



Leipziger Forschungszentrum  
für Zivilisationserkrankungen

# Datomic Fast at Scale?

Alexander Kiel  
[alexanderkiel@gmx.net](mailto:alexanderkiel@gmx.net)  
[@alexander\\_kiel](#)

Universität Leipzig

# Motivation

Can a functional database used in projects  
beyond toy examples?

# LIFE Study

- Cohort Study in Leipzig
- 10.000 Adults
- 10.000 Childs with Parents
- 70 Assessments per Participant
- Central data management and storage

# Central Research DB

- Currently relational Oracle DB
- 5 GB Data in 600 Tables
- 2 Million Rows / 40.000 Spalten
- 36 Million Non-Null Data Points

# The Problem

How many participants do we have for a particular research question?

“Von wie vielen Probanden haben wir eine Soziodemographie und ein Großes Blutbild in der Erwachsenenstudie?”

# i2b2

The screenshot shows the i2b2 Query & Analysis Tool web interface. The browser address bar displays "smarti2b2.org/webclient/#". The page title is "i2b2 Query & Analysis Tool". The interface includes a navigation menu on the left with categories like "Visit Details", "Procedures", "Custom Metadata", "Expression Profiles Data", "Medications", "Clinical Trials", "Diagnoses", "Reports", "Demographics", "Laboratory Tests", and "Providers". Below this is a "Workplace" section showing a list of "i2b2 Patients" with identifiers such as "1000000008 [23 y/o M hispanic]" and "PATIENT:1000000001". At the bottom left, there is a "Previous Queries" section listing various query titles and dates. The main area is the "Query Tool", which features a "Query Name" field and three columns for building a query: "Group 1", "Group 2", and "Group 3". Each column has sub-columns for "Dates", "Occurs > 0x", and "Exclude". A yellow box with the text "drop a term on here" is positioned in the first group. At the bottom of the query tool, there are buttons for "Run Query", "New Query", and "New Group", along with a status indicator showing "0 Groups". A large, bold, red "SLOW!" watermark is overlaid across the center of the interface.

# Functional Database

[dictionary.com](https://www.dictionary.com):

A database which uses a functional language as its query language. Databases would seem [...] **inappropriate** [...] since, a purely functional language would have to **return a new copy** of the **entire database** every time (part of) it was updated.

**But thats exactly what Datomic does!**


# Information Model

- Data stored in form of facts
- Fact: atomic information item
- Entity/Attribute/Value/Transaction (Time)

Entity	Attribute	Value	Time
Sally	likes	pizza	02.05.2013



# Transition DB $\rightarrow$ DB



# Indexes - EAVT

<b>E</b>	<b>A</b>	<b>V</b>	<b>Tx</b>	<b>Op</b>
41	:release/name	“Abbey Road”	1100	assert
42	:release/name	“Magical Mystery Tour”	1007	assert
42	:release/year	1967	1007	assert
42	:release/artistCredit	“The Beatles”	1007	assert
43	:release/name	“Let It Be”	1234	assert

# Indexes - AEVT

<b>A</b>	<b>E</b>	<b>V</b>	<b>Tx</b>	<b>Op</b>
:release/artistCredit	42	"The Beatles"	1007	assert
:release/name	41	"Abbey Road"	1007	assert
:release/name	42	"Magical Mystery Tour"	1007	assert
:release/name	43	"Let It Be"	1234	assert
:release/year	42	1967	1007	assert

# Indexes - AVET

<b>A</b>	<b>V</b>	<b>E</b>	<b>Tx</b>	<b>Op</b>
:release/name	"Abbey Road"	41	1199	assert
:release/name	"Let It Be"	43	1234	assert
:release/name	"Let It Be"	55	2367	assert
:release/name	"Magical Mystery Tour"	42	1007	assert
:release/year	1967	42	1007	assert
:release/year	1984	55	2367	assert

