

# Event Sourcing on the JVM

## QCon San Francisco 2007

[Speakers](#)  
[Schedule](#)  
[Tutorials](#)  
[Tracks](#)  
[Keynotes](#)

[Social Events](#)

[Exhibition](#)  
[Sponsors](#)

[Registration](#)  
[Volunteers](#)  
[Venue](#)

[Travel](#)  
[Hotels](#)

[About QCon](#)

[All QCon events](#)

### **Presentation: "Scaling Domain Driven Design"**

**Track:**  Architecture Quality (day 2)

**Time:** Thursday 16:00 - 17:00

**Location:** Metropolitan I

**Abstract:** We've all heard about Domain Driven Design (DDD), but have you seen it applied in a large system? In the session you'll see how DDD was put to work to help design a high performance application processing 2000 sustained transactions per second with over 30GB of data added daily! It was DDD and the new 'Asynchronous Context Mapping' pattern that made this system possible with near linear scaling. Learn how to apply these techniques in your applications and how to overcome the challenges of DDD, such as reporting. This session makes large scale DDD make sense.

 [Download slides](#)

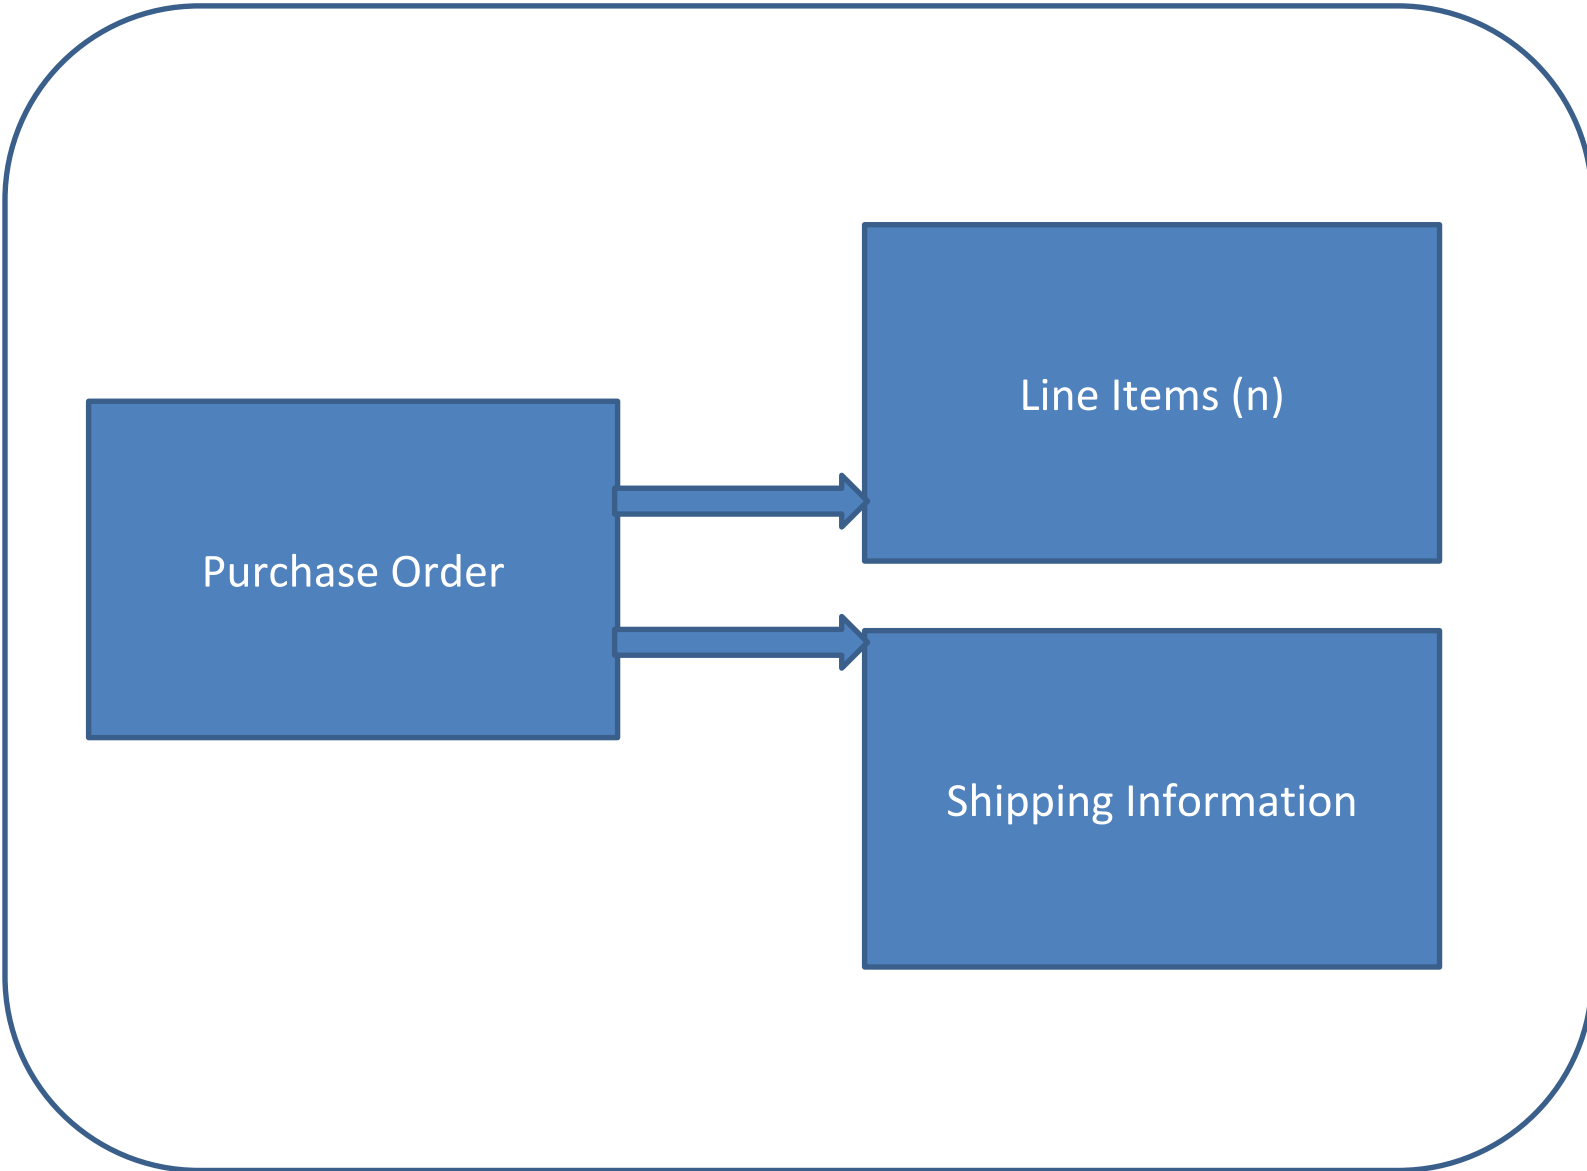
### **Greg Young, IMIS**



writing articles for InfoQ throughout the south east northwest, or floating up kayak.



807.50	481.00
136.00	726.00
36/33	596.00
25/16	622.00
52/6	763.00
21/7	





Cart Created

3 Items  
Added

Shipping  
Information  
Added



481.00	7126.00	5916.00	6222.00	7630.00	217.00
807.50	136.00	5916.00	6222.00	7630.00	217.00
54/53	36/33	92/46	52/6	87/87	

```
graph LR; A[Cart Created] --> B[3 Items Added]; B --> C[1 Item Removed]; C --> D[Shipping Information Added];
```

Cart Created

3 Items  
Added

1 Item  
Removed

Shipping  
Information  
Added



**Replay**






**Skip**






Just functional code.







**Cassandra**



**mongoDB**

**Redis**



**riak**

**CouchDB**

*relax*

There is no “best” storage.

