

O'REILLY®

# Software Architecture

ENGINEERING THE FUTURE OF SOFTWARE

[softwarearchitecturecon.com](http://softwarearchitecturecon.com)

[#oreillysacon](https://twitter.com/oreillysacon)


# From CRUD to Event Sourcing an Investible Stock Universe

O'Reilly Software Architecture Conference

March 18, 2015

[#oreillysacon](https://twitter.com/oreillysacon)

# TIM Group



**Marc Siegel**

Team Lead ([@ms\\_yd](#))



**Brian Roberts**

Senior Developer / Team Lead ([@flicker](#))

# What's in it for you?

- Answer the Unanswerable
  - both now and future

# Vocabulary



1. Domain-Driven Design (DDD)
  - Event
  - Command
  - Aggregate
  
2. Event Sourcing (ES)
  - Projection
  - Read Model

# How does Event Sourcing Work?

Debit Herrn Heinrich Lomer

1876	Januar	1	sh. Cassa	1135
1876	August	1	sh. Cassa	1231 50
				<hr/>
1877	Januar	1	sh. Saldo	935 50
1877	November	2	sh. Waare bei Cassa	522 50
1877	21	sh. Waare bei Cassa	642 50	
1877	Dezember	31	sh. Linsen	33 50
				<hr/>
1878	Januar	1	sh. Saldo	2153 50
1878	Januar	1	sh. Waare bei Cassa	1156 25
1878	Dezember	31	sh. Linsen	3383 50
				<hr/>
1878	Dezember	31	sh. Linsen	191 25
				<hr/>
1879	Januar	1	sh. Saldo	4831 16
1879	Dezember	31	sh. Waare bei Cassa	3031 50
1879	Dezember	31	sh. Linsen	224 35
				<hr/>
1879	Dezember	31	sh. Linsen	242
				<hr/>
1880	Januar	1	sh. Saldo	8591 45
1880	April	1	sh. Linsen	569 45
				<hr/>
1880	April	1	sh. Linsen	11 25

Leipzig

Credit

1876	Januar	1	Saldo von 1875	2406
1876	August	1	sh. Linsen	11
				<hr/>
1877	November	30	sh. Cassa	992
1877	Dezember	31	sh. Linsen	4 50
1877	31	sh. Saldo	1186 30	
				<hr/>
1877	Dezember	31	sh. Saldo	2133 50
				<hr/>
1878	April	22	sh. Cassa	502 50
1878	Dezember	4	sh. Cassa	502 50
1878	31	sh. Linsen	36 50	
1878	Dezember	31	sh. Saldo	581 10
				<hr/>
1878	Dezember	31	sh. Saldo	4831 50
				<hr/>
1879	Januar	3	sh. Cassa	506
1879	11	sh. Cassa	54	
1879	25	sh. Cassa	999 50	
1879	Dezember	25	sh. Cassa	505
1879	Dezember	31	sh. Cassa von geb. Linsen	110 50
1879	"	sh. Linsen	569 45	
1879	"	sh. Saldo	3528 40	
				<hr/>
1880	April	1	sh. Cassa	551

# How does Event Sourcing Work?

(Quotes from Greg Young)

"State transitions are an important part of our problem space and should be modeled within our domain"

# How does Event Sourcing Work?

(Quotes from Greg Young)

"State transitions are an important part of our problem space and should be modeled within our domain"

"Event Sourcing says all state is transient and you only store facts."



# Vocabulary




## 1. Domain-Driven Design (DDD)

- **Event:** *something that happened in the past; a fact; a state transition*
- Command
- Aggregate

## 2. Event Sourcing (ES)


- Projection
- Read Model

# Lifecycle of a Listing (Simplified)




Q: How to represent the meaningful state transitions of the domain as events?


# Listing Lifecycle Events



## Events:




# Determine Current State



Q: How do we determine the state as of a given year? Or as of today?

# Determine Current State




Q: How do we determine the state as of a given year? Or as of today?

A: "When we talk about Event Sourcing, current state is a left-fold of previous behaviors"  
-- Greg Young

# Current State is a Left Fold of Events $\lambda$


- FP: Left Fold aggregates a collection via a function and an initial value
  - Ruby: `[1, 2, 3].inject(0, :+) == 6 # symbol fn name`
  - Scala: `List(1, 2, 3).foldLeft(0)(_ + _) == 6 // anon function`
- Provide an initial state  $s_0$  and a function  $f: (S, E) \Rightarrow S$
- Current State after event  $e_3$  is:
  - $= \text{leftFold}( [e_1, e_2, e_3], s_0, f )$
  - $= f( f( f( s_0, e_1), e_2), e_3 )$

# Replay to Earlier State: 1998



State in 1998 =  $\left[ \right]$ .apply(  
    **TickerListed**  
    ticker: MSFT  
    date: 1986  
)

# Replay to Earlier State: 1998




State in 1998 =  $\left[ \begin{array}{c} \text{MSFT} \\ \text{Listed} \end{array} \right]$






# Replay to Earlier State: 2000




State in 2000 = [ MSFT Listed IPET Listed ]

# Replay to Earlier State: 2001




State in 2001 =  $\left[ \begin{array}{|c|} \hline \text{MSFT} \\ \hline \text{Listed} \\ \hline \end{array} \quad \begin{array}{|c|} \hline \text{IPET} \\ \hline \text{Listed} \\ \hline \end{array} \right] .\text{apply}(\text{TickerDelisted ticker: IPET date: 2000})$

# Replay to Earlier State: 2001




State in 2001 = [ MSFT Listed IPET Delisted ]

# Replay to Earlier State: 2012




State in 2012 =  $\left[ \begin{array}{|c|} \hline \text{MSFT} \\ \hline \text{Listed} \\ \hline \end{array} \quad \begin{array}{|c|} \hline \text{IPET} \\ \hline \text{Delisted} \\ \hline \end{array} \right] .\text{apply}(\text{TickerListed ticker: FB date: 2012})$

# Replay to Earlier State: 2012




State in 2012 = [ MSFT Listed IPET Delisted FB Listed ]

