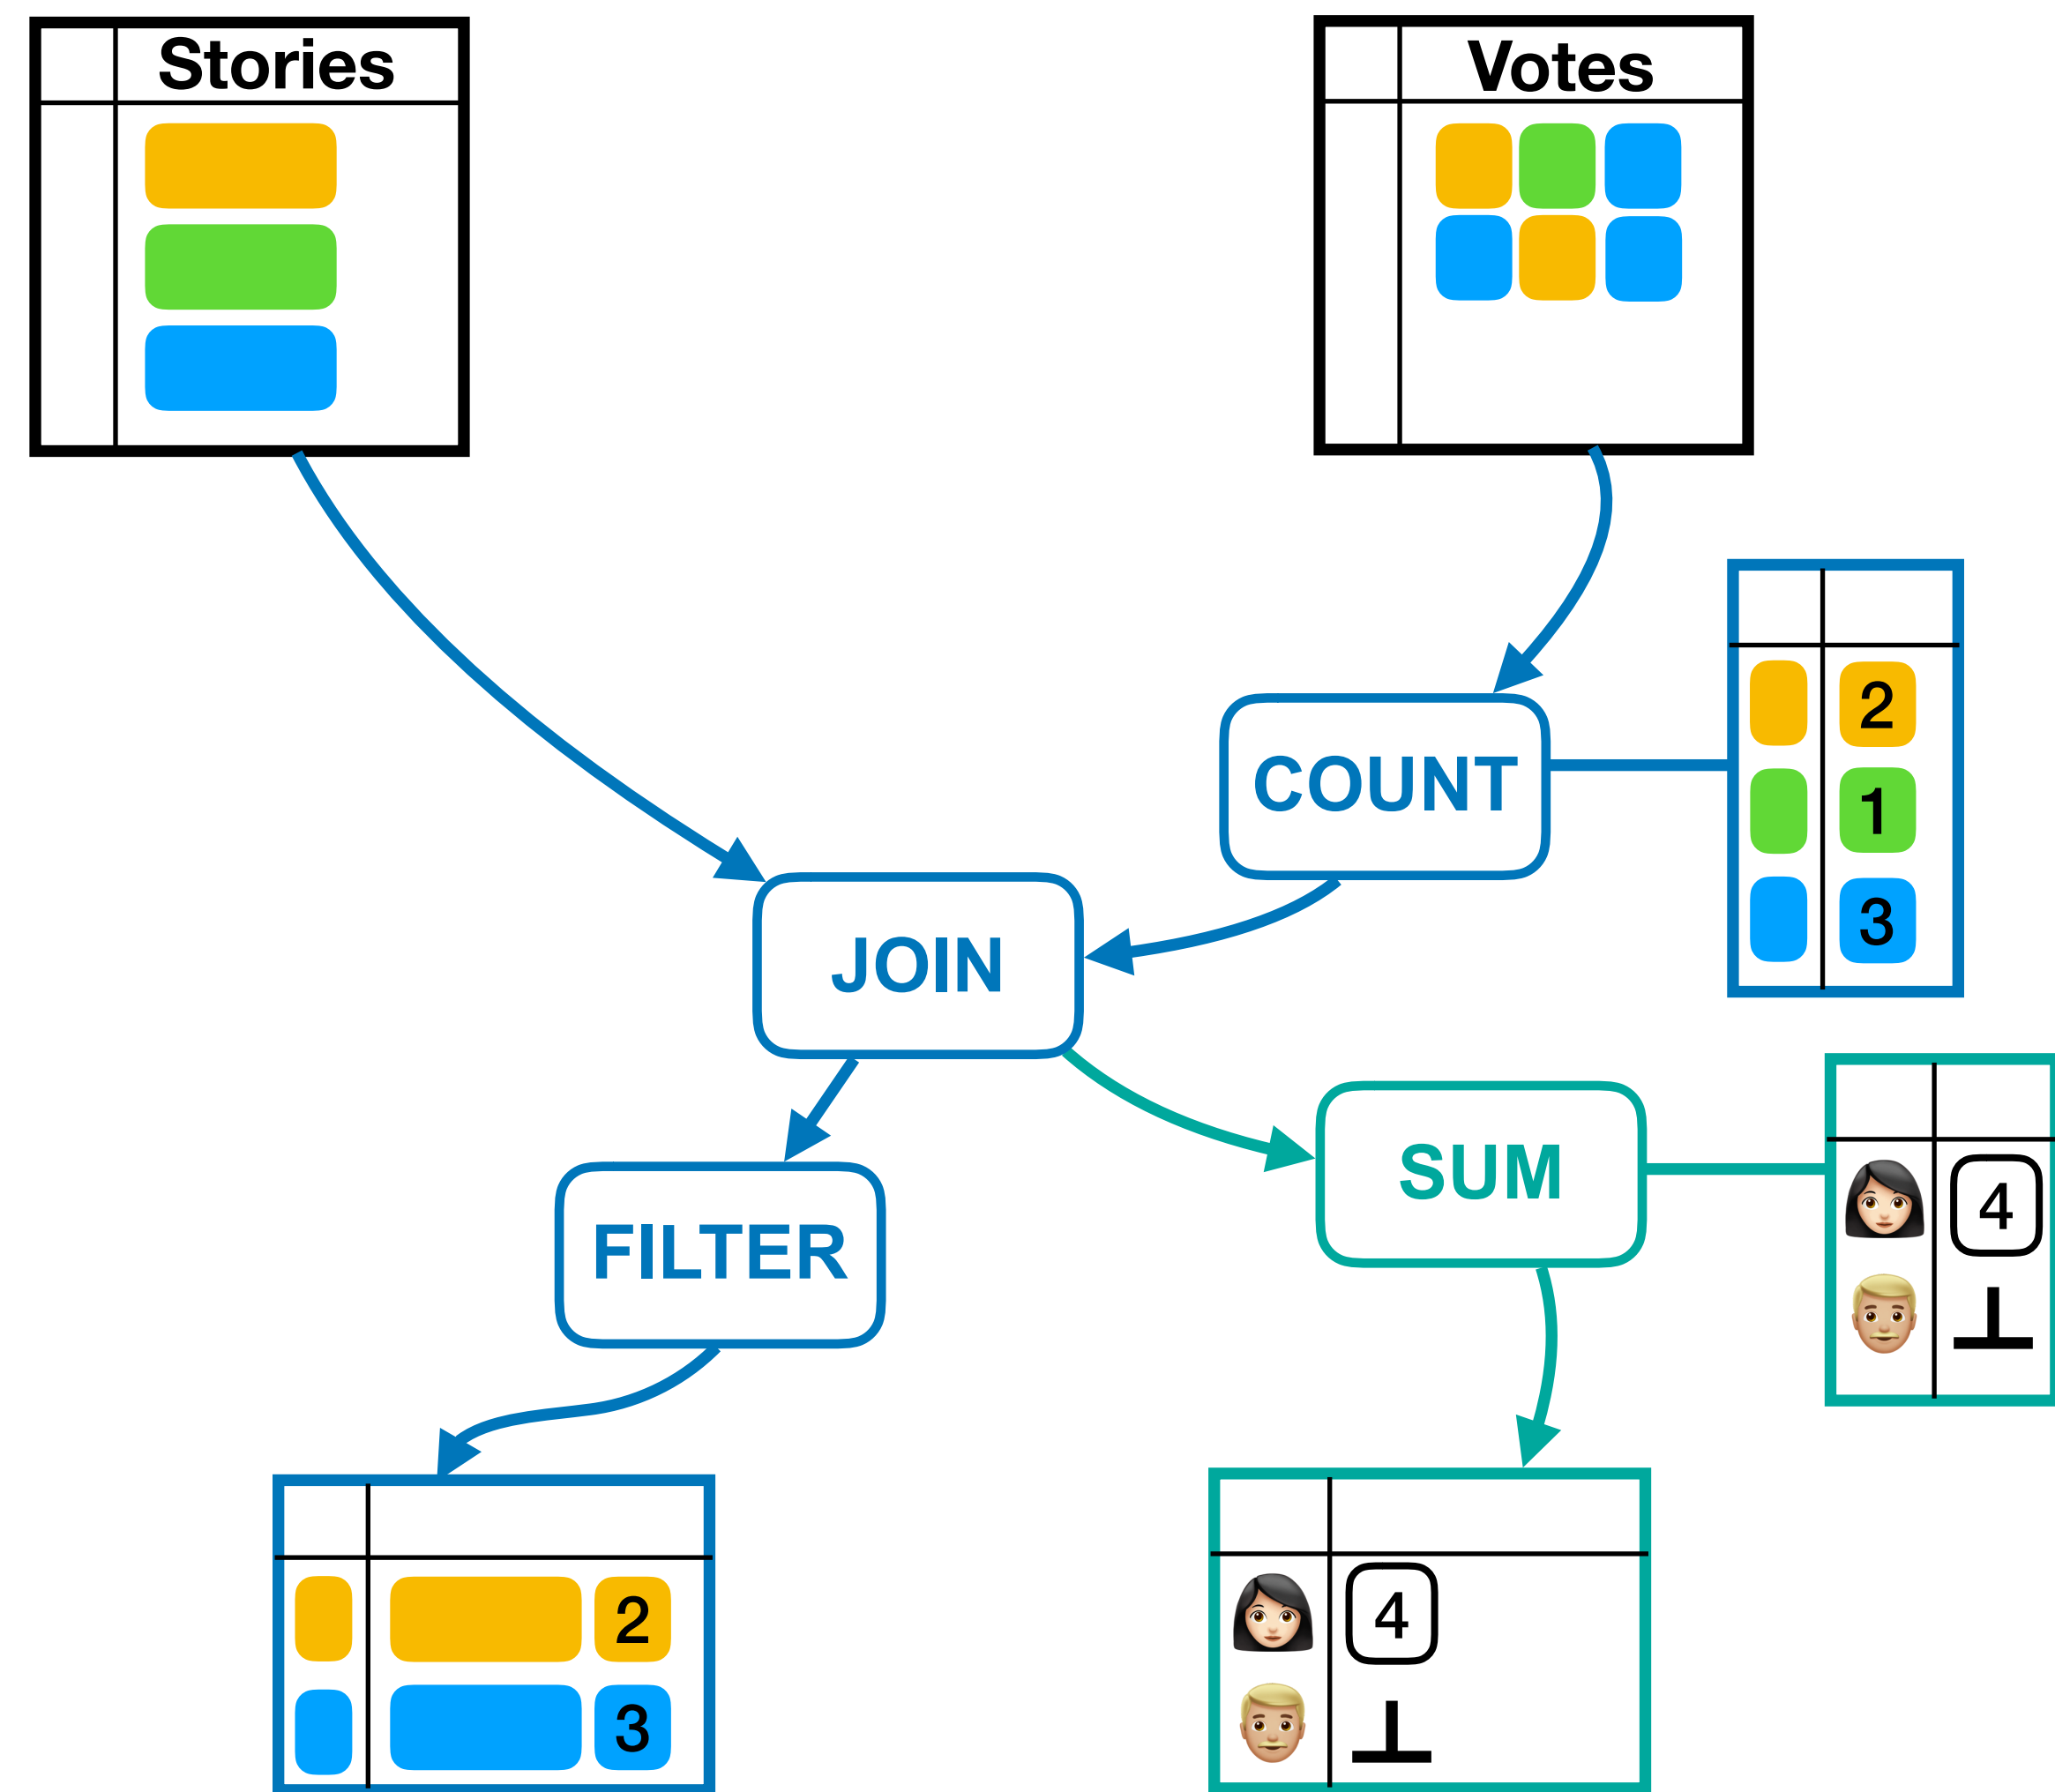


Frontend

# High performance requires concurrency

Process operators concurrently.  
Read from views concurrently.  
Process shards concurrently.