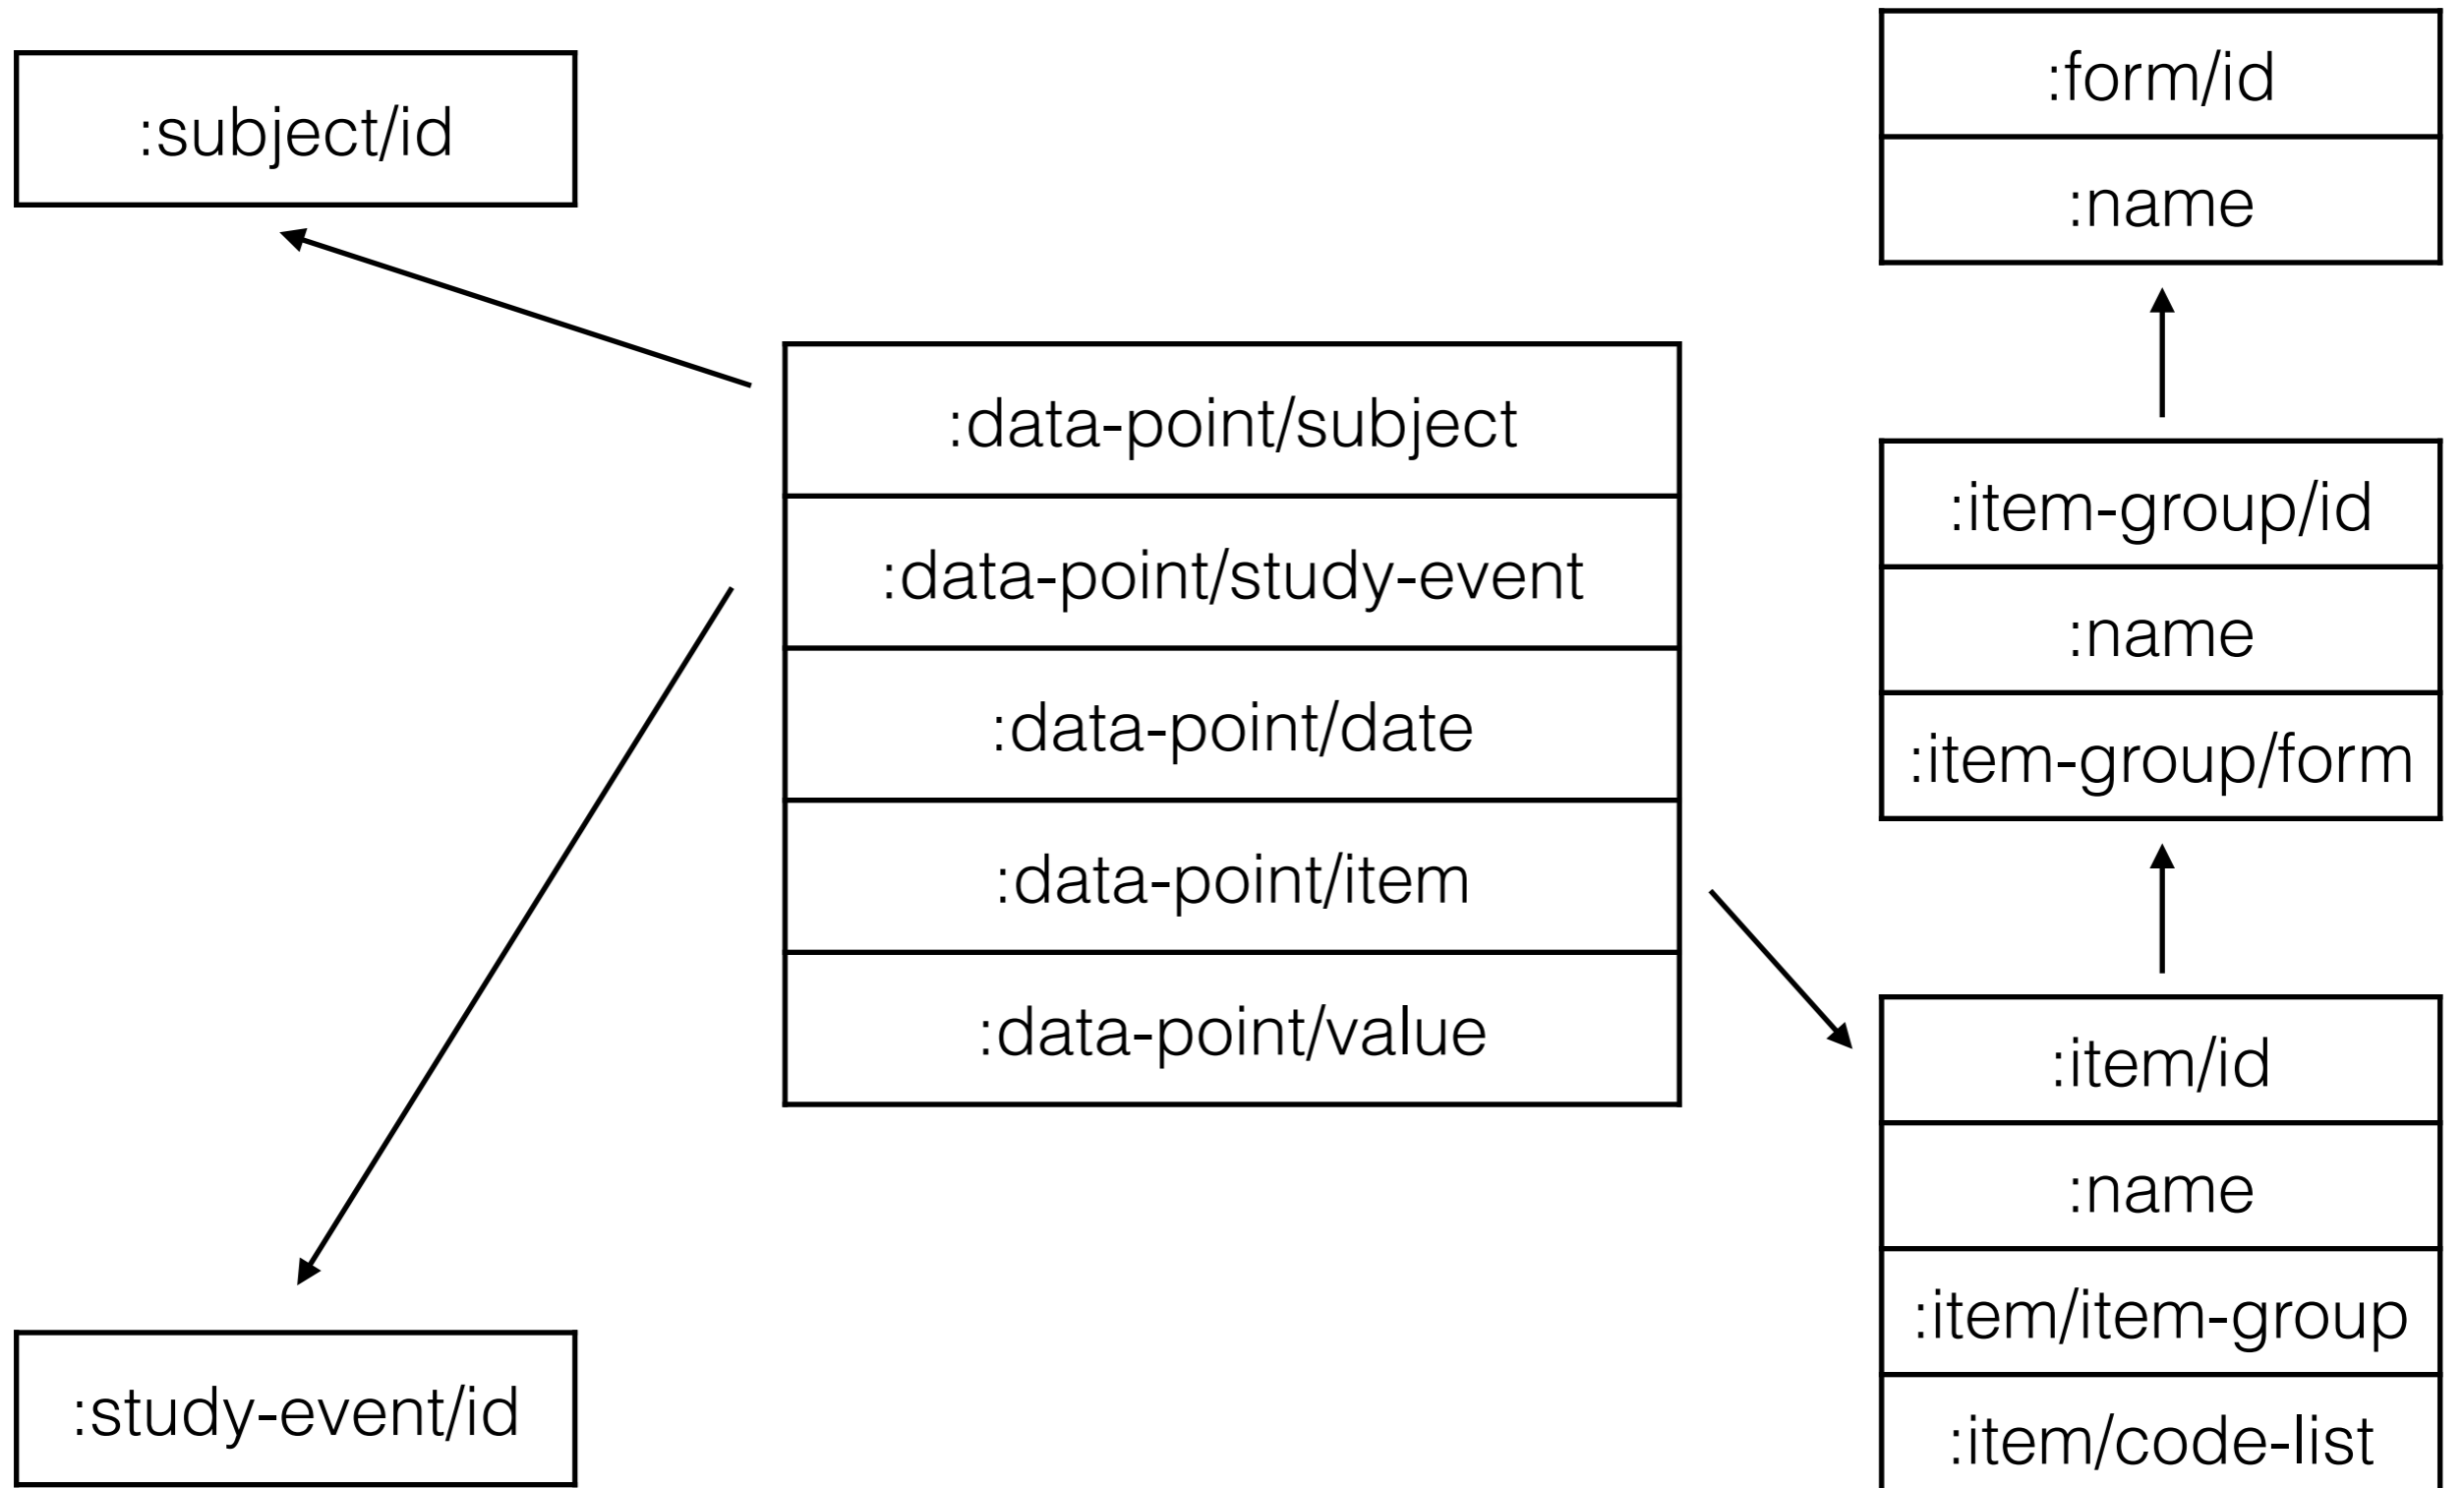
# Indexes - VAET

Release

V	A	E	Tx	Op
100	:release/artists	41	1007	assert
100	:release/artists	42	1007	assert
100	:release/artists	43	1007	assert
200	:release/artists	55	2367	assert