# Review - Only the Events are Stored



## Events:



# Potential Benefits

- Answer the unanswerable (via history replay)
- Debugging of historical states deterministically (via history replay)
- Never Lose Information (write-only store)
- Edit the Past (via new events effective at older times)
- Optimize reads (purpose-built read models)
- Enhanced Analytics (analyze all history as it occurred)



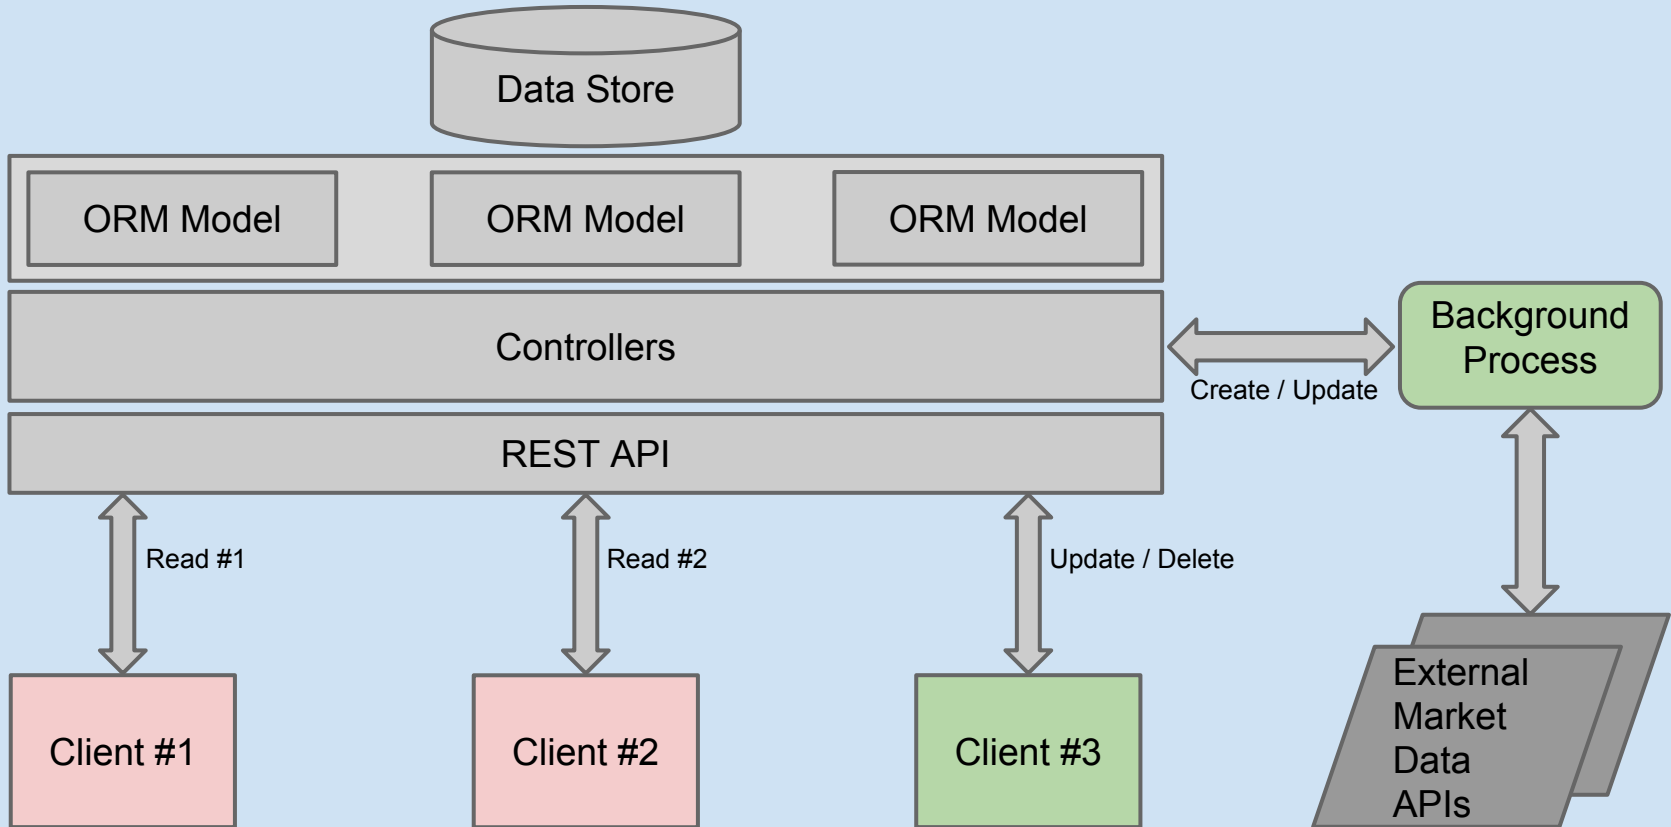
# Potential Drawbacks

- Eventual Consistency
- No built-in querying of domain models (SELECT name WHERE ...)
- Risks of using a new architectural pattern
- Lack of agreement on tools and infrastructure
- Increased storage requirements