**Without global coordination!**

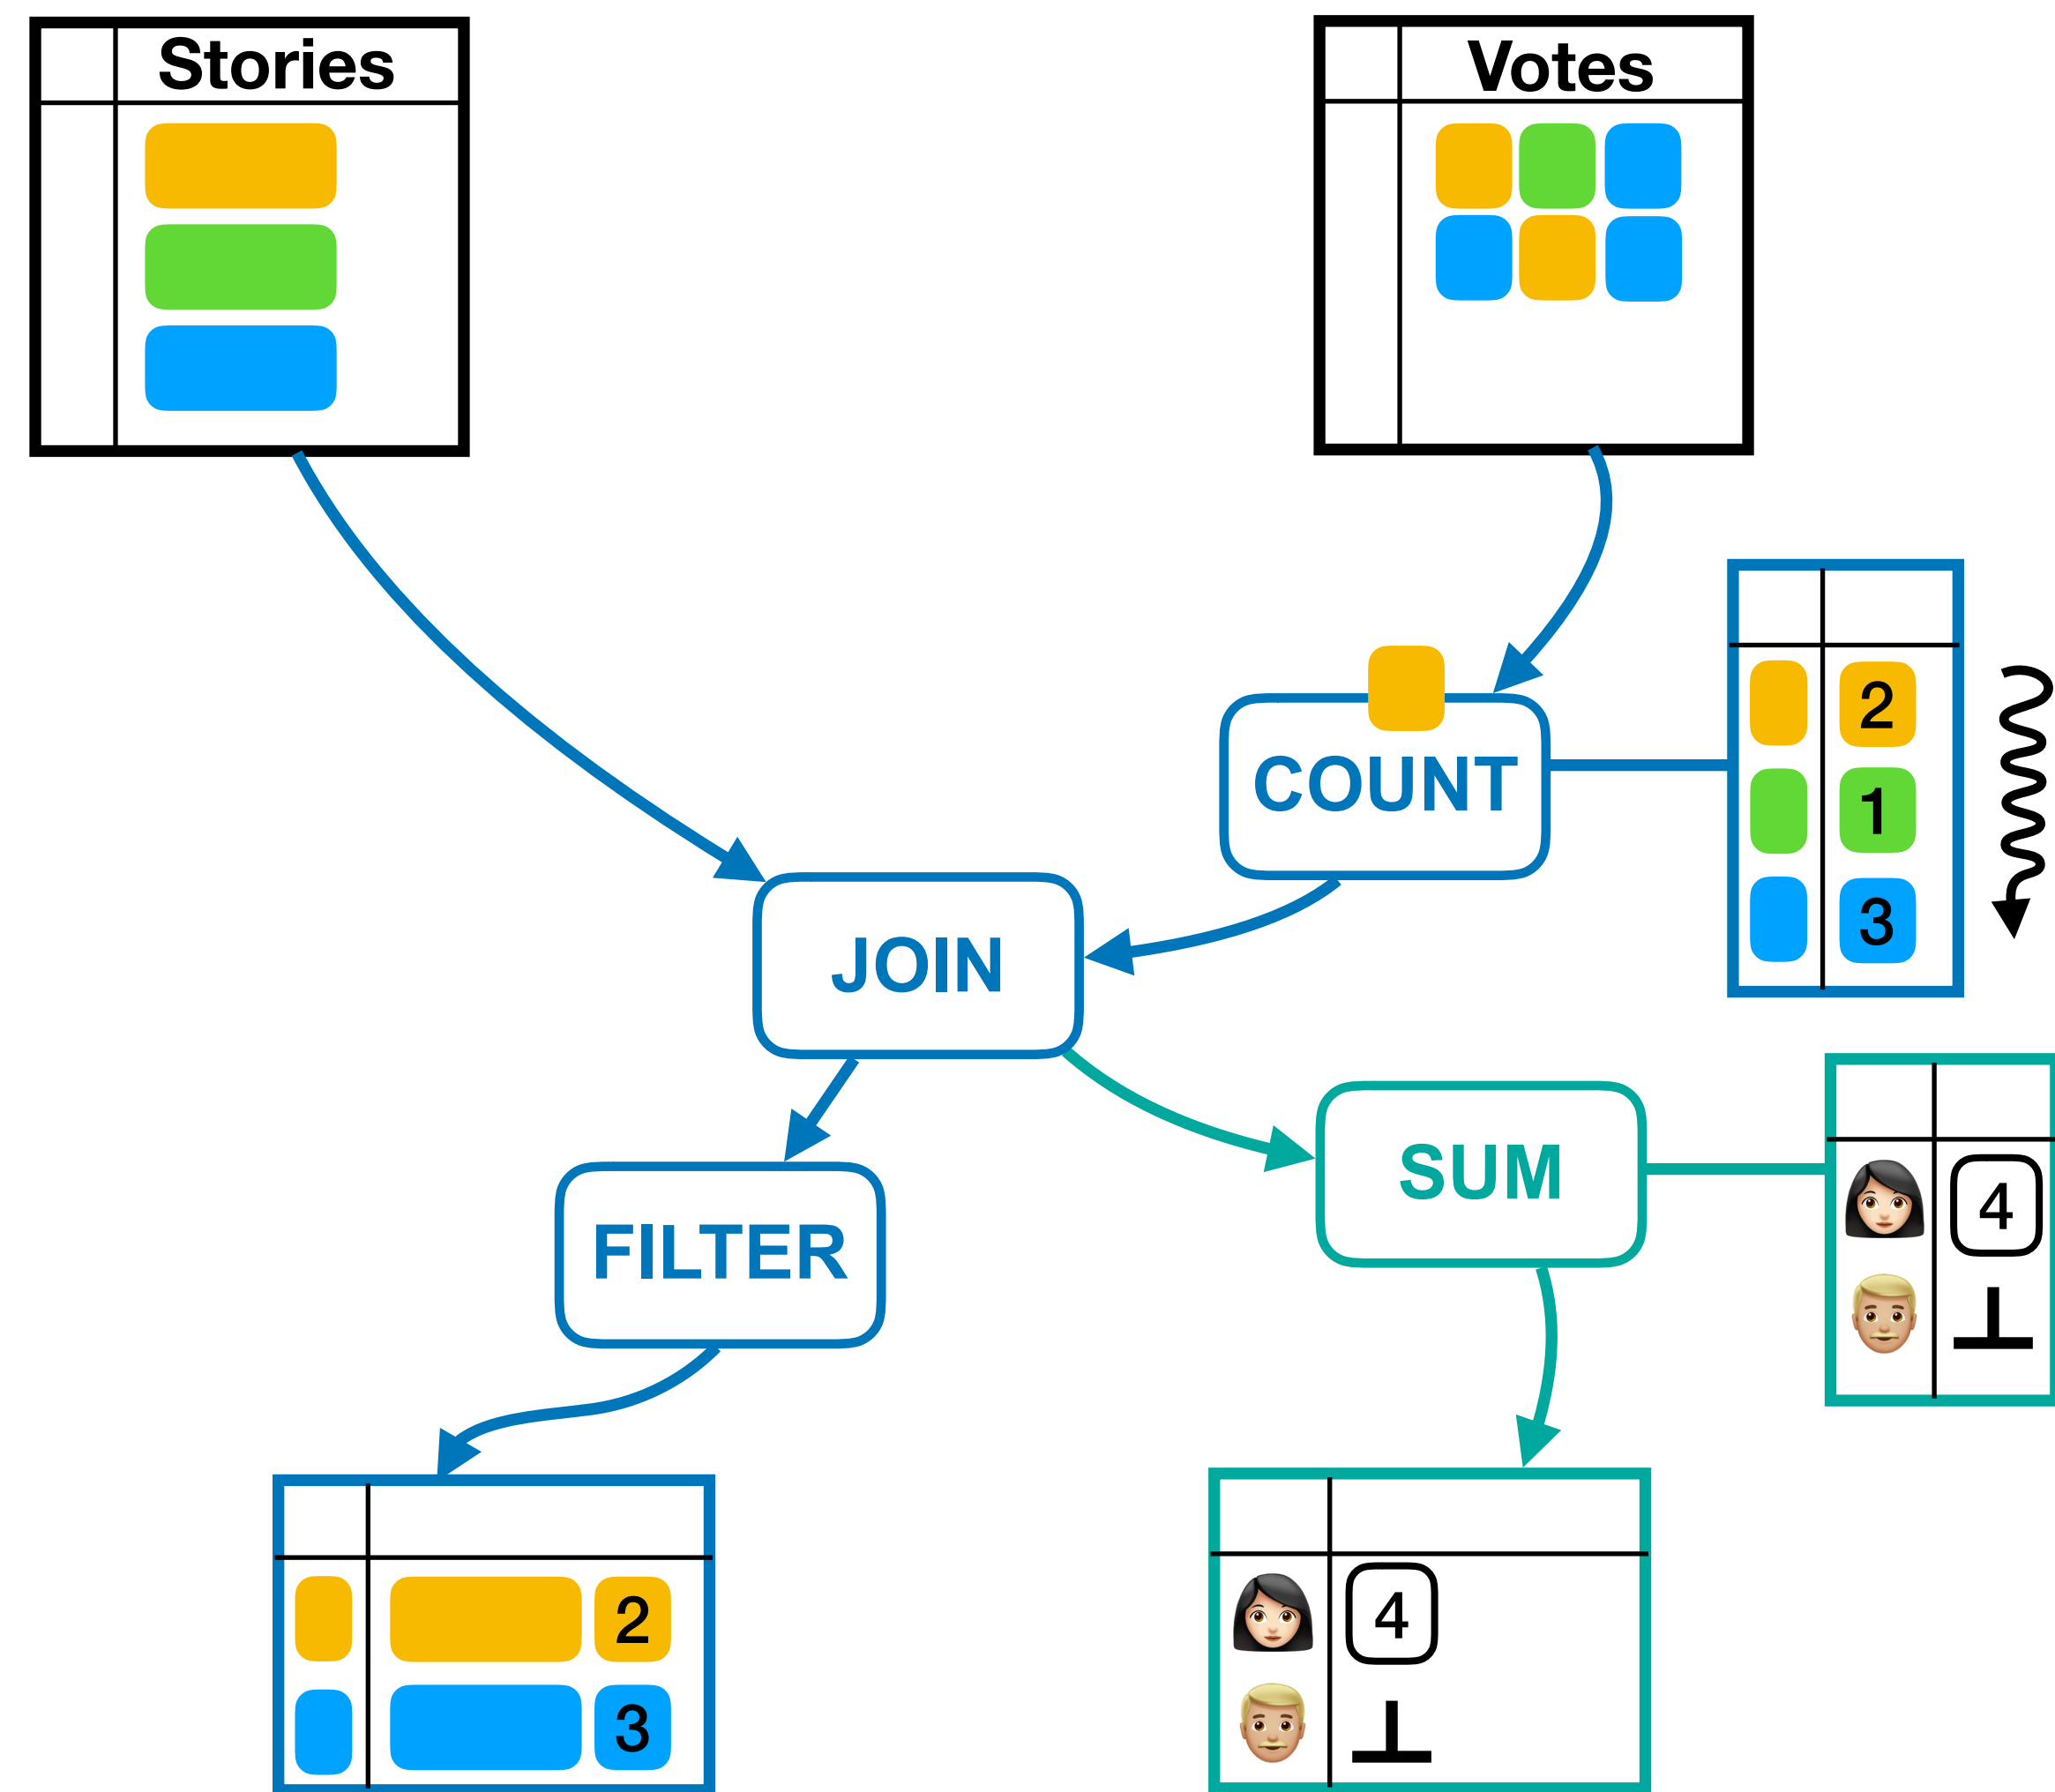


Frontend

# High performance requires concurrency

Process operators concurrently.  
Read from views concurrently.  
Process shards concurrently.