Artist

# Schema Research DB



**Demo**

# Lens Query Language


```
[ [ :form "T00001" ["A1_HAUPT01"] ] ]
```

```
[ [ :form "T00505" ["A1_HAUPT01"] ] ] ]
```




# Lens Query Language

conjunction



[ [<atom> <atom> <atom>] ← disjunction  
[<atom> <atom> <atom>]]



atom:

[<type> <id> [<study-event-id> ...]]

types:

- :form
- :item-group
- :item
- :code-list-item

# Atom Query Execution

## Datomic Datalog Query

```
(q `[:find (distinct ?s)
     :in $ ?i
     :where [?e :data-point/item ?i]
            [?e :data-point/subject ?s]]
 db item-id)
```

# Atom Query Execution

## Datomic Datalog Query

```
(q `[:find (distinct ?s)
     :in $ ?i
     :where [?e :data-point/item ?i]
            [?e :data-point/subject ?s]]
  db item-id)
```

**Database Reference and Arguments**

# Atom Query Execution

## Datomic Datalog Query

```
(q `[:find (distinct ?s)
     :in $ ?i
     :where [?e :data-point/item ?i]
            [?e :data-point/subject ?s]]
 db item-id)
```

## Unification of Clauses

# Atom Query Execution

## Datomic Datalog Query

```
(q `[:find (distinct ?s)
    :in $ ?i
    :where [?e :data-point/item ?i]
           [?e :data-point/subject ?s]]
db item-id)
```

**Entity**

# Atom Query Execution

## Datomic Datalog Query

```
(q `[:find (distinct ?s)
    :in $ ?i
    :where [?e :data-point/item ?i]
           [?e :data-point/subject ?s]]
 db item-id)
```

**Attribute**

# Atom Query Execution

## Datomic Datalog Query

```
(q `[:find (distinct ?s)
    :in $ ?i
    :where [?e :data-point/item ?i]
           [?e :data-point/subject ?s]]
db item-id)
```

Value

# Atom Query Execution

[?e :data-point/item ?i]

1

<b>V</b>	<b>A</b>	<b>E</b>
100	:data-point/item	23
100	:data-point/item	42
100	:data-point/item	96

Reverse Index



# Atom Query Execution

[?e :data-point/item ?i]

Reverse Index

1

V	A	E
100	:data-point/item	23
100	:data-point/item	42
100	:data-point/item	96

[?e :data-point/subject ?s]