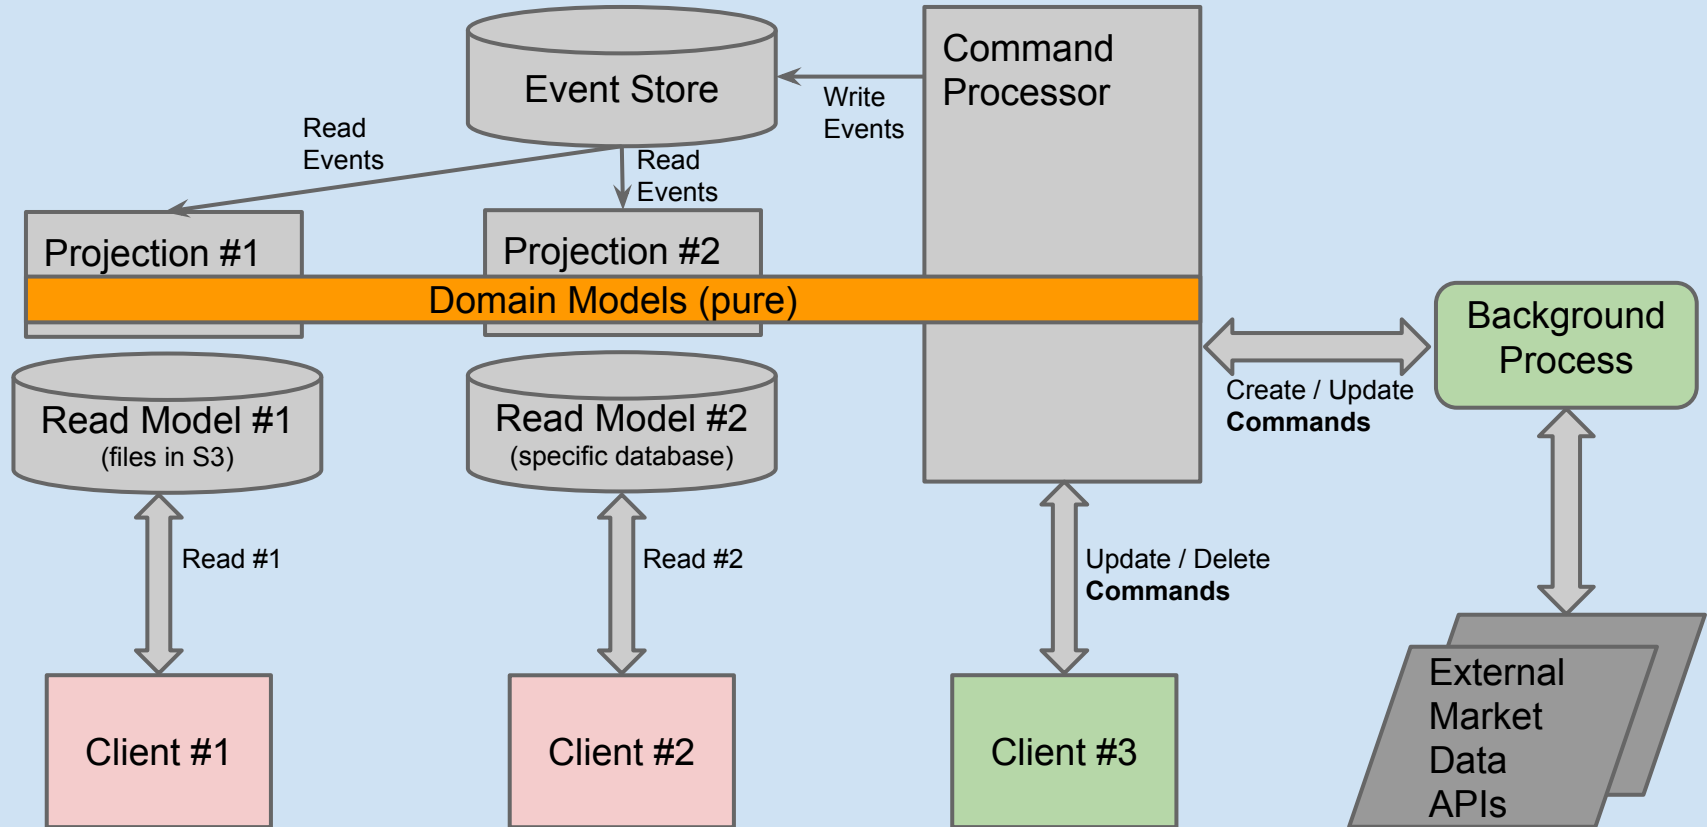
**What is different from CRUD?**



# CRUD Micro-Service



# Event-sourced Micro-Service



# Vocabulary



## 1. Domain-Driven Design (DDD)

- Event: *something that happened in the past; a fact; a state transition*
- **Command:** *a request for a state transition in the domain*
- Aggregate

## 2. Event Sourcing (ES)

- Projection
- Read Model

# Aside: Domain Model is Pure?



- FP: A "pure" function doesn't cause any side effects
  - No reads or writes that modify the world
  - No altering a mutable data structure
  - Substitute  $f(x)$  for its result without changing meaning of program
- An event-sourced domain can be two pure functions
  - `process(currentState, command) => e1, e2, e3`
  - `apply (currentState, event) => nextState`
- Separates the logic of the model from interactions with any changing state in the world

# What is different from CRUD?

(Quotes from Greg Young)

"The model that a client needs for the data in a distributed system is screen-based and different than the domain model."