**Without global coordination!**

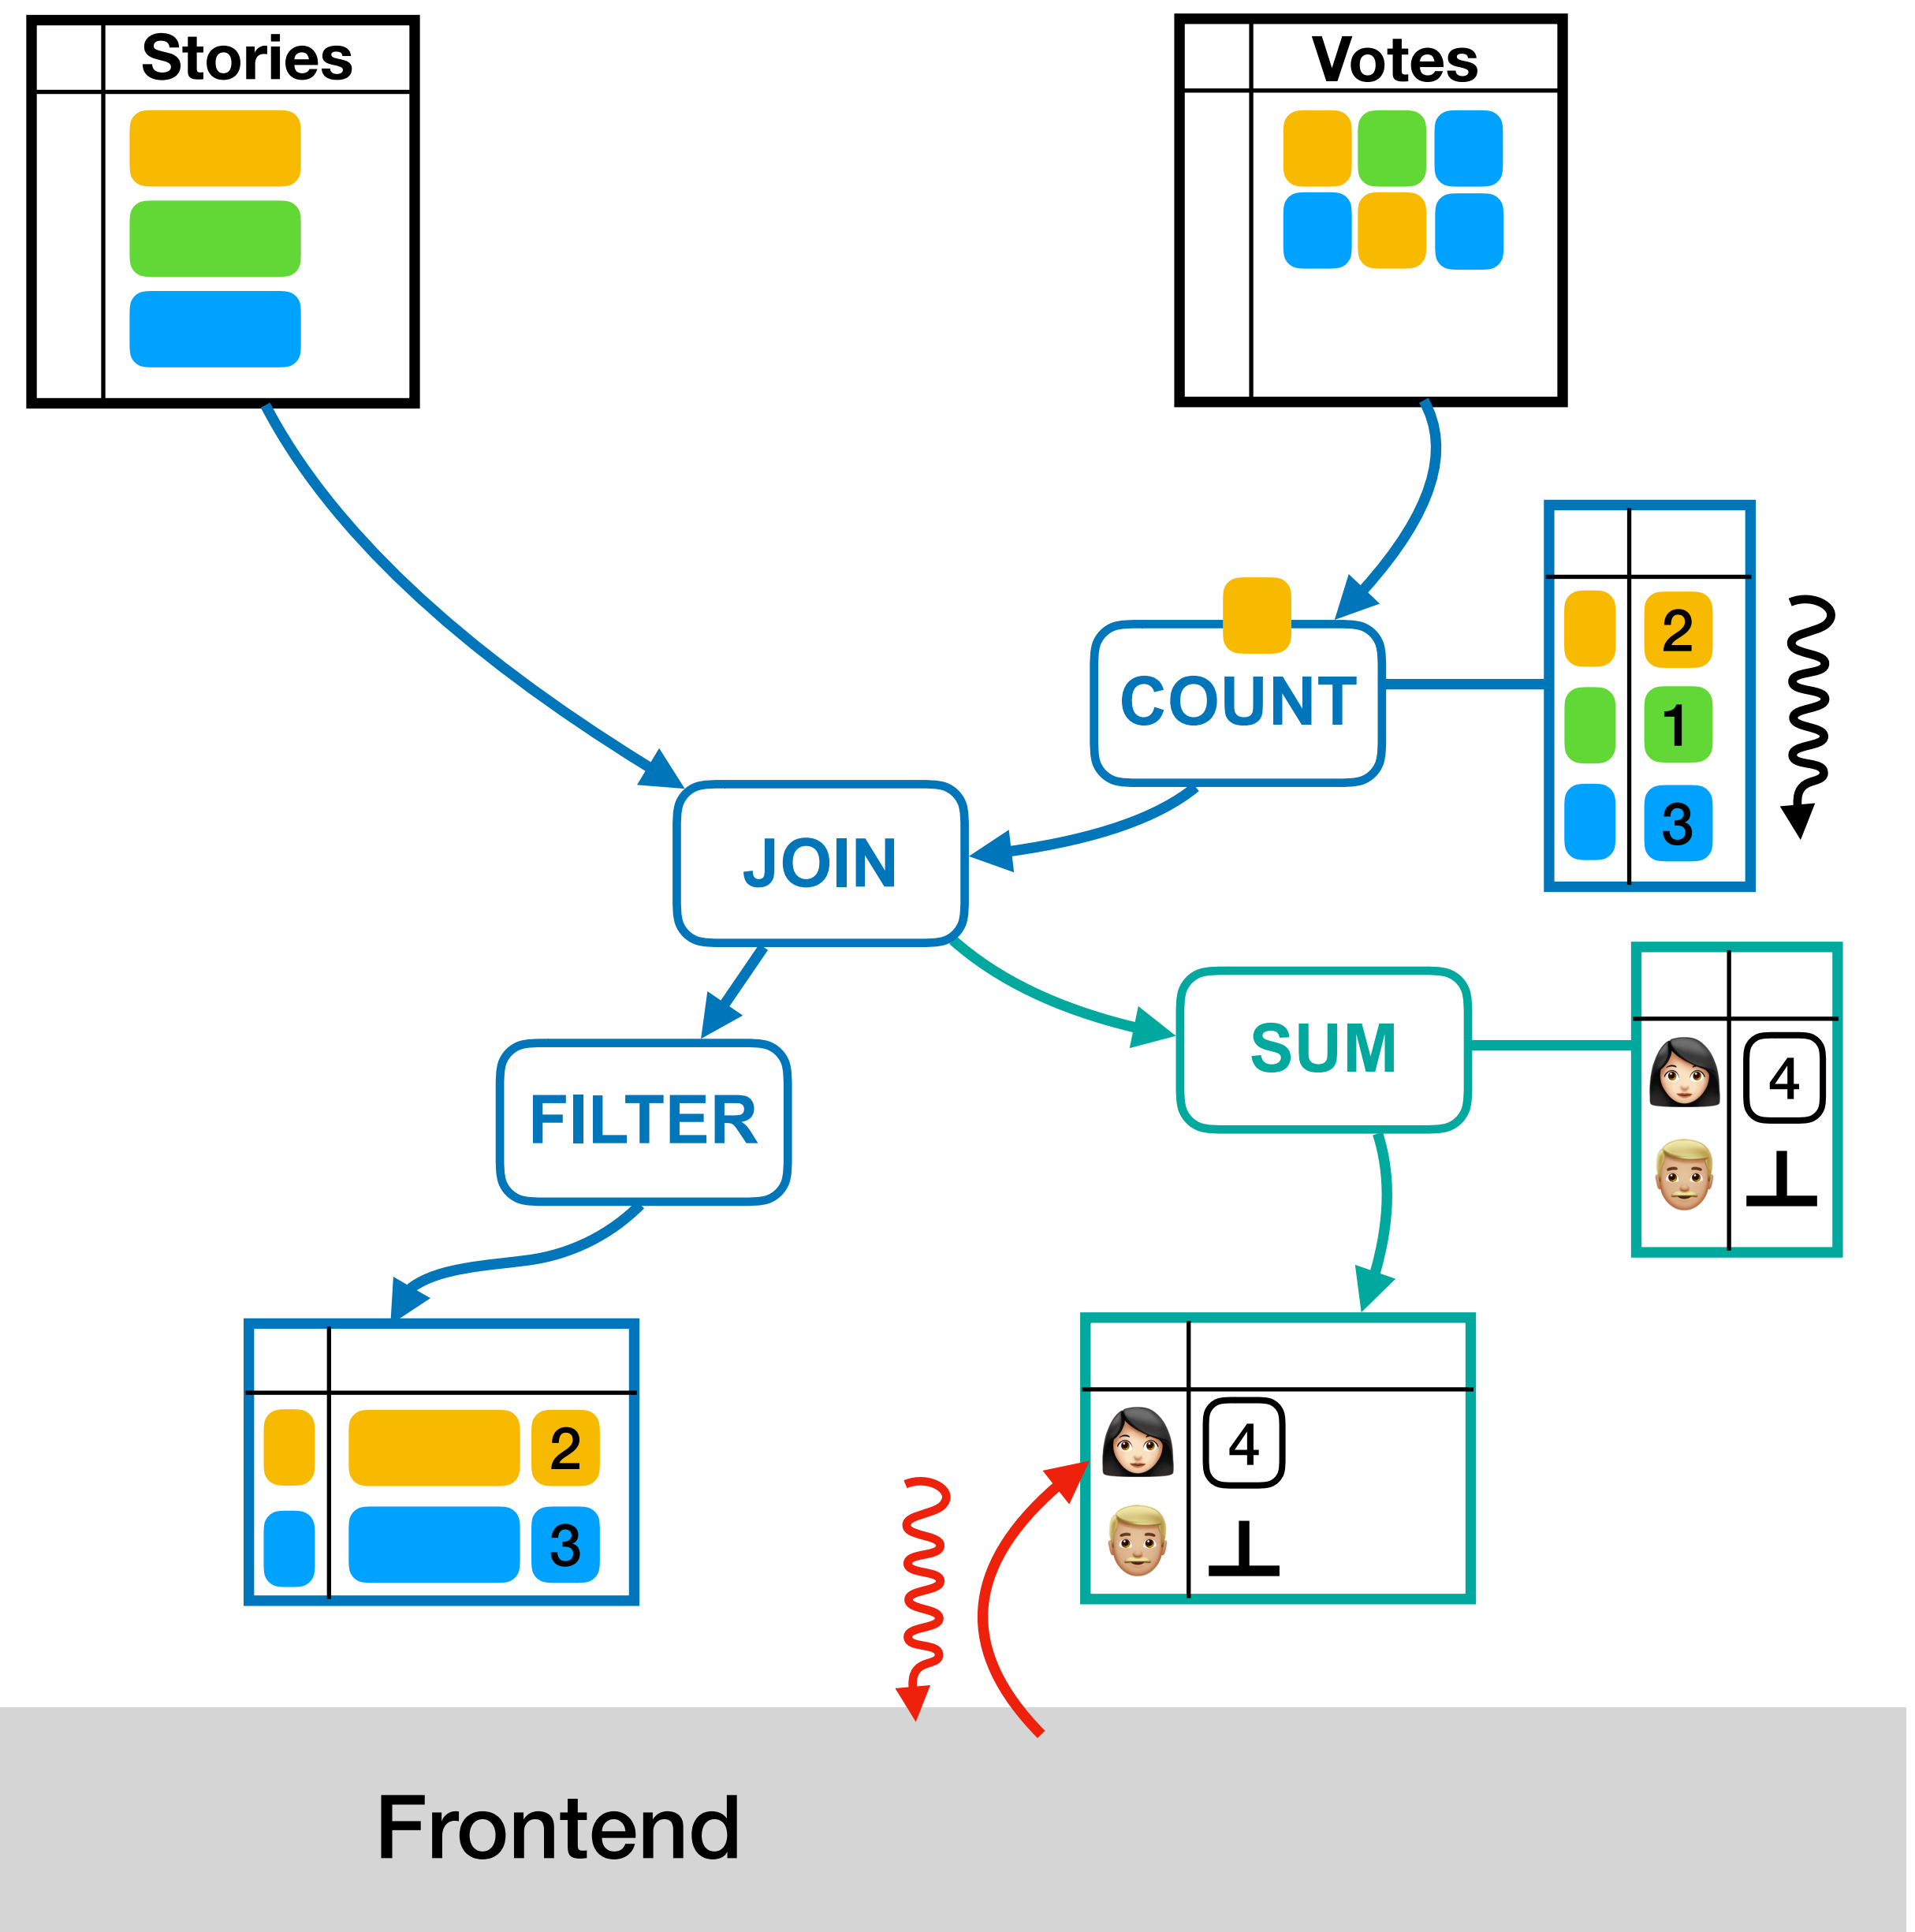


Frontend

# High performance requires concurrency

Process operators concurrently.  
Read from views concurrently.  
Process shards concurrently.

**Without global coordination!**



# Challenges implementing partially-stateful data-flow

# Challenges implementing partially-stateful data-flow

1. Concurrent upqueries and forward processing — races!

Must maintain **correctness** under concurrency!

# Challenges implementing partially-stateful data-flow

1. Concurrent upqueries and forward processing — races!

Must maintain **correctness** under concurrency!


# Correctness under concurrency

**Goal:** upquery restores state as if present all along.




# Correctness under concurrency

**Goal:** upquery restores state as if present all along.




# Correctness under concurrency

**Goal:** upquery restores state as if present all along.




# Correctness under concurrency

**Goal:** upquery restores state as if present all along.




# Correctness under concurrency

**Goal:** upquery restores state as if present all along.



# Correctness under concurrency


**Goal:** upquery restores state as if present all along.



Upquery response is a **snapshot** of state

# Correctness under concurrency

**Goal:** upquery restores state as if present all along.




Upquery response is a **snapshot** of state

includes **2** **1**

does **not** include **3**

# Correctness under concurrency

**Goal:** upquery restores state as if present all along.



Upquery response is a **snapshot** of state

includes **2** **1**

does **not** include 

**Solution:** Maintain **order** of upquery response and surrounding updates, despite lack of global coordination.

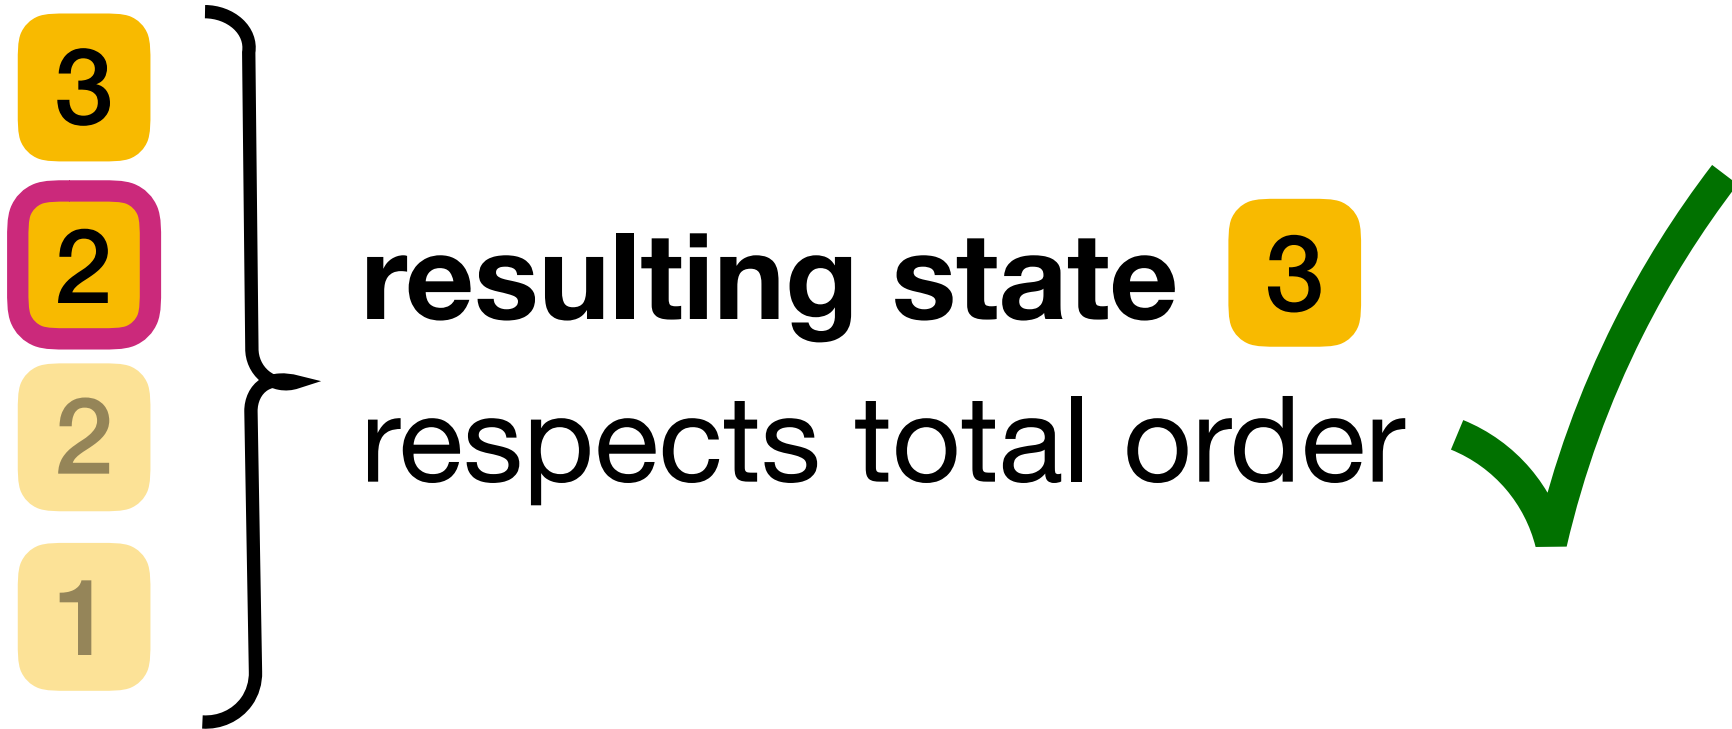
# Upquery responses in total order with updates

**Goal:** upquery restores state as if present all along.



# Upquery responses in total order with updates

**Goal:** upquery restores state as if present all along.



# Upquery responses in total order with updates

**Goal:** upquery restores state as if present all along.

3  
2  
2  
1


resulting state 3  
respects total order ✓

2  
3  
2  
1

resulting state 2  
violates total order ✗

# Upquery responses in total order with updates

**Goal:** upquery restores state as if present all along.



More complex cases: merged upquery responses, evictions (**Paper**).

# Challenges implementing partially-stateful data-flow

1. Concurrent upqueries and forward processing — races!

Must maintain **correctness** under concurrency!