Attr. Index

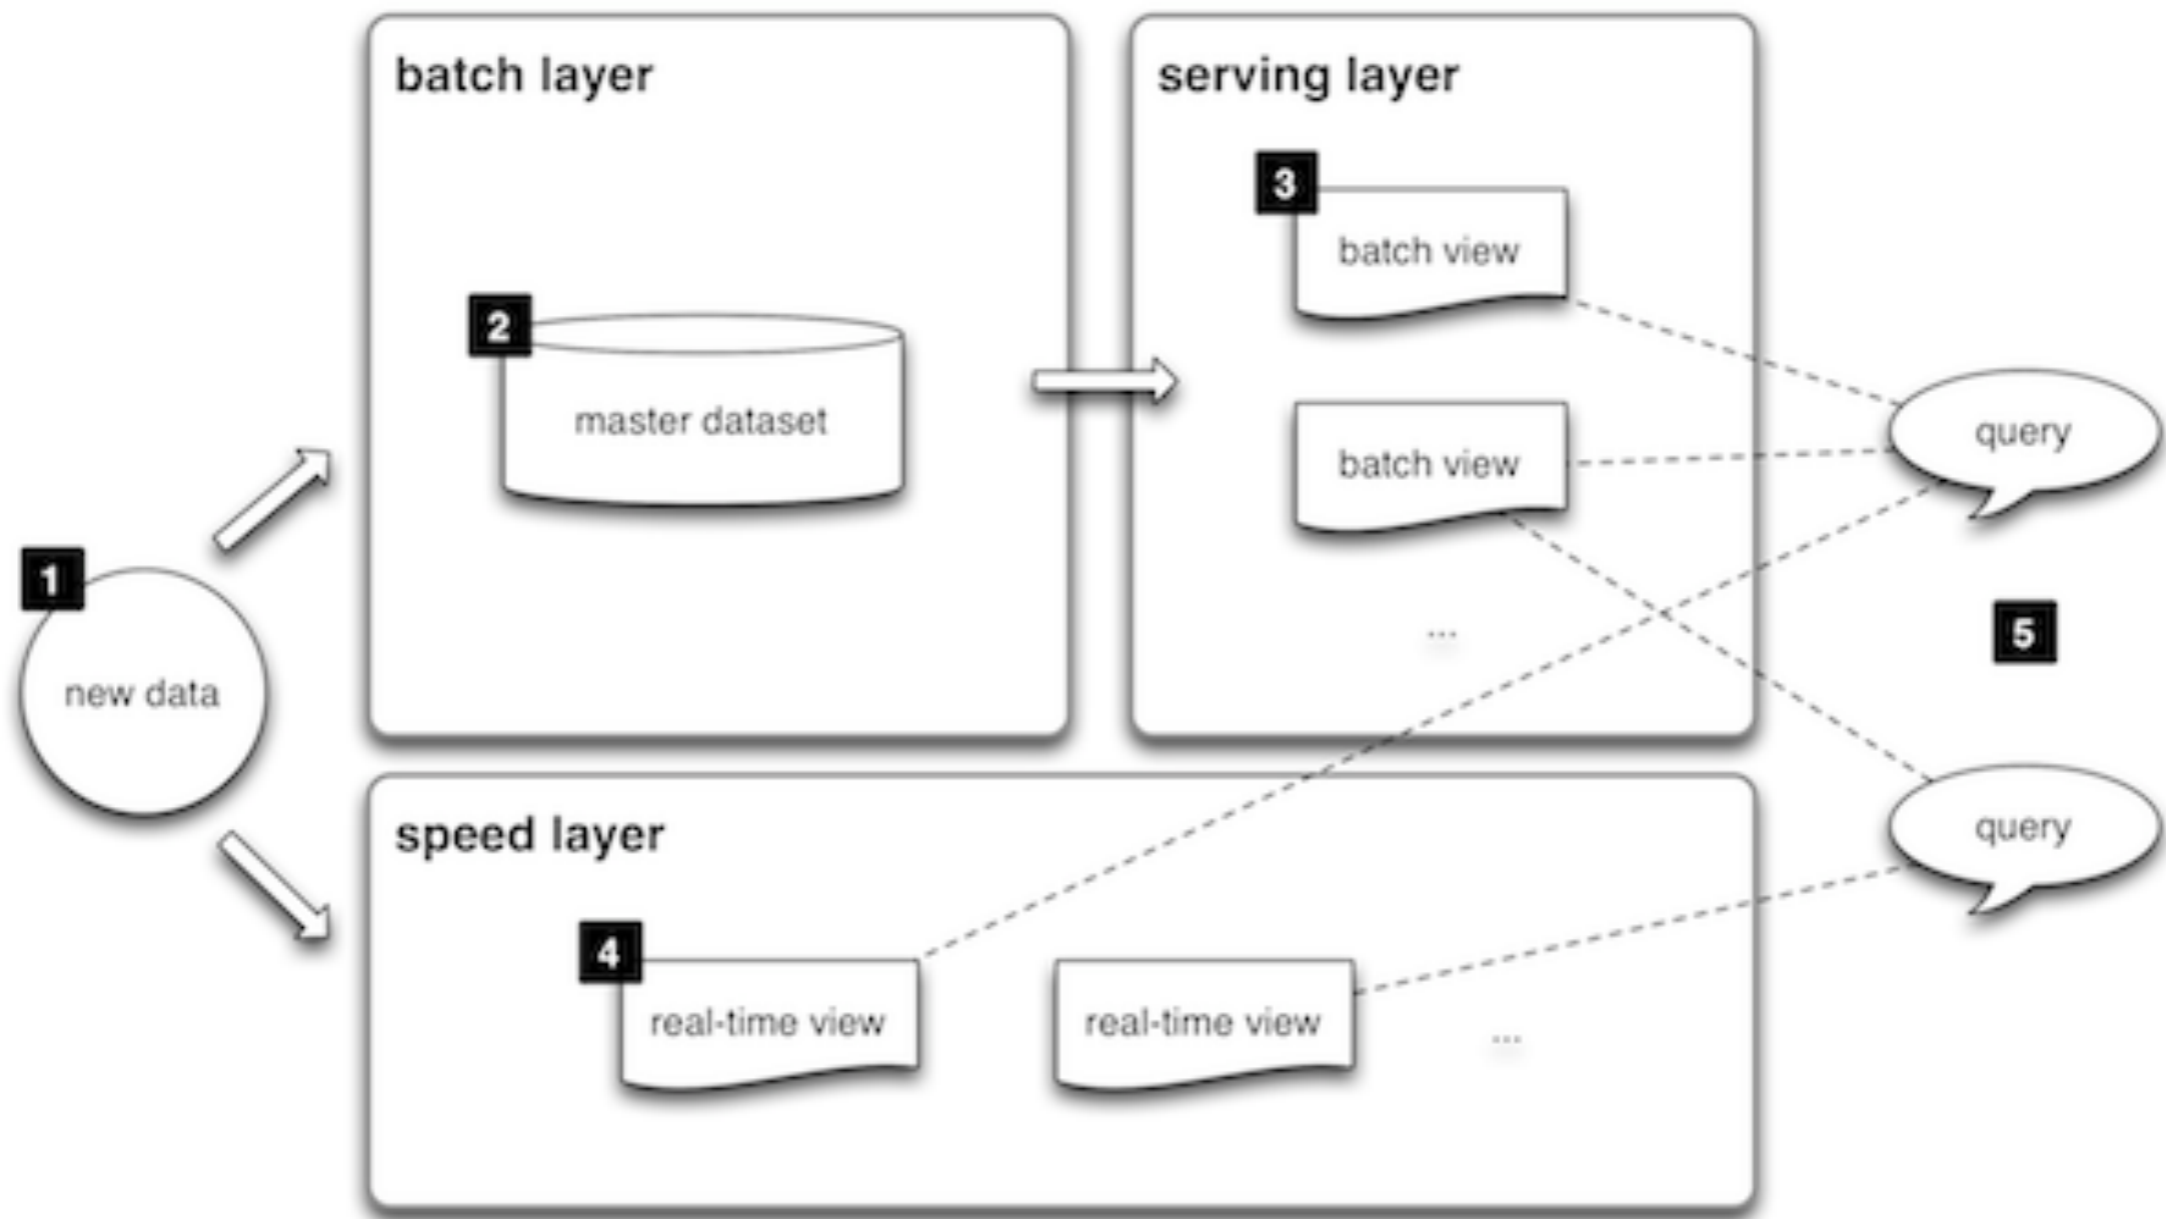
2

A	E	V
:data-point/subject	23	1
:data-point/subject	42	1
:data-point/subject	96	2

→ subjects

# Lambda Architecture

Nathan Marz



# Batch-View Schema

:subject-stat/form

:subject-stat/subjects +

:subject-stat/item

:subject-stat/subjects +

:subject-stat/item-group

:subject-stat/subjects +

:subject-stat/code-list-item

:subject-stat/subjects +

**+ = cardinality/many**

# Batch-View Queries

No Datalog - just Clojure

```
(-> (entity bv [ :subject-stat/form 123 ])  
    ( :subject-stat/subjects ) )
```

# Batch-View Queries

No Datalog - just Clojure

```
(-> (entity bv [:subject-stat/form 123])  
    (:subject-stat/subjects))
```

**Simple Reference to Stat Entity**

# Batch-View Queries

No Datalog - just Clojure

```
(-> (entity bv [:subject-stat/form 123])  
    (:subject-stat/subjects))
```

**entity function returns map of attributes**

# Batch-View Queries

No Datalog - just Clojure

```
(-> (entity bv [:subject-stat/form 123])  
    (:subject-stat/subjects))
```

**Batch-View Database**

# Batch-View Queries

No Datalog - just Clojure

```
(-> (entity bv [:subject-stat/form 123])  
    (:subject-stat/subjects))
```

**Lookup Ref**



# Batch-View Queries

No Datalog - just Clojure

```
(-> (entity bv [:subject-stat/form 123])  
    (:subject-stat/subjects) )
```

**Navigation of Maps by Keyword**

# Batch-View Queries

(entity bv [:subject-stat/form 123])

1

Reverse Index

V	A	E
123	:subject-stat-form	42

(:subject-stat/subjects)