# Trade-offs

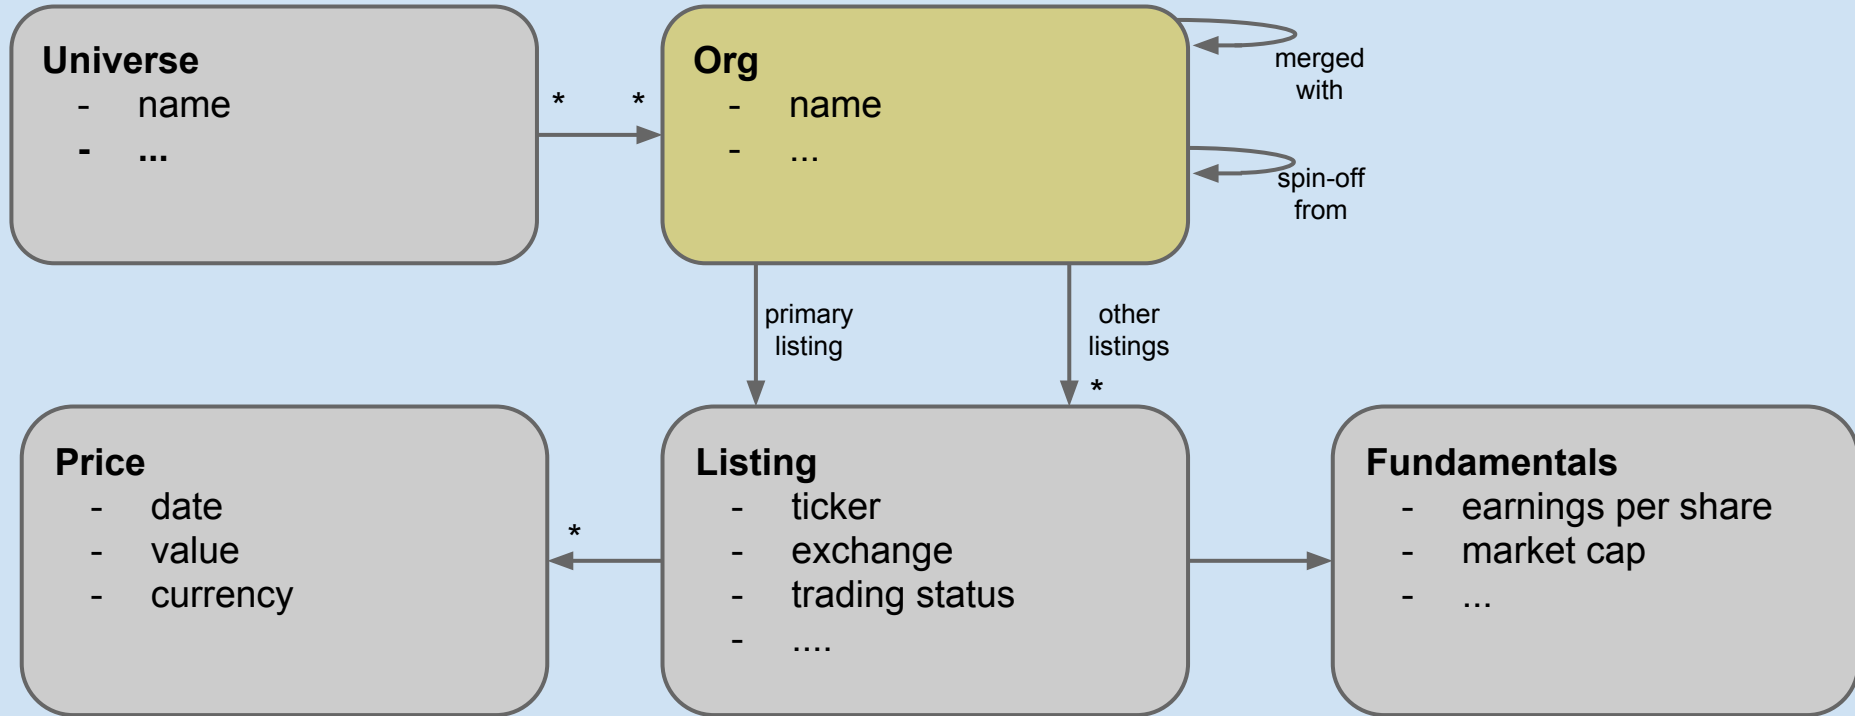
- ACID vs. Eventual Consistency
  - CRUD w/ ACID database
    - Once a row is written, subsequent reads reflect it
    - But: no help with domain-level consistency!
  - Event Sourced
    - Once event is written, subsequent reads reflect it
    - Projections eventually consistent
- Up-front costs
  - Domain modeling is hard!