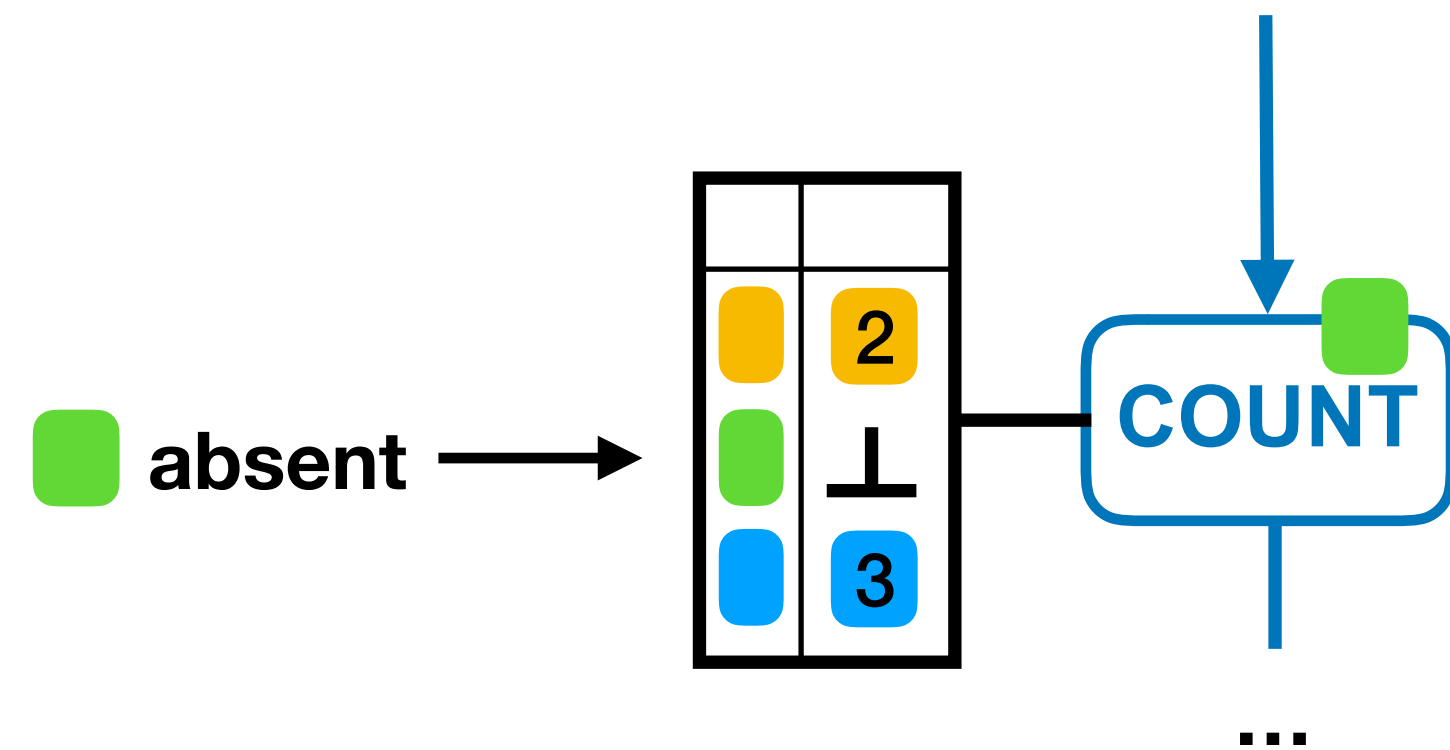
2. Update processing may require absent state

# Challenges implementing partially-stateful data-flow

1. Concurrent upqueries and forward processing — races!

Must maintain **correctness** under concurrency!

2. Update processing may require absent state




# Challenges implementing partially-stateful data-flow

## 1. Concurrent upqueries and forward processing — races!

Must maintain **correctness** under concurrency!

## 2. Update processing may require absent state




Drop updates that touch absent state, future upquery repeats them.

# Challenges implementing partially-stateful data-flow

## 1. Concurrent upqueries and forward processing — races!

Must maintain **correctness** under concurrency!

## 2. Update processing may require absent state (see Paper)



Drop updates that touch absent state, future upquery repeats them.

# Noria implementation



# Noria implementation

```
1 /* base tables */
2 CREATE TABLE stories
3   (id int, author int, title text, url text);
4 CREATE TABLE votes (user int, story_id int);
5 CREATE TABLE users (id int, username text);
6 /* internal view: vote count per story */
7 CREATE INTERNAL VIEW VoteCount AS
8   SELECT story_id, COUNT(*) AS vcount
9     FROM votes GROUP BY story_id;
10 /* external view: story details */
11 CREATE VIEW StoriesWithVC AS
12   SELECT id, author, title, url, vcount
13     FROM stories
14     JOIN VoteCount ON VoteCount.story_id = stories.id
15     WHERE stories.id = ?;
```

# Noria implementation

```
1 /* base tables */
2 CREATE TABLE stories
3   (id int, author int, title text, url text);
4 CREATE TABLE votes (user int, story_id int);
5 CREATE TABLE users (id int, username text);
6 /* internal view: vote count per story */
7 CREATE INTERNAL VIEW VoteCount AS
8   SELECT story_id, COUNT(*) AS vcount
9     FROM votes GROUP BY story_id;
10 /* external view: story details */
11 CREATE VIEW StoriesWithVC AS
12   SELECT id, author, title, url, vcount
13     FROM stories
14     JOIN VoteCount ON VoteCount.story_id = stories.id
15     WHERE stories.id = ?;
```




MySQL adapter

# Noria implementation

```
1 /* base tables */
2 CREATE TABLE stories
3   (id int, author int, title text, url text);
4 CREATE TABLE votes (user int, story_id int);
5 CREATE TABLE users (id int, username text);
6 /* internal view: vote count per story */
7 CREATE INTERNAL VIEW VoteCount AS
8   SELECT story_id, COUNT(*) AS vcount
9     FROM votes GROUP BY story_id;
10 /* external view: story details */
11 CREATE VIEW StoriesWithVC AS
12   SELECT id, author, title, url, vcount
13     FROM stories
14     JOIN VoteCount ON VoteCount.story_id = stories.id
15     WHERE stories.id = ?;
```

Transform

## Data-flow graph




MySQL adapter

# Noria implementation

```
1 /* base tables */
2 CREATE TABLE stories
3   (id int, author int, title text, url text);
4 CREATE TABLE votes (user int, story_id int);
5 CREATE TABLE users (id int, username text);
6 /* internal view: vote count per story */
7 CREATE INTERNAL VIEW VoteCount AS
8   SELECT story_id, COUNT(*) AS vcount
9     FROM votes GROUP BY story_id;
10 /* external view: story details */
11 CREATE VIEW StoriesWithVC AS
12   SELECT id, author, title, url, vcount
13     FROM stories
14     JOIN VoteCount ON VoteCount.story_id = stories.id
15     WHERE stories.id = ?;
```

Transform

## Data-flow graph




MySQL adapter

# Noria implementation

```
1 /* base tables */
2 CREATE TABLE stories
3   (id int, author int, title text, url text);
4 CREATE TABLE votes (user int, story_id int);
5 CREATE TABLE users (id int, username text);
6 /* internal view: vote count per story */
7 CREATE INTERNAL VIEW VoteCount AS
8   SELECT story_id, COUNT(*) AS vcount
9     FROM votes GROUP BY story_id;
10 /* external view: story details */
11 CREATE VIEW StoriesWithVC AS
12   SELECT id, author, title, url, vcount
13     FROM stories
14     JOIN VoteCount ON VoteCount.story_id = stories.id
15     WHERE stories.id = ?;
```

Transform

## Data-flow graph



MySQL adapter

- 45k lines of Rust + 15k libraries
- RocksDB for base table storage
- ZooKeeper for leader election

# Evaluation

1. Can Noria improve a real web application's performance?
2. How does Noria compare to alternative approaches?
3. Can Noria change queries without downtime?

# Evaluation

1. Can Noria improve a real web application's performance?
  2. How does Noria compare to alternative approaches?
  3. Can Noria change queries without downtime?
- 

**Setup** Amazon EC2 c5.4xlarge instance (16 vCPUs)

Open-loop clients, measuring latency & throughput

# Evaluation

1. Can Noria improve a real web application's performance?
2. How does Noria compare to alternative approaches?
3. Can Noria change queries without downtime?

---

**Setup** Amazon EC2 c5.4xlarge instance (16 vCPUs)

Open-loop clients, measuring latency & throughput













---

multi-machine experiments } see Paper  
comparison with differential dataflow }















# Case study: Lobsters (<http://lobste.rs>)

**L** Lobsters [Recent](#) [Comments](#) [Search](#) Login

- ▲ **Falling in love with Rust** programming rust *dtrace.org*  
40  via blake 11 hours ago | cached | 7 comments
- ▲ **FreeBSD Desktop - Part 16 - Configuration - Pause Any Application** freebsd illumos linux *vermaden.wordpress.com*  
6  authored by vermaden 6 hours ago | cached | 4 comments
- ▲ **You Think the Visual Studio Code binary you use is a Free Software? Think again** law privacy *carlchenet.com*  
61  authored by chaica 31 hours ago | cached | 30 comments
- ▲ **Using Make — writing less Makefile** programming *text.causal.agency*  
27  authored by causal\_agent 23 hours ago | cached | 11 comments
- ▲ **LLVM 7.0.0 Release** release *lists.llvm.org*  
4  via colin 43 minutes ago | cached | no comments
- ▲ **Kit programming language** programming *kitlang.org*  
27  via btbytes 25 hours ago | cached | 21 comments
- ▲ **Why Aren't More Users More Happy With Our VMs? Part 2** performance *tratt.net*  
5  via edd 4 hours ago | cached | no comments
- ▲ **Spleen - Monospaced bitmap fonts** design *cambus.net*  
1  authored by fcambus 3 minutes ago | cached | no comments
- ▲ **The Singleton module in Ruby - Part I** ruby *medium.com*  
1  authored by mehdi-farsi 14 minutes ago | cached | no comments
- ▲ **You Can't Always Hash Pointers in C** c *nullprogram.com*  
3  via calvin 4 hours ago | cached | 1 comment
- ▲ **Learn Go with Seam Carving and Rockets** go *getstream.io*  
1  authored by tschellenbach 18 minutes ago | cached | no comments
- ▲ **Times Newer Roman is a sneaky font designed to make your essays look longer** education graphics *theverge.com*  
9  via Ricardus 15 hours ago | cached | 5 comments

# Case study: Lobsters (<http://lobste.rs>)













**L** Lobsters [Recent](#) [Comments](#) [Search](#) [Login](#)

- ▲ **Falling in love with Rust** [programming](#) [rust](#) *dtrace.org*  
40  via blake 11 hours ago | cached | 7 comments
- ▲ **FreeBSD Desktop - Part 16 - Configuration - Pause Any Application** [freebsd](#) [illumos](#) [linux](#) *vermaden.wordpress.com*  
6  authored by vermaden 6 hours ago | cached | 4 comments
- ▲ **You Think the Visual Studio Code binary you use is a Free Software? Think again** [law](#) [privacy](#) *carlchenet.com*  
61  authored by chaica 31 hours ago | cached | 30 comments
- ▲ **Using Make — writing less Makefile** [programming](#) *text.causal.agency*  
27  authored by causal\_agent 23 hours ago | cached | 11 comments
- ▲ **LLVM 7.0.0 Release** [release](#) *lists.llvm.org*  
4  via colin 43 minutes ago | cached | no comments
- ▲ **Kit programming language** [programming](#) *kitlang.org*  
27  via btbytes 25 hours ago | cached | 21 comments
- ▲ **Why Aren't More Users More Happy With Our VMs? Part 2** [performance](#) *tratt.net*  
5  via edd 4 hours ago | cached | no comments
- ▲ **Spleen - Monospaced bitmap fonts** [design](#) *cambus.net*  
1  authored by fcambus 3 minutes ago | cached | no comments
- ▲ **The Singleton module in Ruby - Part I** [ruby](#) *medium.com*  
1  authored by mehdi-farsi 14 minutes ago | cached | no comments
- ▲ **You Can't Always Hash Pointers in C** [c](#) *nullprogram.com*  
3  via calvin 4 hours ago | cached | 1 comment
- ▲ **Learn Go with Seam Carving and Rockets** [go](#) *getstream.io*  
1  authored by tschellenbach 18 minutes ago | cached | no comments
- ▲ **Times Newer Roman is a sneaky font designed to make your essays look longer** [education](#) [graphics](#) *theverge.com*  
9  via Ricardus 15 hours ago | cached | 5 comments

► Ruby-on-Rails application with MySQL backend

# Case study: Lobsters (<http://lobste.rs>)

**L** Lobsters [Recent](#) [Comments](#) [Search](#) [Login](#)

- ▲ **Falling in love with Rust** [programming](#) [rust](#) *dtrace.org*  
40  via blake 11 hours ago | [cached](#) | 7 comments
- ▲ **FreeBSD Desktop - Part 16 - Configuration - Pause Any Application** [freebsd](#) [illumos](#) [linux](#) *vermaden.wordpress.com*  
6  authored by vermaden 6 hours ago | [cached](#) | 4 comments
- ▲ **You Think the Visual Studio Code binary you use is a Free Software? Think again** [law](#) [privacy](#) *carlchenet.com*  
61  authored by chaica 31 hours ago | [cached](#) | 30 comments
- ▲ **Using Make — writing less Makefile** [programming](#) *text.causal.agency*  
27  authored by causal\_agent 23 hours ago | [cached](#) | 11 comments
- ▲ **LLVM 7.0.0 Release** [release](#) *lists.llvm.org*  
4  via colin 43 minutes ago | [cached](#) | no comments
- ▲ **Kit programming language** [programming](#) *kitlang.org*  
27  via btbytes 25 hours ago | [cached](#) | 21 comments
- ▲ **Why Aren't More Users More Happy With Our VMs? Part 2** [performance](#) *tratt.net*  
5  via edd 4 hours ago | [cached](#) | no comments
- ▲ **Spleen - Monospaced bitmap fonts** [design](#) *cambus.net*  
1  authored by fcambus 3 minutes ago | [cached](#) | no comments
- ▲ **The Singleton module in Ruby - Part I** [ruby](#) *medium.com*  
1  authored by mehdi-farsi 14 minutes ago | [cached](#) | no comments
- ▲ **You Can't Always Hash Pointers in C** [c](#) *nullprogram.com*  
3  via calvin 4 hours ago | [cached](#) | 1 comment
- ▲ **Learn Go with Seam Carving and Rockets** [go](#) *getstream.io*  
1  authored by tschellenbach 18 minutes ago | [cached](#) | no comments
- ▲ **Times Newer Roman is a sneaky font designed to make your essays look longer** [education](#) [graphics](#) *theverge.com*  
9  via Ricardus 15 hours ago | [cached](#) | 5 comments