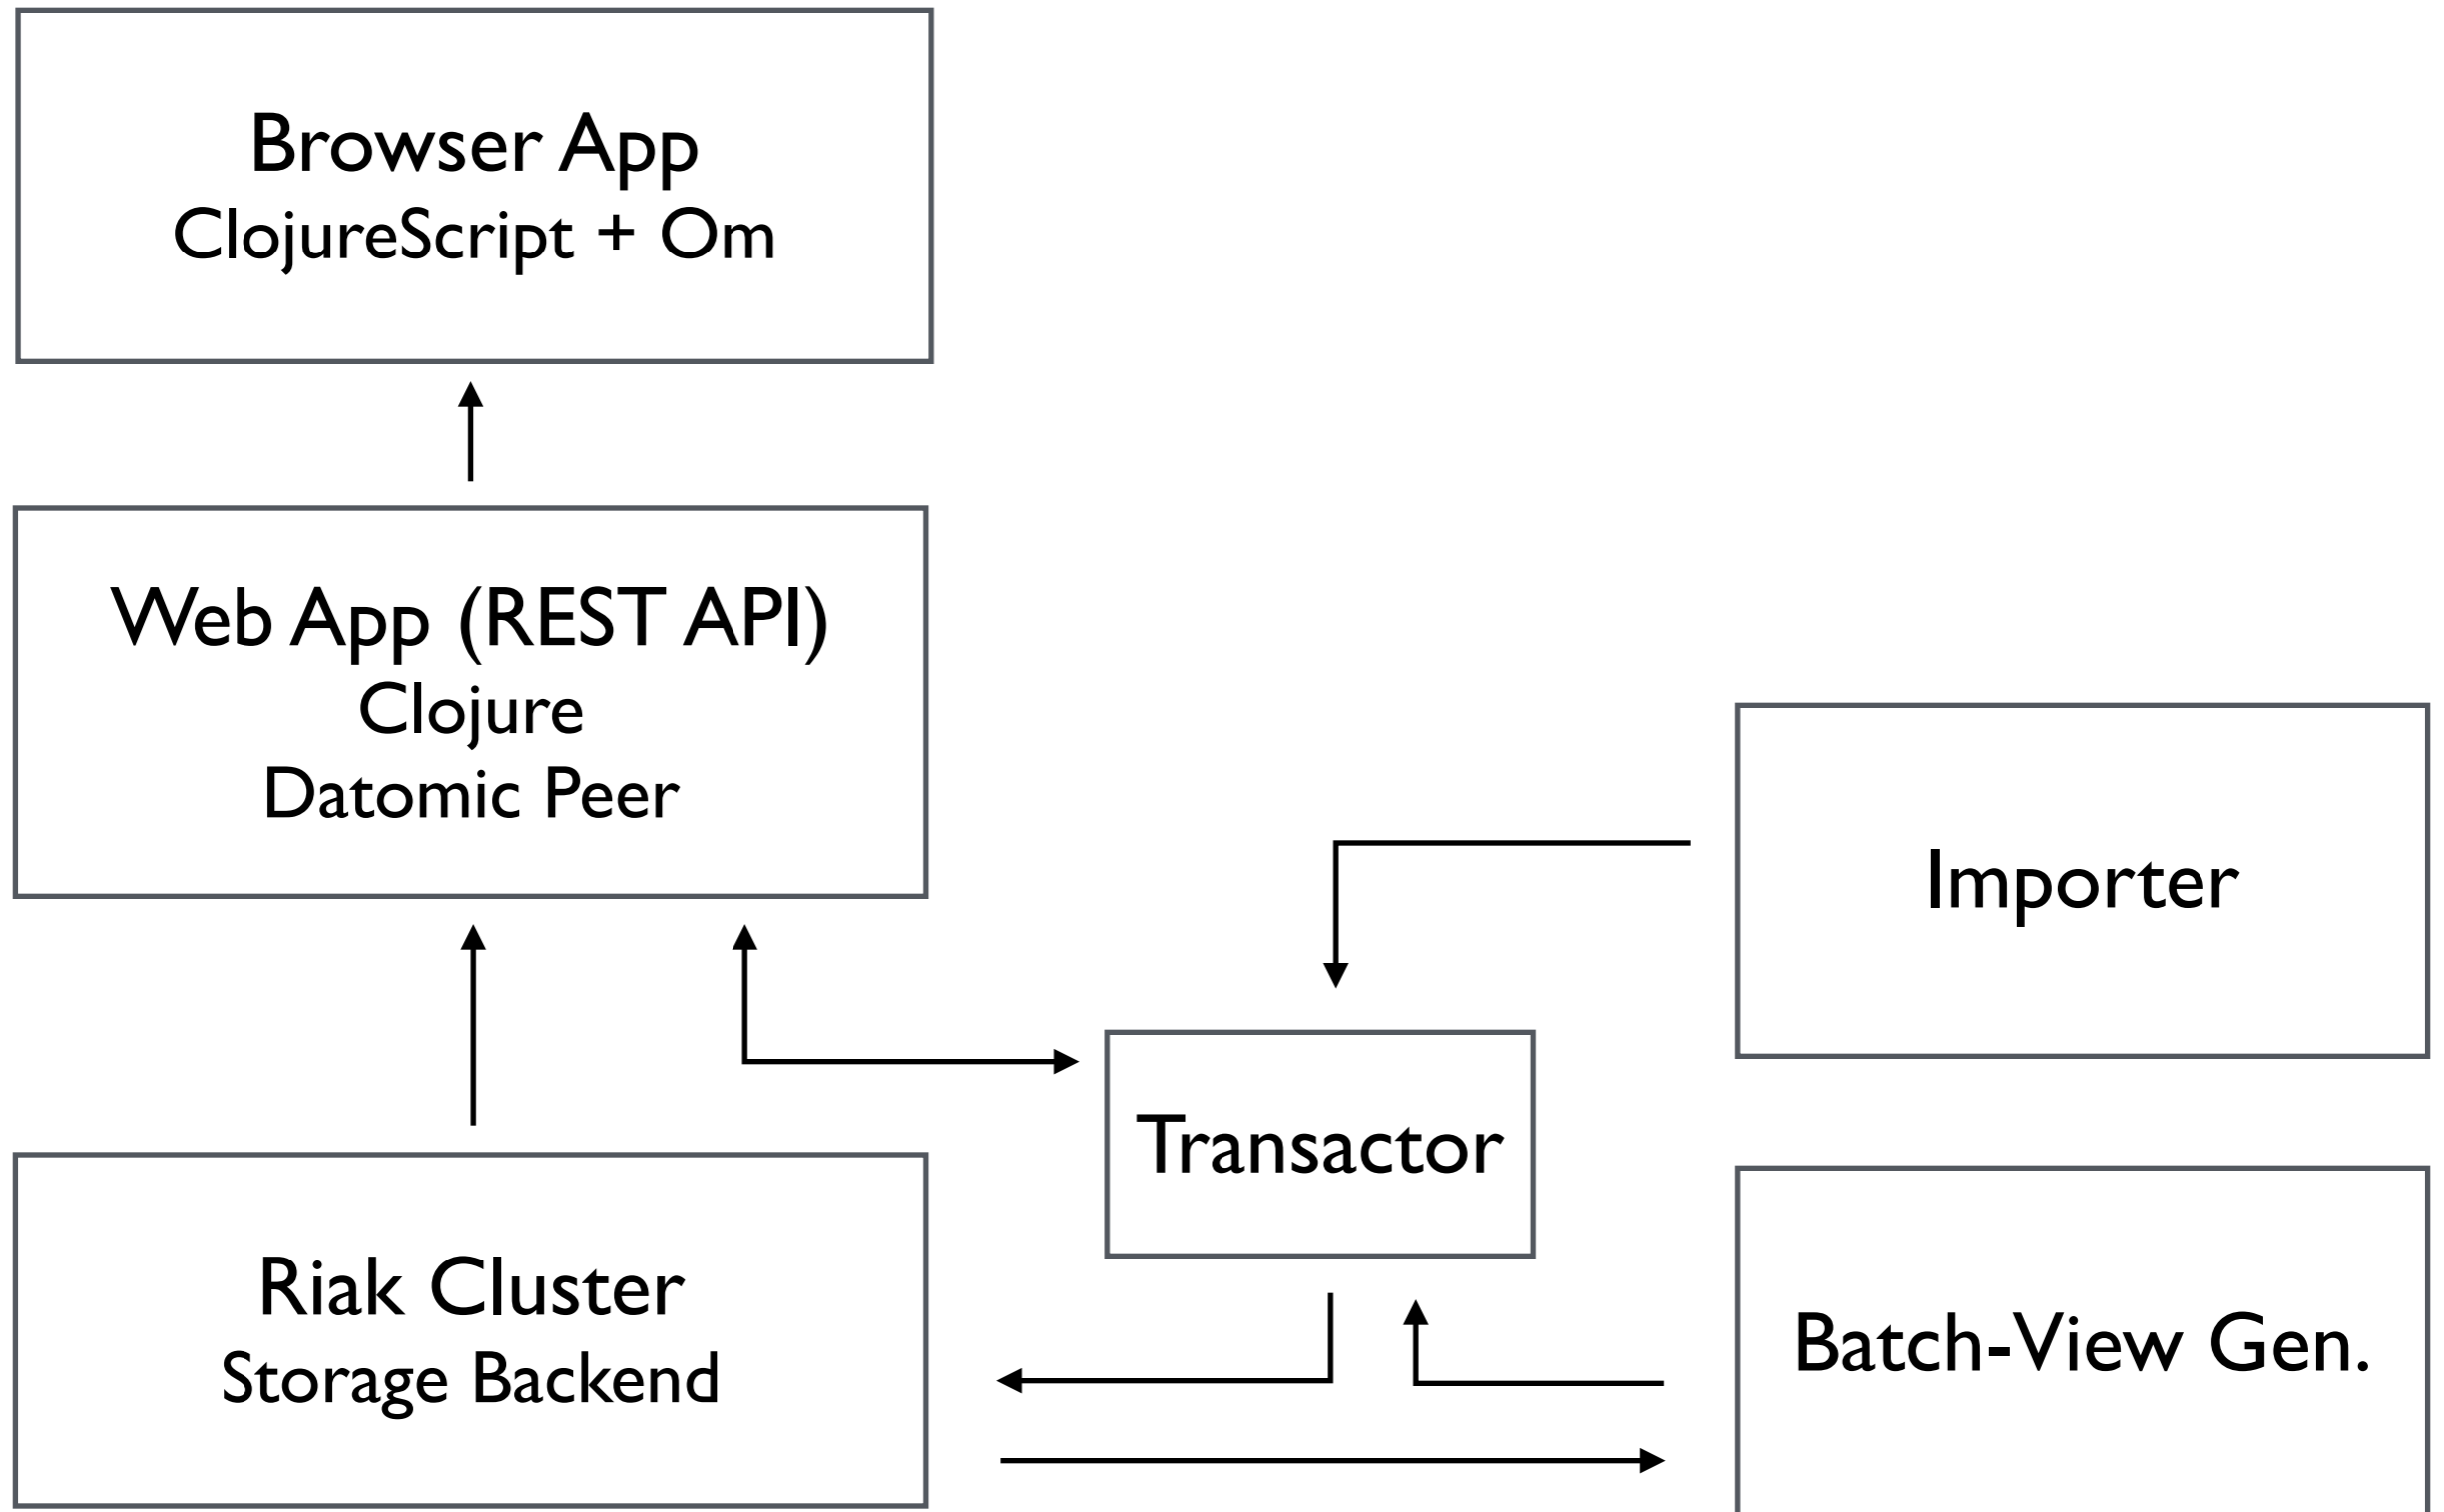
2

Attr. Index

A	E	V
:subject-stat/subjects	42	1
:subject-stat/subjects	42	2
:subject-stat/subjects	42	3

→ subjects

# Lens Architecture



**Thanks**