# CRUD Models of Market Data

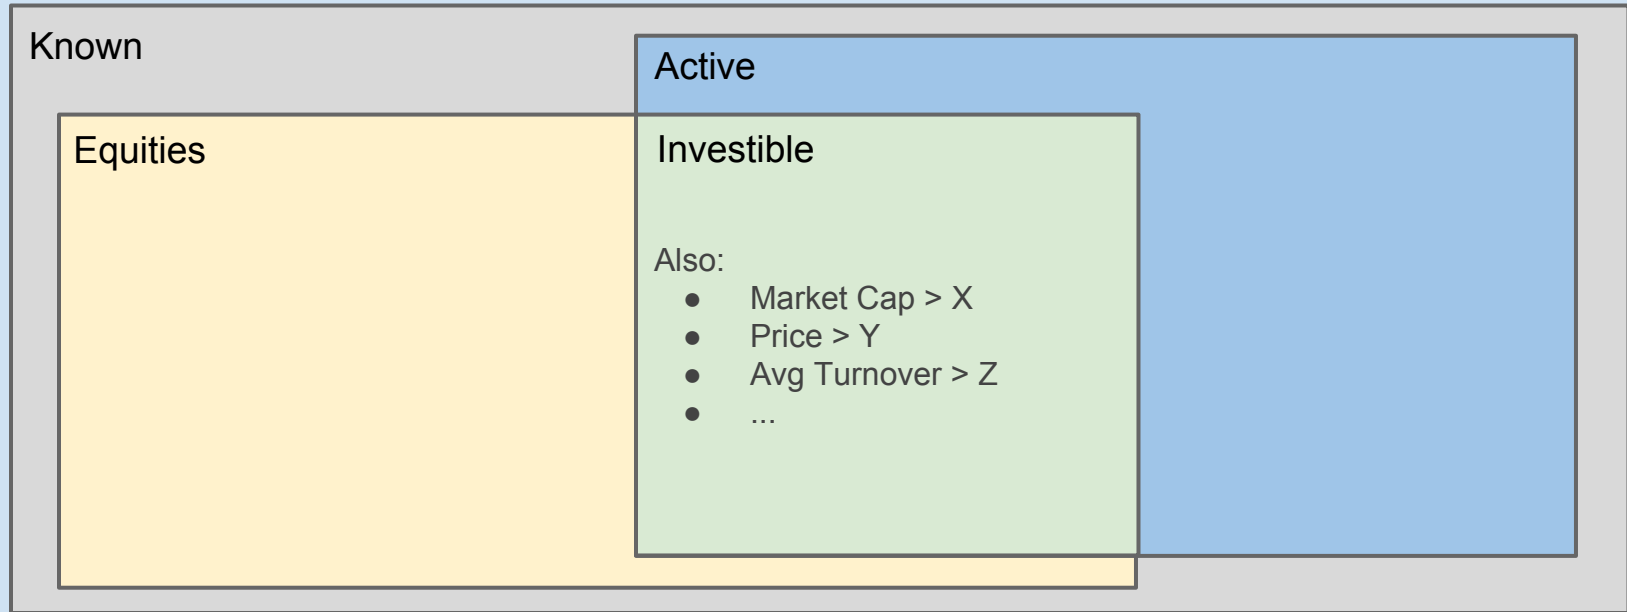
← 1-9  
WALL ST



# CRUD Models of Market Data

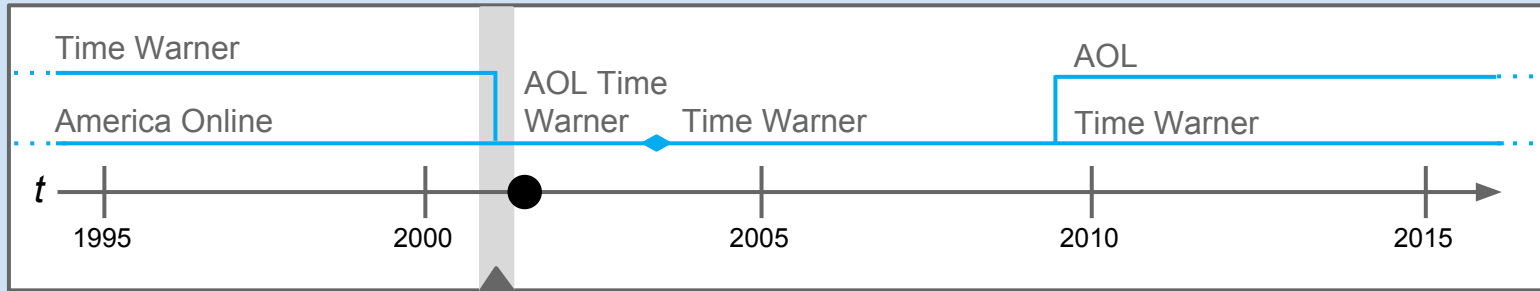


# Answerable?



Q: Was AOL in **Investible Universe** in 1995?

# 2001 - Merger of AOL/Time Warner



**Org**  
name: Time Warner

**Listing**  
ticker: TWX  
trading status: Listed  
....

**Org**  
name: America Online

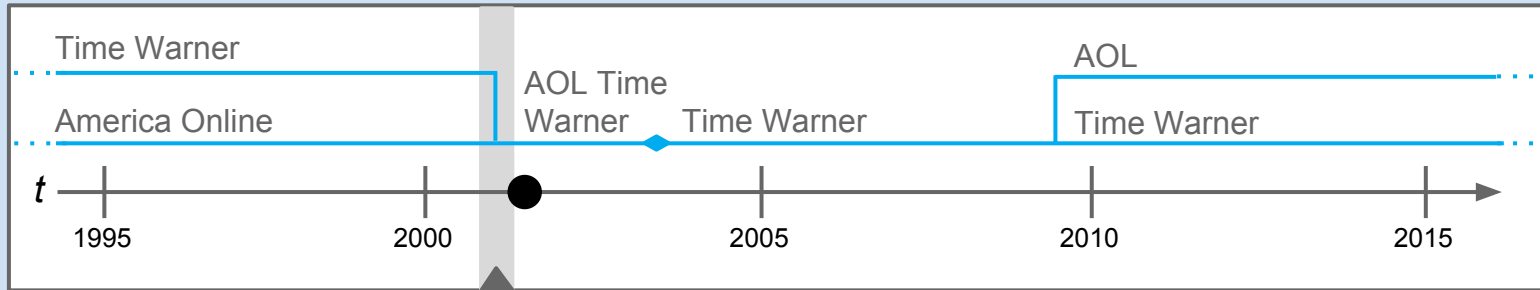
**Listing**  
ticker: AOL  
trading status: Listed  
....

```
Orgs  
.findByName("Time Warner")  
.mergeWith("America Online")
```

```
Orgs  
.findByName("America Online")  
.setName("AOL Time Warner")  
.addListing("TWX")
```

```
Listings  
.findByTicker("TWX")  
.setTradingStatus("Delisted")
```

# 2001 - Merger of AOL/Time Warner




```
Orgs  
.findByName("Time Warner")  
.mergeWith("America Online")
```

```
Orgs  
.findByName("America Online")  
.setName("AOL Time Warner")  
.addListing("TWX")
```

```
Listings  
.findByTicker("TWX")  
.setTradingStatus("Delisted")
```