- ▶ Ruby-on-Rails application with MySQL backend
- ▶ Hand-optimized by developers to pre-compute aggregations

# Case study: Lobsters (<http://lobste.rs>)

**L** Lobsters Recent Comments Search Login

- ▲ **Falling in love with Rust** programming rust *dtrace.org*  
40 via blake 11 hours ago | cached | 7 comments
- ▲ **FreeBSD Desktop - Part 16 - Configuration - Pause Any Application** freebsd illumos linux *vermaden.wordpress.com*  
6 authored by vermaden 6 hours ago | cached | 4 comments
- ▲ **You Think the Visual Studio Code binary you use is a Free Software? Think again** law privacy *carlchenet.com*  
61 authored by chaica 31 hours ago | cached | 30 comments
- ▲ **Using Make — writing less Makefile** programming *text.causal.agency*  
27 authored by causal\_agent 23 hours ago | cached | 11 comments
- ▲ **LLVM 7.0.0 Release** release *lists.llvm.org*  
4 via colin 43 minutes ago | cached | no comments
- ▲ **Kit programming language** programming *kitlang.org*  
27 via btbytes 25 hours ago | cached | 21 comments
- ▲ **Why Aren't More Users More Happy With Our VMs? Part 2** performance *tratt.net*  
5 via edd 4 hours ago | cached | no comments
- ▲ **Spleen - Monospaced bitmap fonts** design *cambus.net*  
1 authored by fcambus 3 minutes ago | cached | no comments
- ▲ **The Singleton module in Ruby - Part I** ruby *medium.com*  
1 authored by mehdi-farsi 14 minutes ago | cached | no comments
- ▲ **You Can't Always Hash Pointers in C** c *nullprogram.com*  
3 via calvin 4 hours ago | cached | 1 comment
- ▲ **Learn Go with Seam Carving and Rockets** go *getstream.io*  
1 authored by tschellenbach 18 minutes ago | cached | no comments
- ▲ **Times Newer Roman is a sneaky font designed to make your essays look longer** education graphics *theverge.com*  
9 via Ricardus 15 hours ago | cached | 5 comments

- ▶ Ruby-on-Rails application with MySQL backend
- ▶ Hand-optimized by developers to pre-compute aggregations
- ▶ Noria data-flow with 235 operators, 35 views


# Case study: Lobsters (<http://lobste.rs>)

**L** Lobsters [Recent](#) [Comments](#) [Search](#) [Login](#)

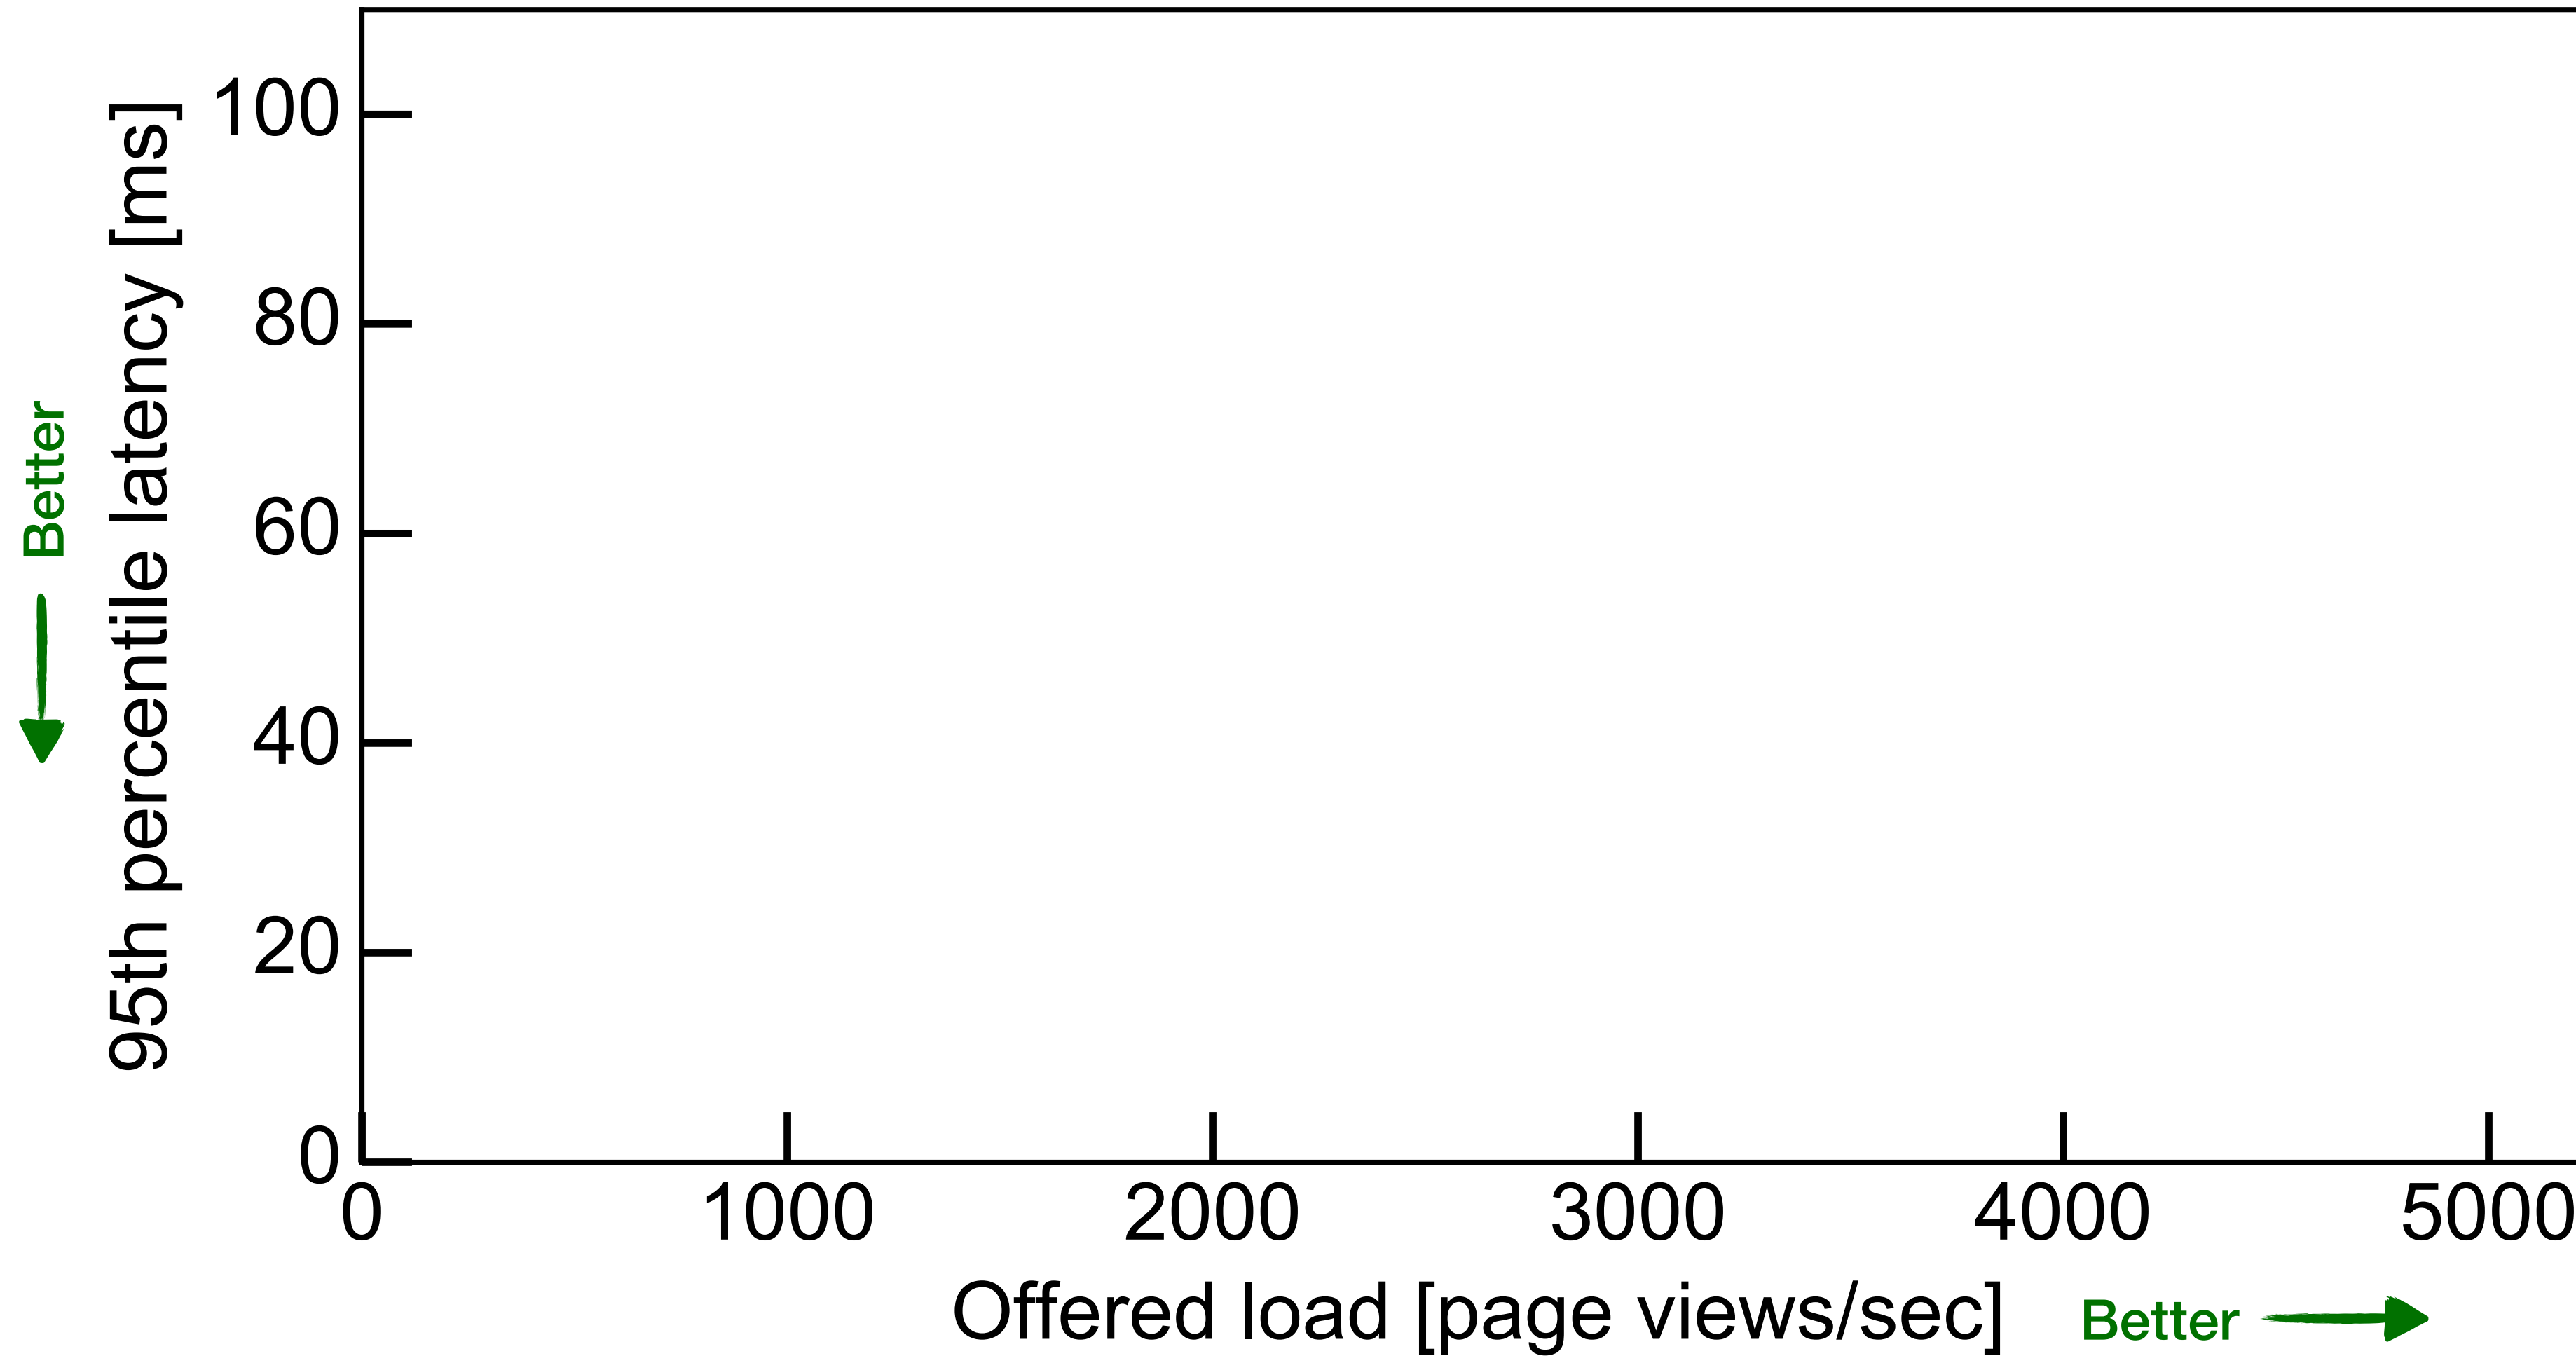
- ▲ **Falling in love with Rust** [programming](#) [rust](#) *dtrace.org*  
40 via blake 11 hours ago | cached | 7 comments
- ▲ **FreeBSD Desktop - Part 16 - Configuration - Pause Any Application** [freebsd](#) [illumos](#) [linux](#) *vermaden.wordpress.com*  
6 authored by vermaden 6 hours ago | cached | 4 comments
- ▲ **You Think the Visual Studio Code binary you use is a Free Software? Think again** [law](#) [privacy](#) *carlchenet.com*  
61 authored by chaica 31 hours ago | cached | 30 comments
- ▲ **Using Make — writing less Makefile** [programming](#) *text.causal.agency*  
27 authored by causal\_agent 23 hours ago | cached | 11 comments
- ▲ **LLVM 7.0.0 Release** [release](#) *lists.llvm.org*  
4 via colin 43 minutes ago | cached | no comments
- ▲ **Kit programming language** [programming](#) *kitlang.org*  
27 via btbytes 25 hours ago | cached | 21 comments
- ▲ **Why Aren't More Users More Happy With Our VMs? Part 2** [performance](#) *tratt.net*  
5 via edd 4 hours ago | cached | no comments
- ▲ **Spleen - Monospaced bitmap fonts** [design](#) *cambus.net*  
1 authored by fcambus 3 minutes ago | cached | no comments
- ▲ **The Singleton module in Ruby - Part I** [ruby](#) *medium.com*  
1 authored by mehdi-farsi 14 minutes ago | cached | no comments
- ▲ **You Can't Always Hash Pointers in C** [c](#) *nullprogram.com*  
3 via calvin 4 hours ago | cached | 1 comment
- ▲ **Learn Go with Seam Carving and Rockets** [go](#) *getstream.io*  
1 authored by tschellenbach 18 minutes ago | cached | no comments
- ▲ **Times Newer Roman is a sneaky font designed to make your essays look longer** [education](#) [graphics](#) *theverge.com*  
9 via Ricardus 15 hours ago | cached | 5 comments

- ▶ Ruby-on-Rails application with MySQL backend
- ▶ Hand-optimized by developers to pre-compute aggregations
- ▶ Noria data-flow with 235 operators, 35 views
- ▶ Emulate production load


# Can Noria improve Lobsters' performance?



# Can Noria improve Lobsters' performance?




# Can Noria improve Lobsters' performance?






# Can Noria improve Lobsters' performance?

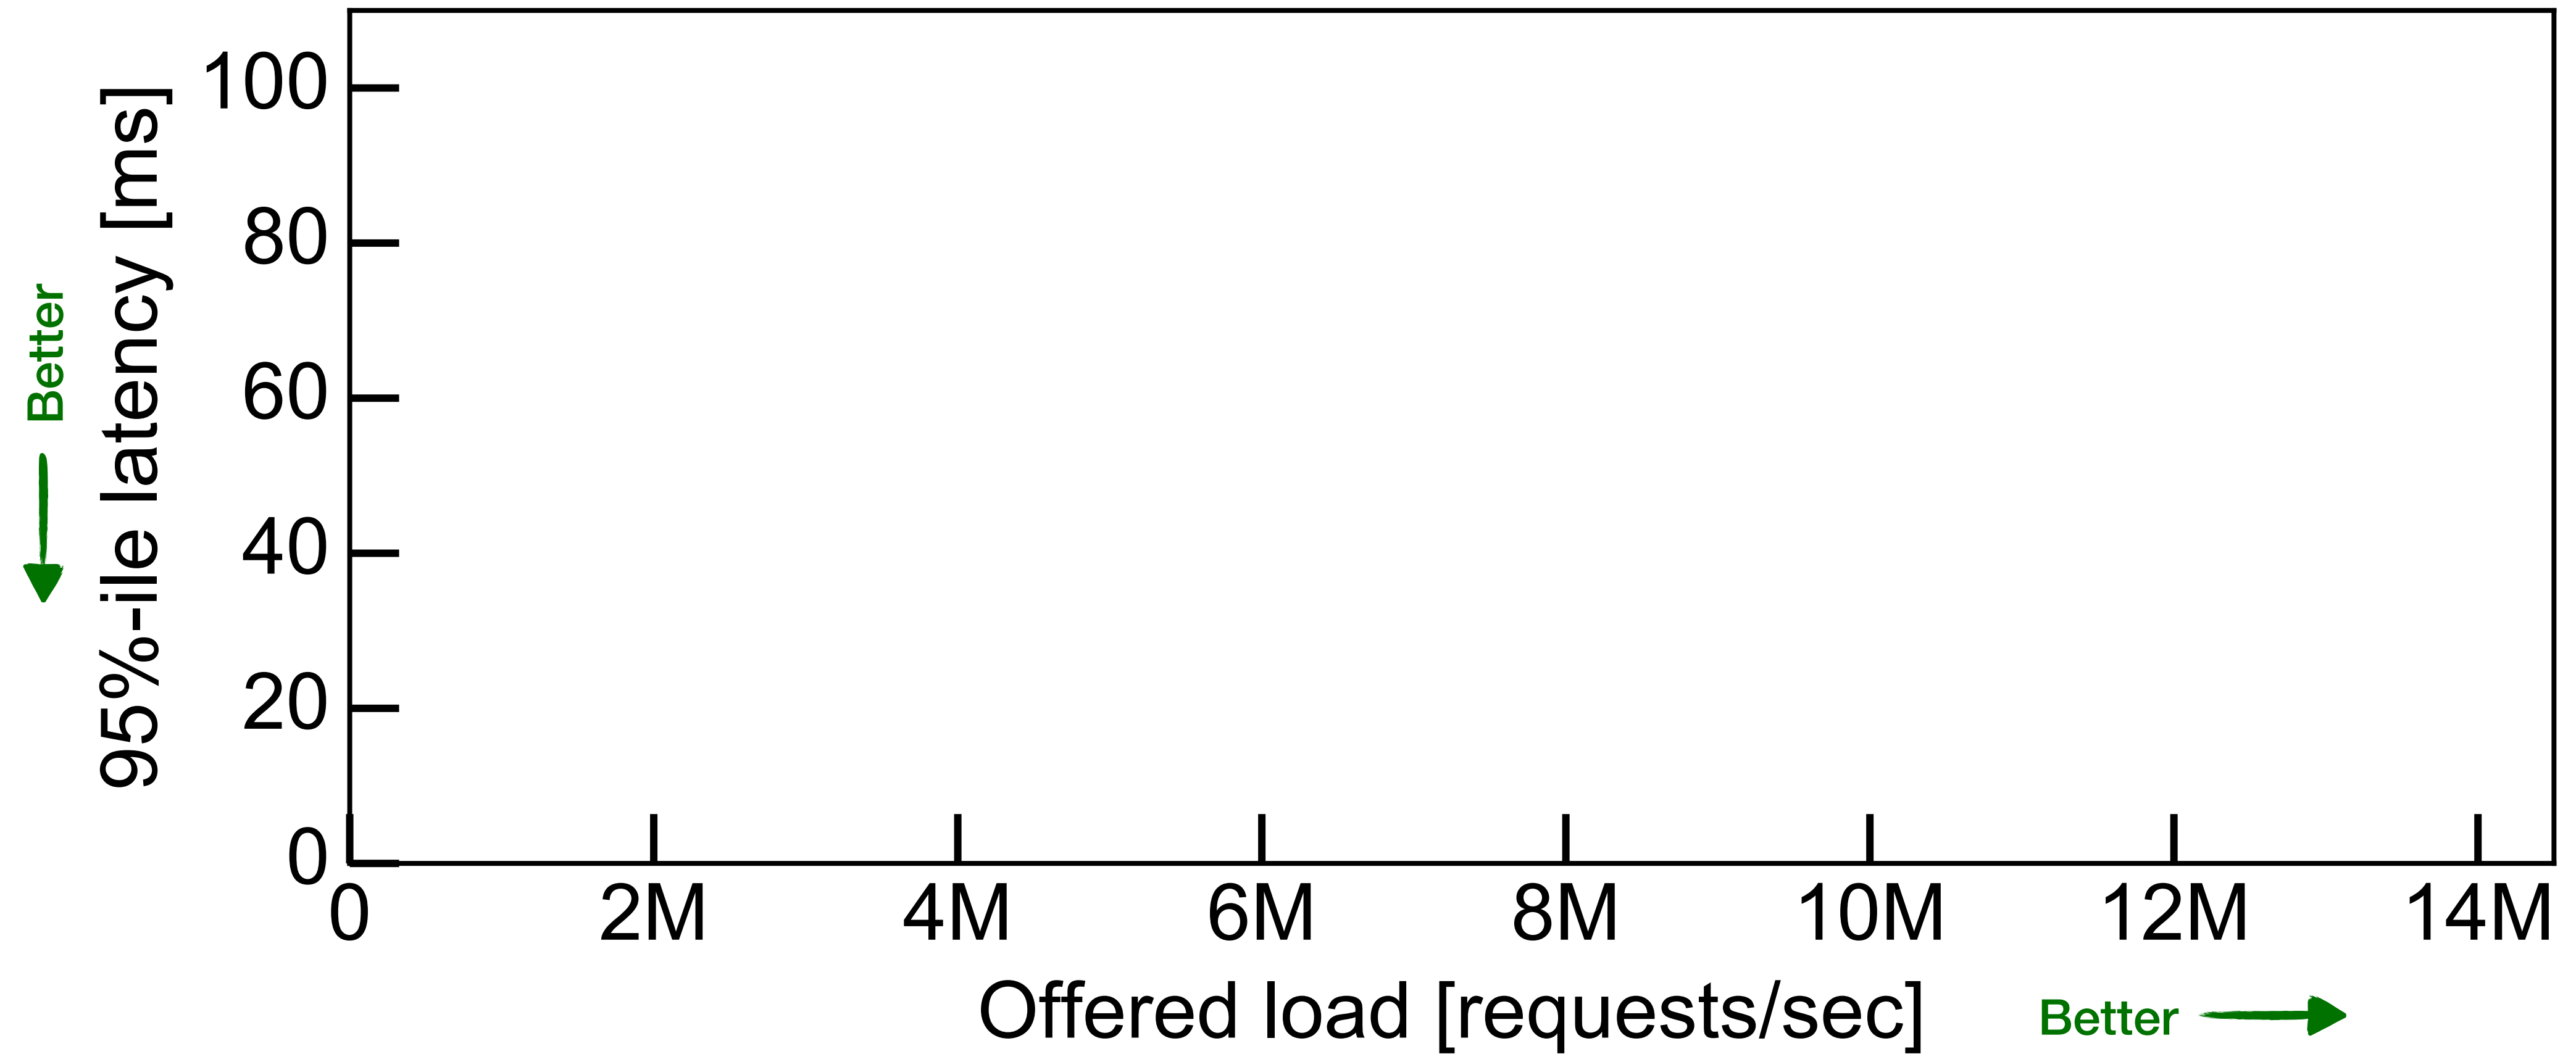


# Can Noria improve Lobsters' performance?




Noria with **natural queries** supports **5x** MySQL's throughput.