# 2003 - Name / Ticker Change



**Org**  
name: Time Warner

merged with

**Org**  
name: AOL Time Warner

```
Org  
.findByName("AOL Time Warner")  
.setName("Time Warner")
```

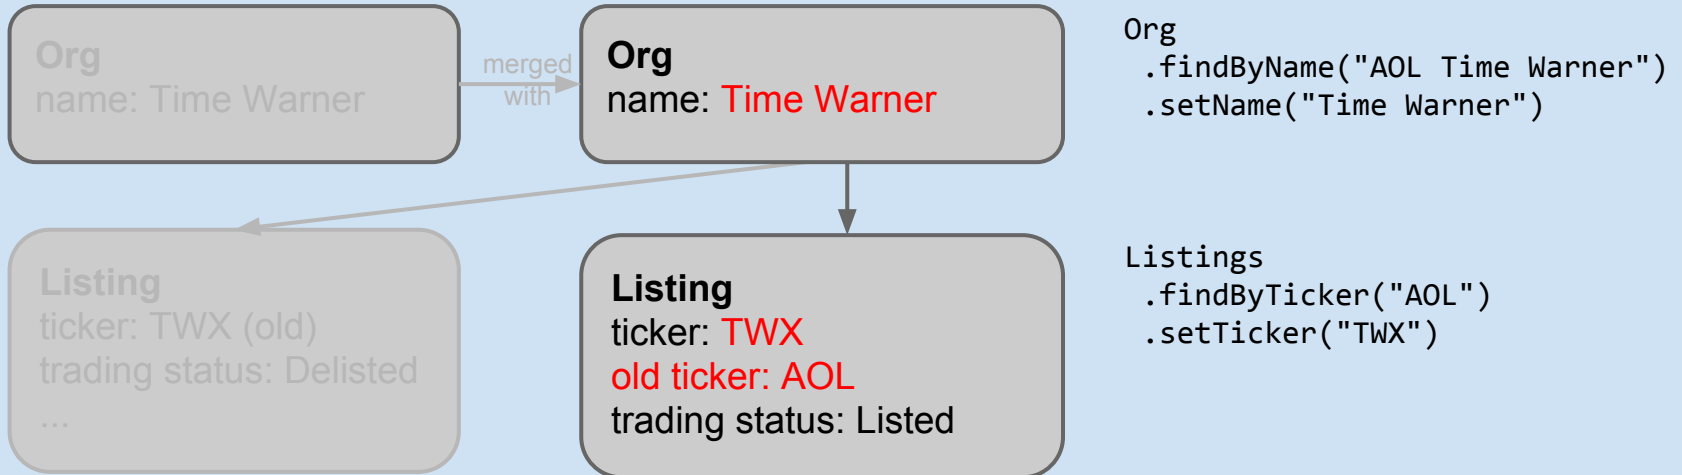

**Listing**  
ticker: TWX (old)  
trading status: Delisted  
...

**Listing**  
ticker: AOL  
trading status: Listed  
....

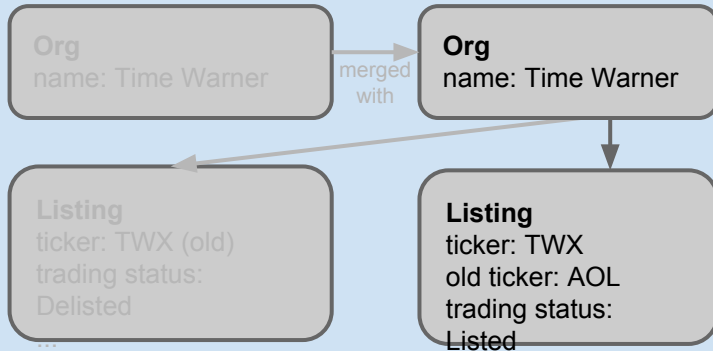

```
Listings  
.findByTicker("AOL")  
.setTicker("TWX")
```



# 2003 - Name / Ticker Change



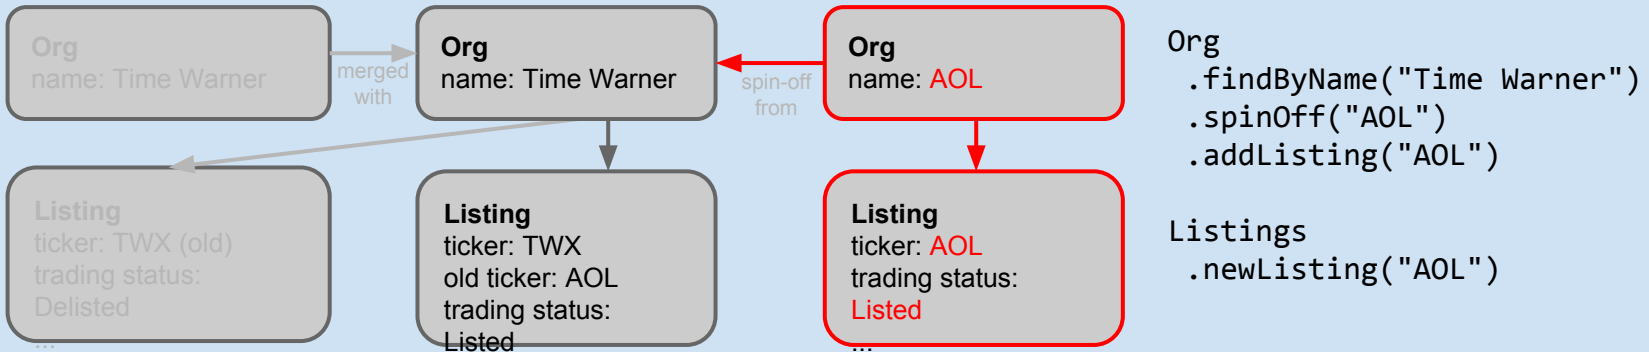

# 2009 - AOL Spinoff



```
Org  
  .findByName("Time Warner")  
  .spinOff("AOL")  
  .addListing("AOL")
```

```
Listings  
  .newListing("AOL")
```

# 2009 - AOL Spinoff



# Answerable with CRUD Models?

Q: Was AOL in **Investible Universe** in 1995?

A: No **and** Yes

- **No** current AOL org (didn't exist at time)
- **Yes** former America Online (now Time Warner)

Complexity in query, requires previous states

- Query against version columns with date ranges?
- Query against previous versions tables?

# CRUD Models - How to Update?

- Update price of TWX on 1995-01-03
  - Original: ~~\$56.22~~ Correct: \$52.62
- Complexity in query
  - Wrong: `Listings.findByTicker("TWX")` // this is AOL!
  - Right: Complex historical query...
- Complexity in update
  - Org Primary Listing - any change?
  - Org Universe membership - any changes?
  - Support *Two-Dimensional Time* aka *as-of* query?


# Problems



# Main Problem

Unanswerable questions

- Time travel intractable
- Past not always reproducible





# More Problems

- **Correctness**
  - Divergent interpretations of data
- **Availability**
  - How often can data be unavailable to clients?
- **Performance**
  - How fast must operations complete?
- **Determinism**
  - Reproducing prior states for reporting, debugging, etc.
- **Auditability**
  - Who changed what when and why?

# Problems - Correctness

- Need a new definition of e.g. adjusted price
  - **Old:** unadjusted \* splits \* spin-offs
  - **New:** unadjusted \* splits \* spin-offs \* dividends
- But...
  - Some client systems still need the old definition
  - CRUD data store didn't store the individual factors
- Common Solutions
  - Add past to relational model? Reprocess?

# Problems - Availability

- How long can data be unavailable?
  - **Not long** - End-user client
  - **Hours to Days** - Reporting
- But...
  - Most stringent of client requirements applies to all
  - Cascading failures: unavailability propagates
- Common solutions
  - bulk-heading, circuit-breakers, more servers
  - more complex than necessary?