# How does Noria compare to alternatives?




# How does Noria compare to alternatives?




- ▶ Zipf-distributed story ID, 95% reads, 5% writes
- ▶ No TX, all in-memory

# How does Noria compare to alternatives?




- ▶ Zipf-distributed story ID, 95% reads, 5% writes
- ▶ No TX, all in-memory

# How does Noria compare to alternatives?




- ▶ Zipf-distributed story ID, 95% reads, 5% writes
- ▶ No TX, all in-memory

# How does Noria compare to alternatives?



- ▶ Zipf-distributed story ID, 95% reads, 5% writes
- ▶ No TX, all in-memory

# How does Noria compare to alternatives?

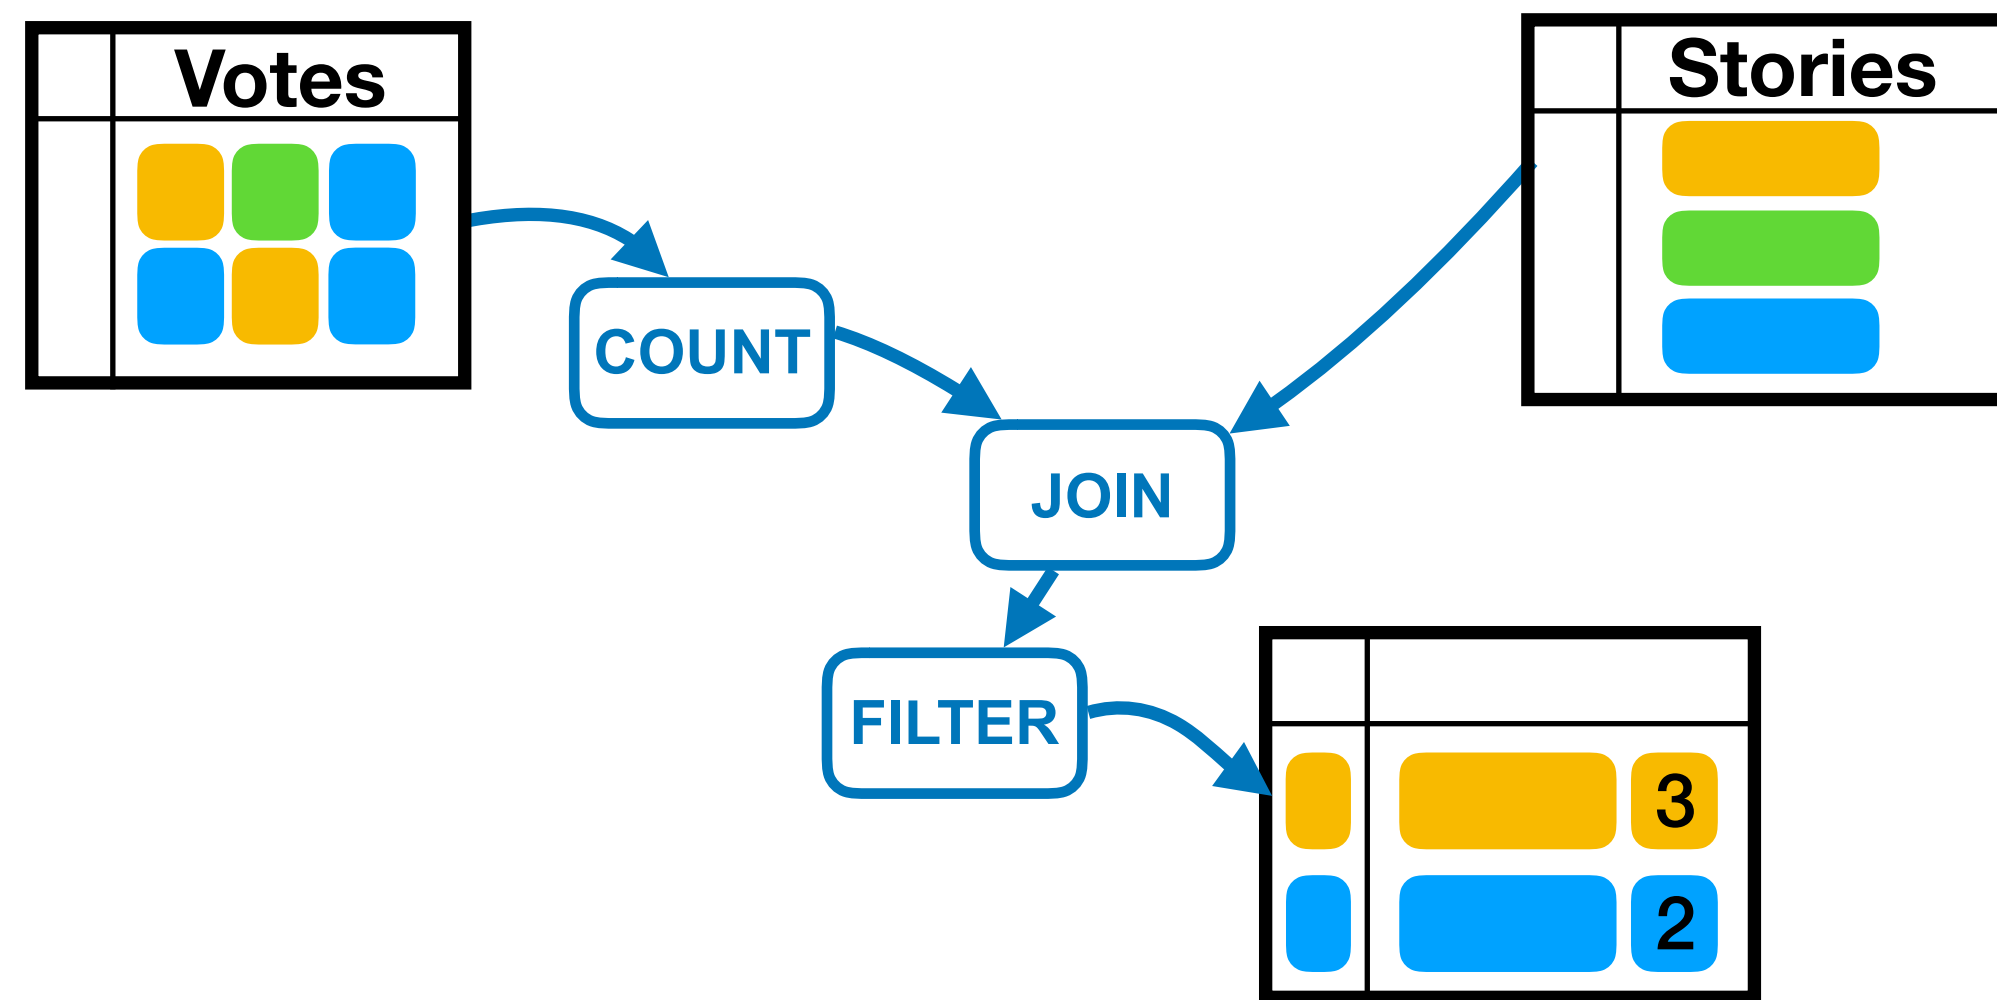


- ▶ Zipf-distributed story ID, 95% reads, 5% writes
- ▶ No TX, all in-memory

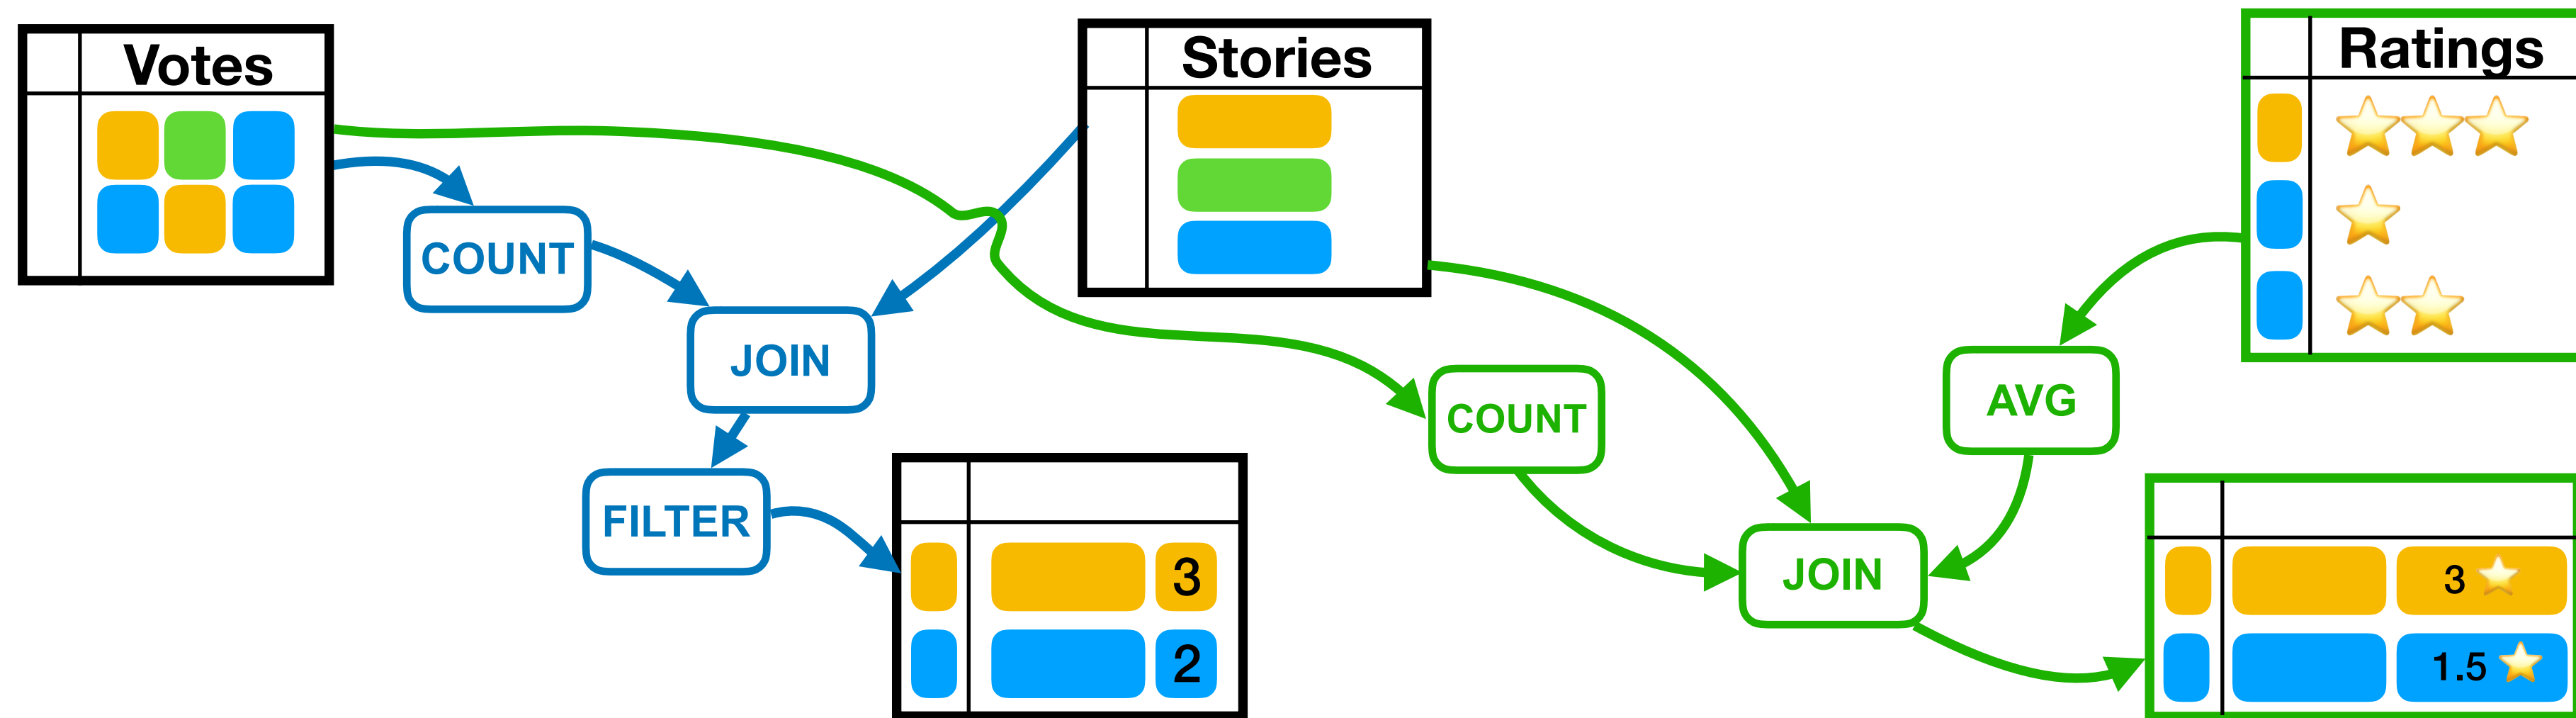
Noria outperforms an **in-memory key-value store** and simplifies its interface.