# Problems - Performance

- How fast must operations complete?
  - Writes need to keep up with input
  - Reads have varying requirements
- But..
  - Due to contention on Shared Mutable State, badly performing Reads can impact everything else
- Common Solutions
  - Caching, sharding, more server resources
  - Trade-off with ACID consistency

# Problems - Determinism

- Reproducing prior states
  - Reporting Consistently on a Past period
    - Apply adjustments **only** to end of the period
  - Debugging
    - Reproduce state of data in past
- But..
  - Not easy with Shared Mutable State!
- Common Solutions
  - Versioned rows, audit tables, database snapshots

# Problems - Auditability

- Why did data change?
  - Attribution (source of data)
  - Security (who did it)
  - History (what and when was previous value)
- But..
  - Not easy with Shared Mutable State!
- Common Solutions
  - Versioned rows, audit tables

# Event Sourced Models of Market Data



Creditreform

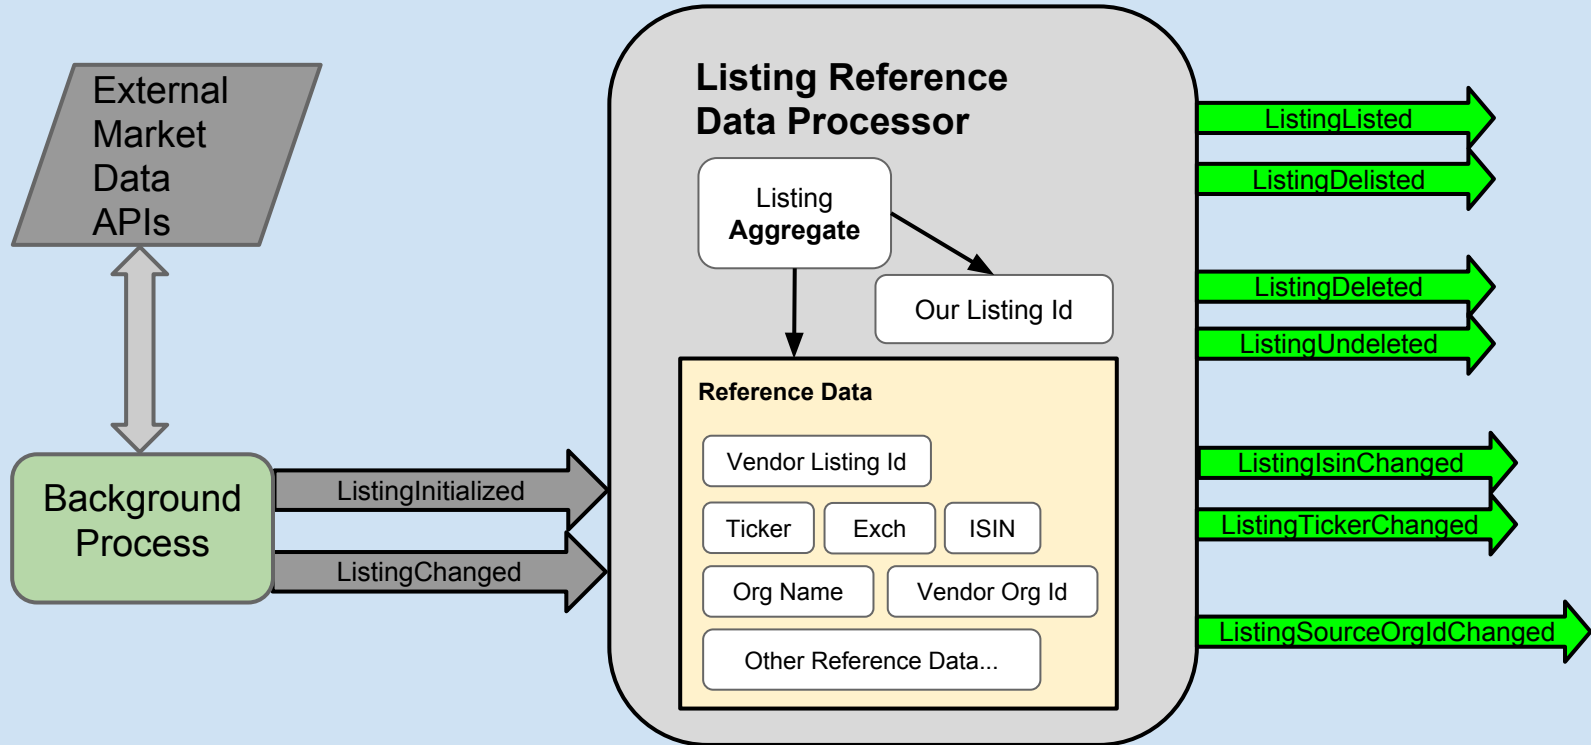
GANTZ

GANTZ

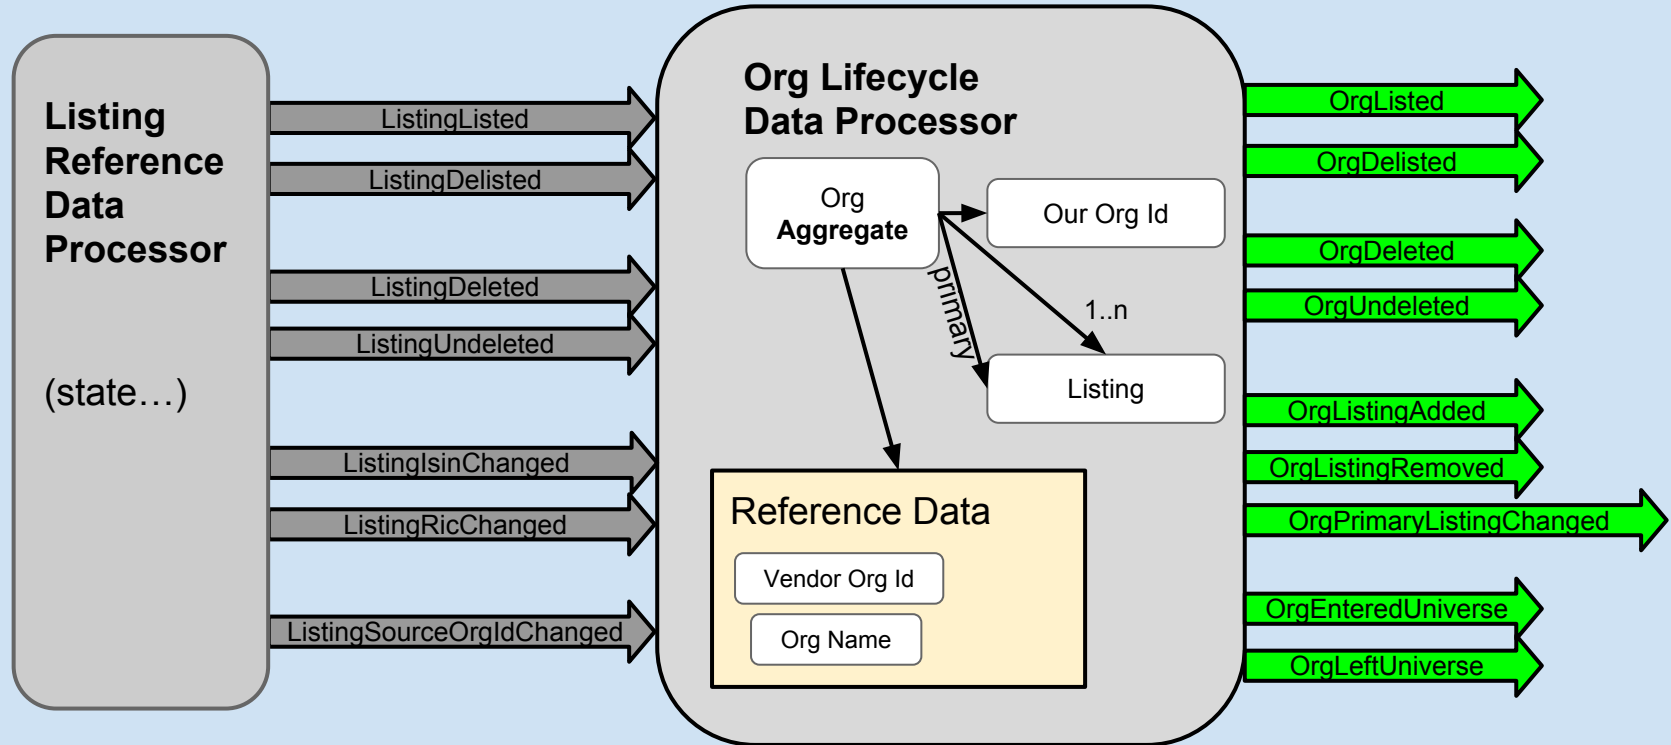
GANTZ



# Event Sourced Listings



# Event Sourced Orgs



# Vocabulary



## 1. Domain-Driven Design (DDD)

- Event: *something that happened in the past; a fact; a state transition*
- Command: *a request for a state transition in the domain*
- **Aggregate**: *domain objects in a transactionally-consistent unit*

## 2. Event Sourcing (ES)

- Projection
- Read Model

# Resulting Read Models



# Stock Universes - Read Model

- Generate as CSV files in S3, clients retrieve via REST API

GET /universes/1995-01-03

```
Org Id,Known Universe?,Active Universe?,Investible Universe?,Primary Listing Id, ... Ticker  
"49498",true,true,true,"0x00100b000b569402","US8873173038",..."AOL"
```

- Problems?
  - **Correctness:** Interpretation only for this use case
  - **Availability:** No impact on other use cases
  - **Performance:** No read-side calculations
  - **Determinism:** Can re-generate from event source data
  - **Auditability:** Available in event source data
- Consistency is at domain level -- entire history of universes in this case
  - Generate entire history to S3 bucket, API switches buckets atomically

# Answerable via Event Sourcing?

Q: Was AOL in **Investible Universe** in 1995?

```
GET /universes/1995-01-03
```

```
Org Id, Known Universe?, Active Universe?, Investible Universe?, Primary Listing Id, ..., Ticker  
"49498", true, true, true, "0x00100b000b569402", "US8873173038", ..., "AOL"
```

A: **Yes!**      former America Online (now Time Warner)

Trivial query against purpose-built Read Model

# Org History - Read Model

- Build directly from indexed event stream, clients retrieve via REST API

GET /orgs?ticker=TRW

```