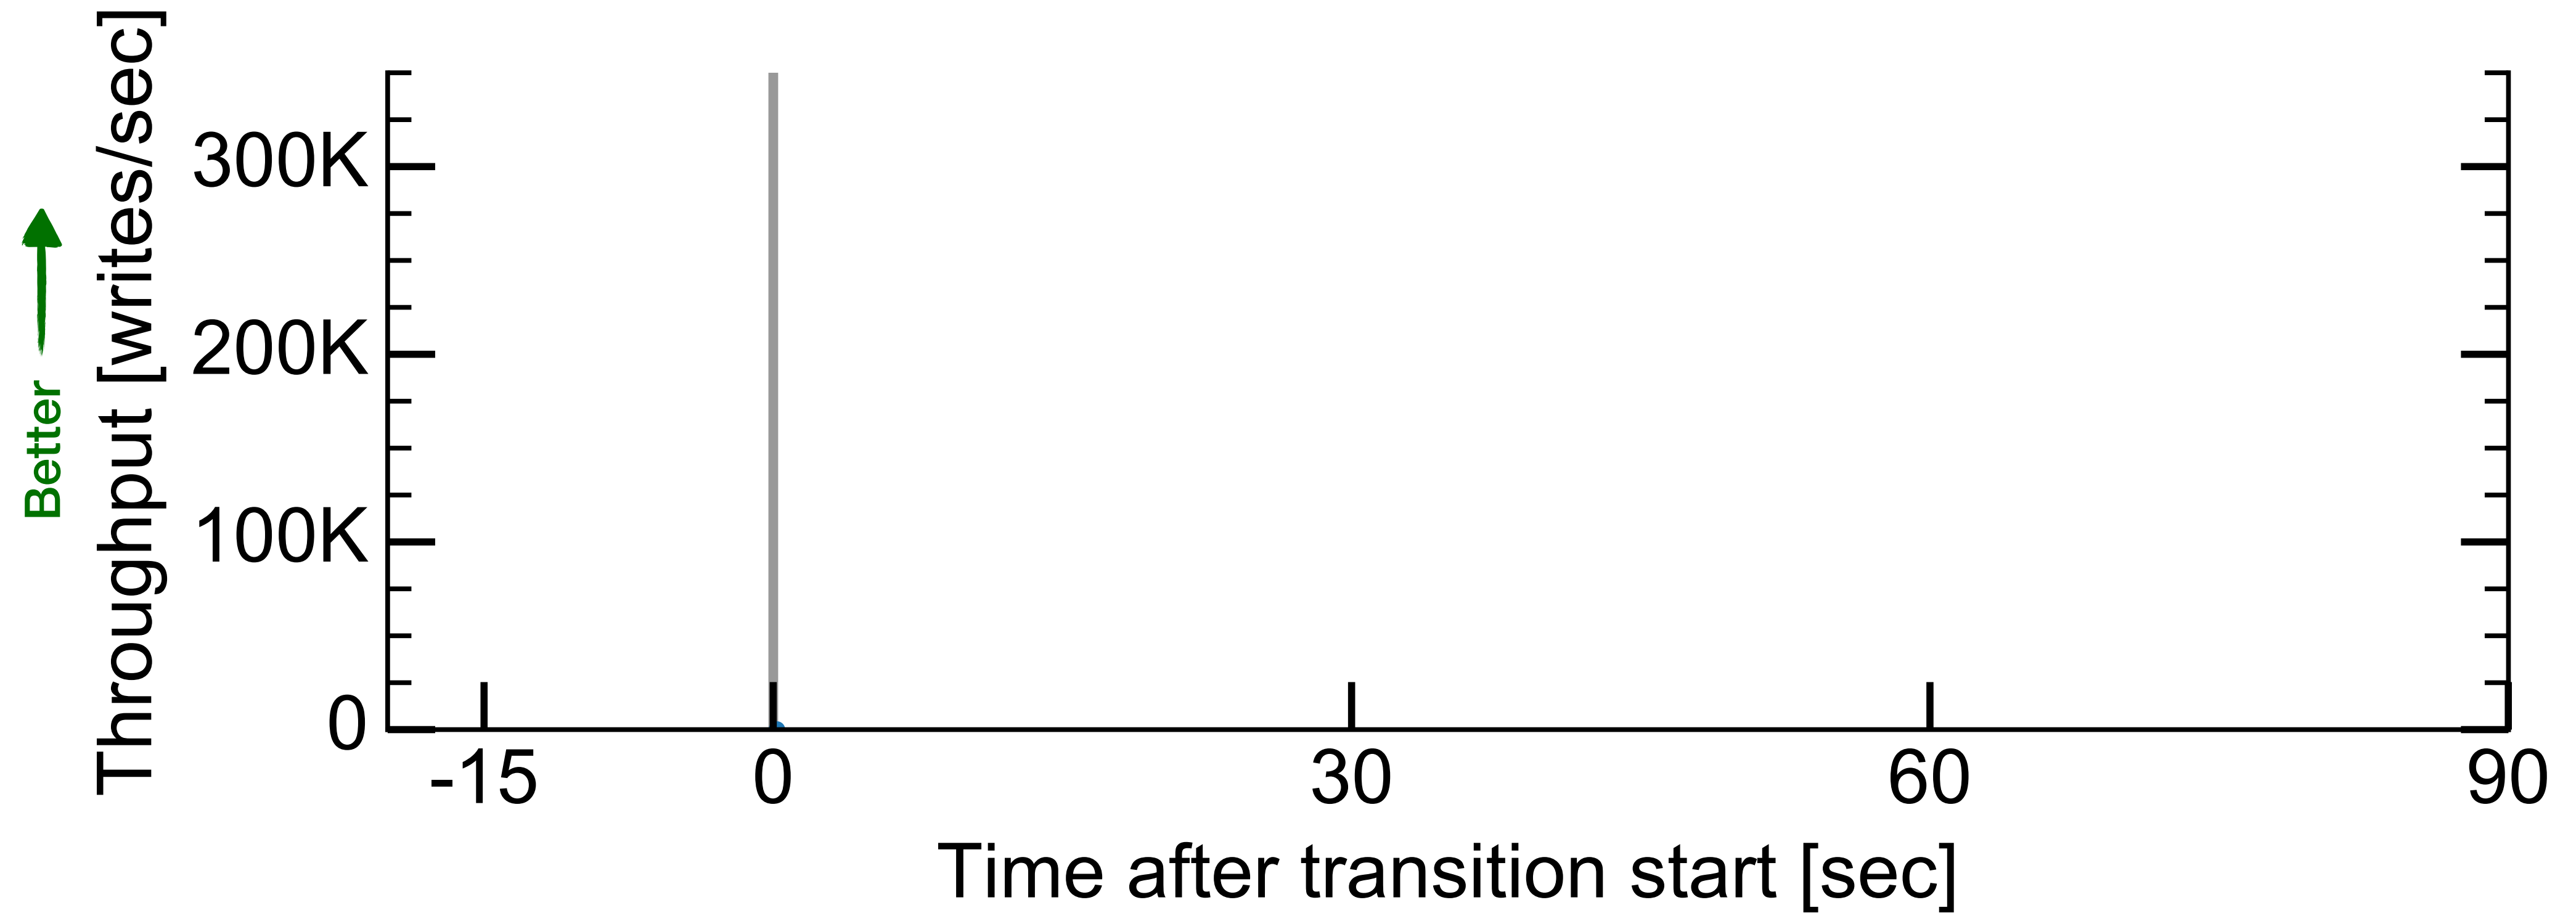
# Can Noria change queries without downtime?




# Can Noria change queries without downtime?




# Can Noria change queries without downtime?




# Can Noria change queries without downtime?



# Can Noria change queries without downtime?




# Can Noria change queries without downtime?



↑  
new table & query added


► Zipf-distributed story ID, 95% reads;  
2M existing votes at transition

# Can Noria change queries without downtime?



- ▶ Zipf-distributed story ID, 95% reads; 2M existing votes at transition
- ▶ Old view reads are live throughout

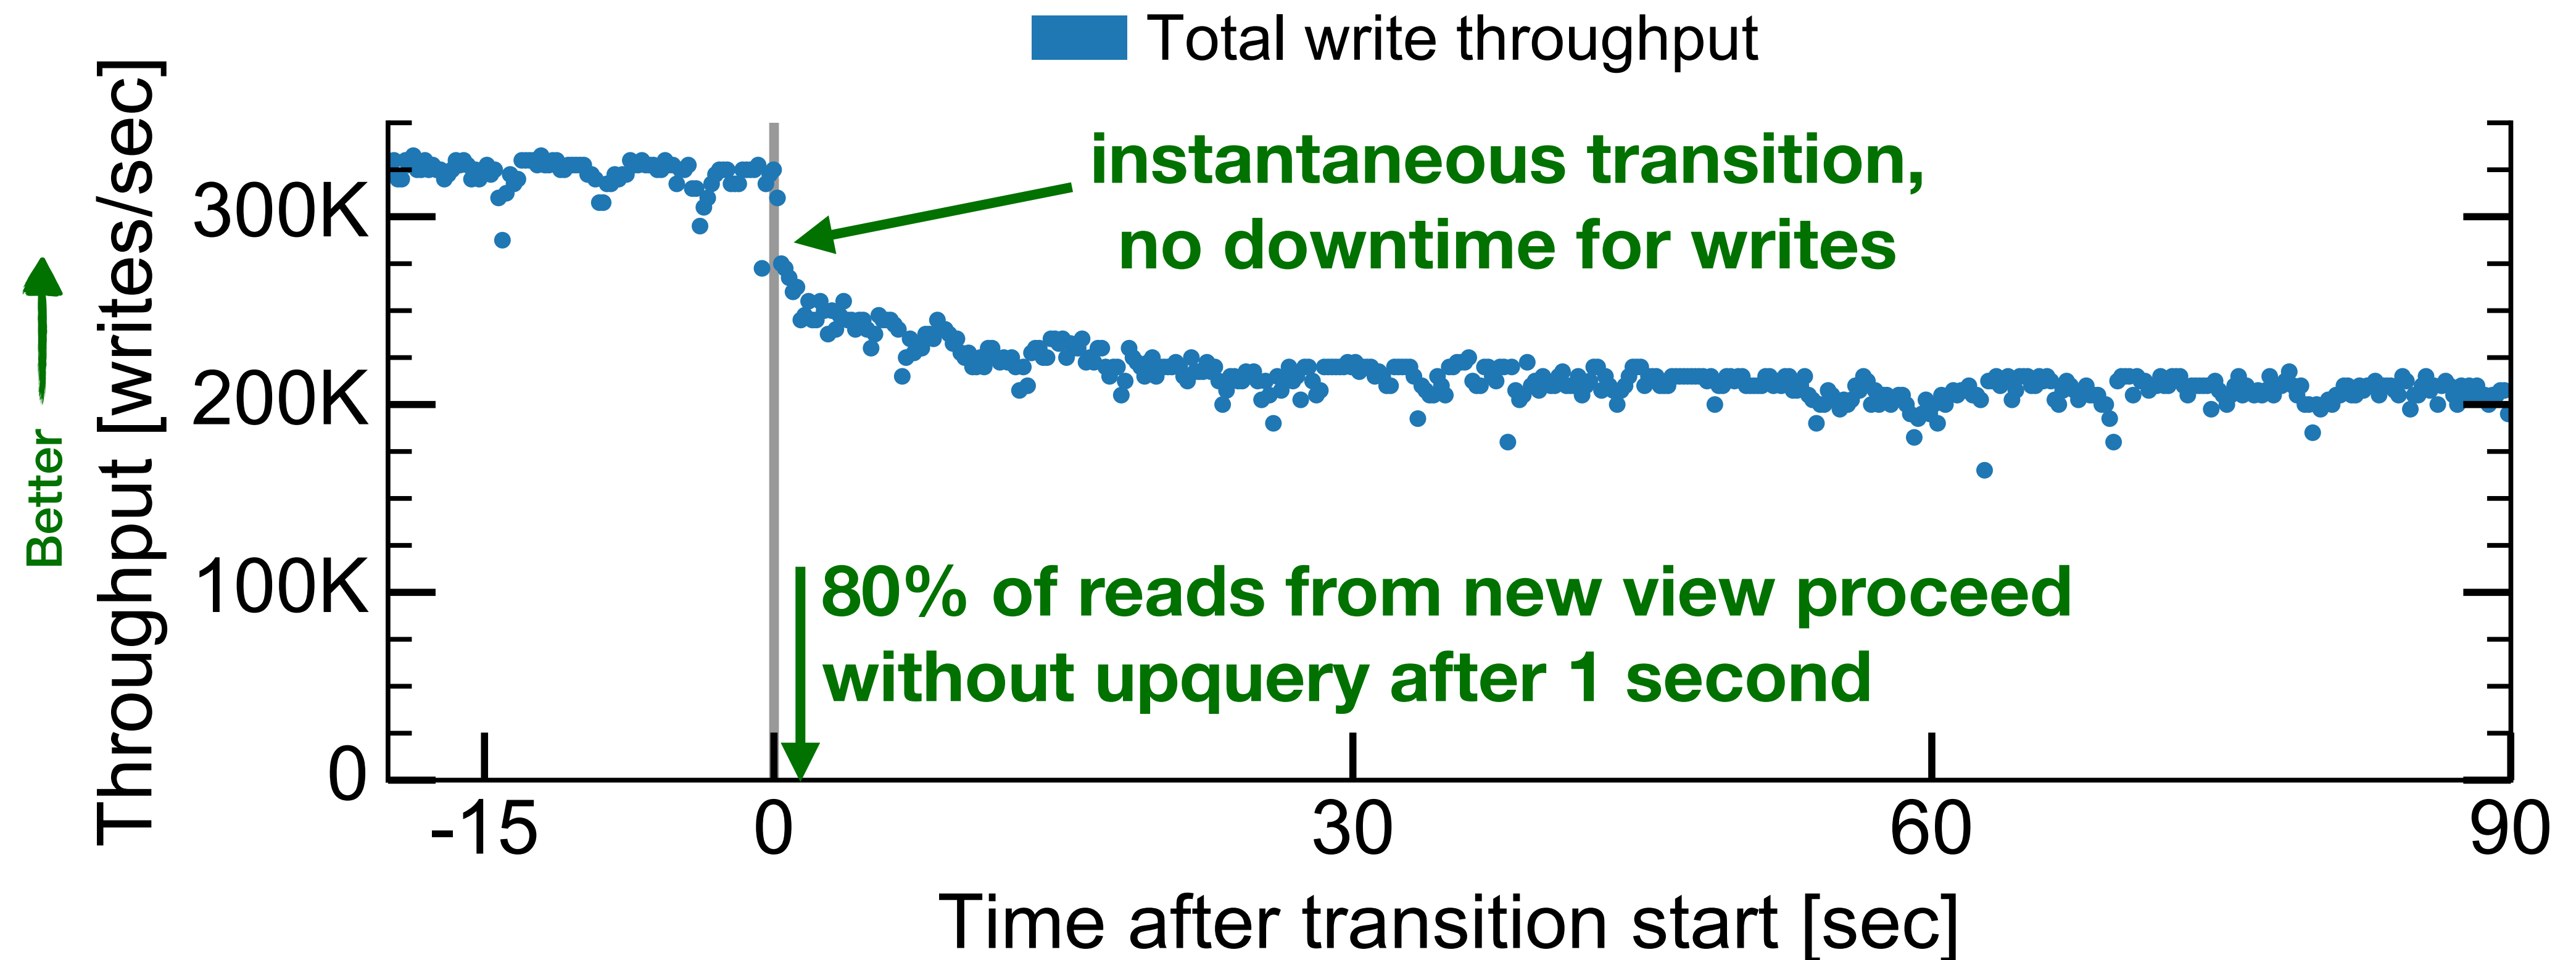
# Can Noria change queries without downtime?



- ▶ Zipf-distributed story ID, 95% reads; 2M existing votes at transition
- ▶ Old view reads are live throughout




# Can Noria change queries without downtime?



- ▶ Zipf-distributed story ID, 95% reads; 2M existing votes at transition
- ▶ Old view reads are live throughout

# Can Noria change queries without downtime?



**Noria achieves downtime-free query change with partial state.**

- ▶ Zipf-distributed story ID, 95% reads; 2M existing votes at transition
- ▶ Old view reads are live throughout

# Noria – Summary

- New partially-stateful data-flow model.
- Noria: new web application backend based on data-flow.
- Partial state saves space and allows live change.
- Supports high throughput on one or more machines.
- Open source, try it out!

# Noria – Summary

- New partially-stateful data-flow model.
- Noria: new web application backend based on data-flow.
- Partial state saves space and allows live change.
- Supports high throughput on one or more machines.
- Open source, try it out!

<https://pdos.csail.mit.edu/noria>

(see our demo at poster #37 today!)