[ { eventType: "ListingListed", listing_id: "1", ticker: "AOL", processedAt: "...", effectiveAt: "...", ...}
  { eventType: "OrgListed", ... },
  { eventType: "ListingAdded", listing_id: "1", ...}, ...
  { eventType: "ListingTickerChanged", listing_id: "1", old_ticker: "AOL", ticker: "TWX" ...}, ... ]
```

- Problems?
  - **Correctness:** Interpretation only for this use case
  - **Availability:** No impact on other use cases
  - **Performance:** No read-side calculations
  - **Determinism:** Reads directly from event source data
  - **Auditability:** Available in event source data
- Consistency is at domain level -- entire history of single org in this case
  - Generate entire history on-the-fly directly from source (indexed!)



# Vocabulary



## 1. Domain-Driven Design (DDD)

- **Event:** *something that happened in the past; a fact; a state transition*
- **Command:** *a request for a state transition in the domain*
- **Aggregate:** *domain objects in a transactionally-consistent unit*

## 2. Event Sourcing (ES)

- **Projection:** *to derive current state from the stream of events*
- **Read Model:** *a model of current state designed to answer a query*


# Read Model vs Cache

A **cache** is a query intermediary. It holds a previous response until invalidated, then queries again.

A **read model** does not query. Applying new events to it changes the answer it returns.

Example: count of people in room

- both may return 10 when the answer is now 11 or 9 - staleness
- cache counts the people - slow and may require locking the doors
- read model applies the entrance/exit events - like a click counter - no impact on people nor doors



A FEW CENTURIES FROM NOW, ALL THE ENGLISH OF THE PAST 400 YEARS WILL SOUND EQUALLY OLD-TIMEY AND INTERCHANGEABLE.


# Conclusions

- Answer the "unanswerable"
- Avoid impacts on other use cases
- Keep facts that may answer future questions

# Questions?



Follow up later



**Marc Siegel**

[@ms\\_yd](#)



**Brian Roberts**

[@flicker](#)



"So long,

and thanks for  
all the fish."

# Events vs Audit Tables or Versioned Rows

An **audit table** or **row version column** use shared mutable state as the source of truth, and additionally store some history. Inconsistencies are hard to fix, edits of the past are challenging, new use cases can be challenging.

An **event store** is an append-only list of immutable facts. What has occurred is recorded, and can be replayed to interpret according to a future use case.

## Example: count of people in room

- audit table is a logbook in the room - query may be complex, facts may not be consistent, depends on keeping it up to date
- versioned rows is a logbook carried by each person - same issues as audit table
- event store is "just the facts", only interpretation changes