 WILEY

# Succeeding with **Object Databases**

A Practical  
Look at Today's  
Implementations  
with Java<sup>™</sup>  
and XML

Akmal B. Chaudhri  
Roberto Zicari



A single data model is almost never appropriate. 3+  
is common!



twitter

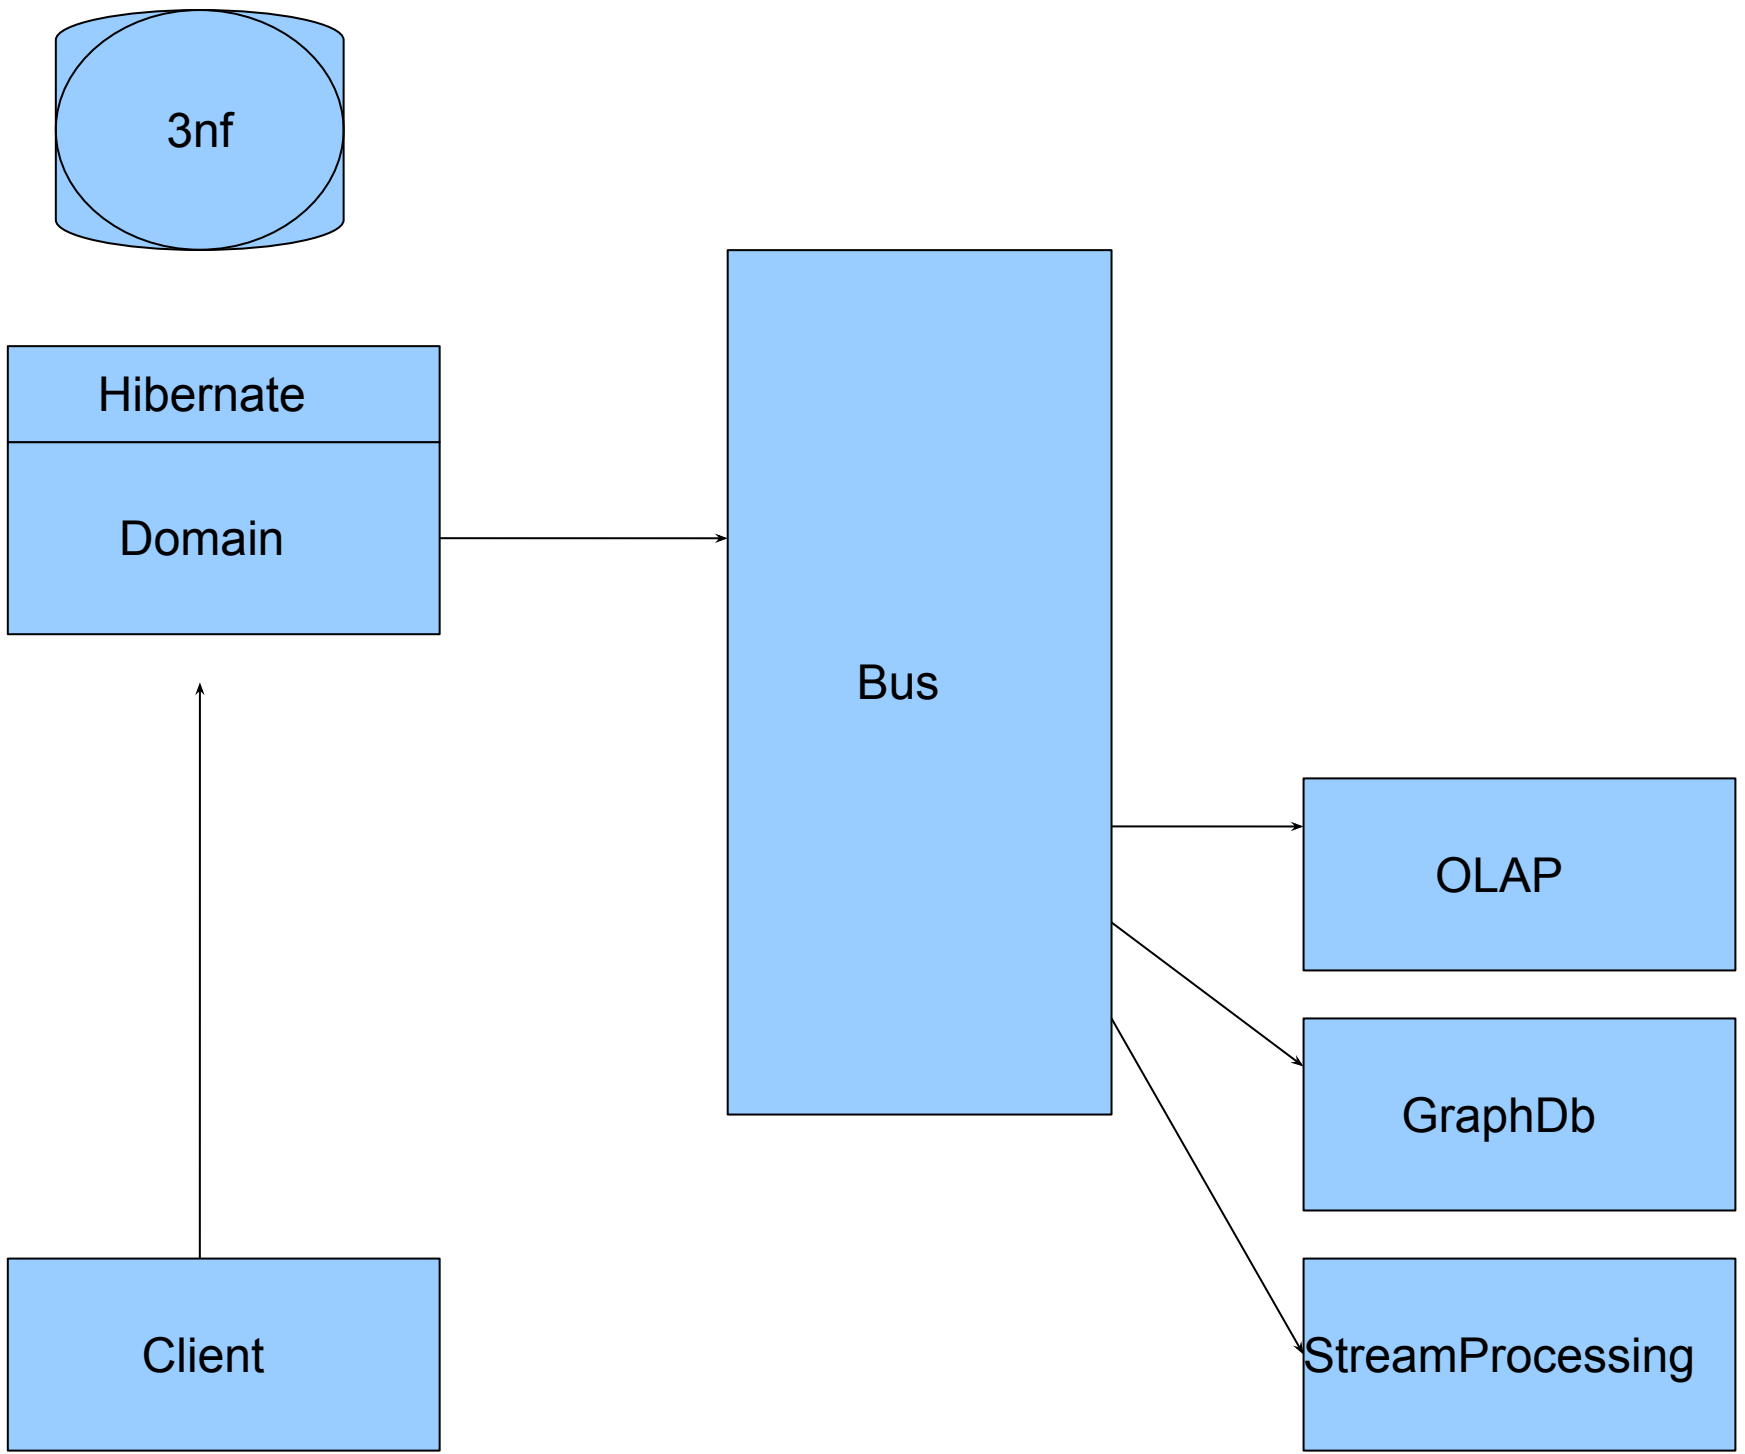


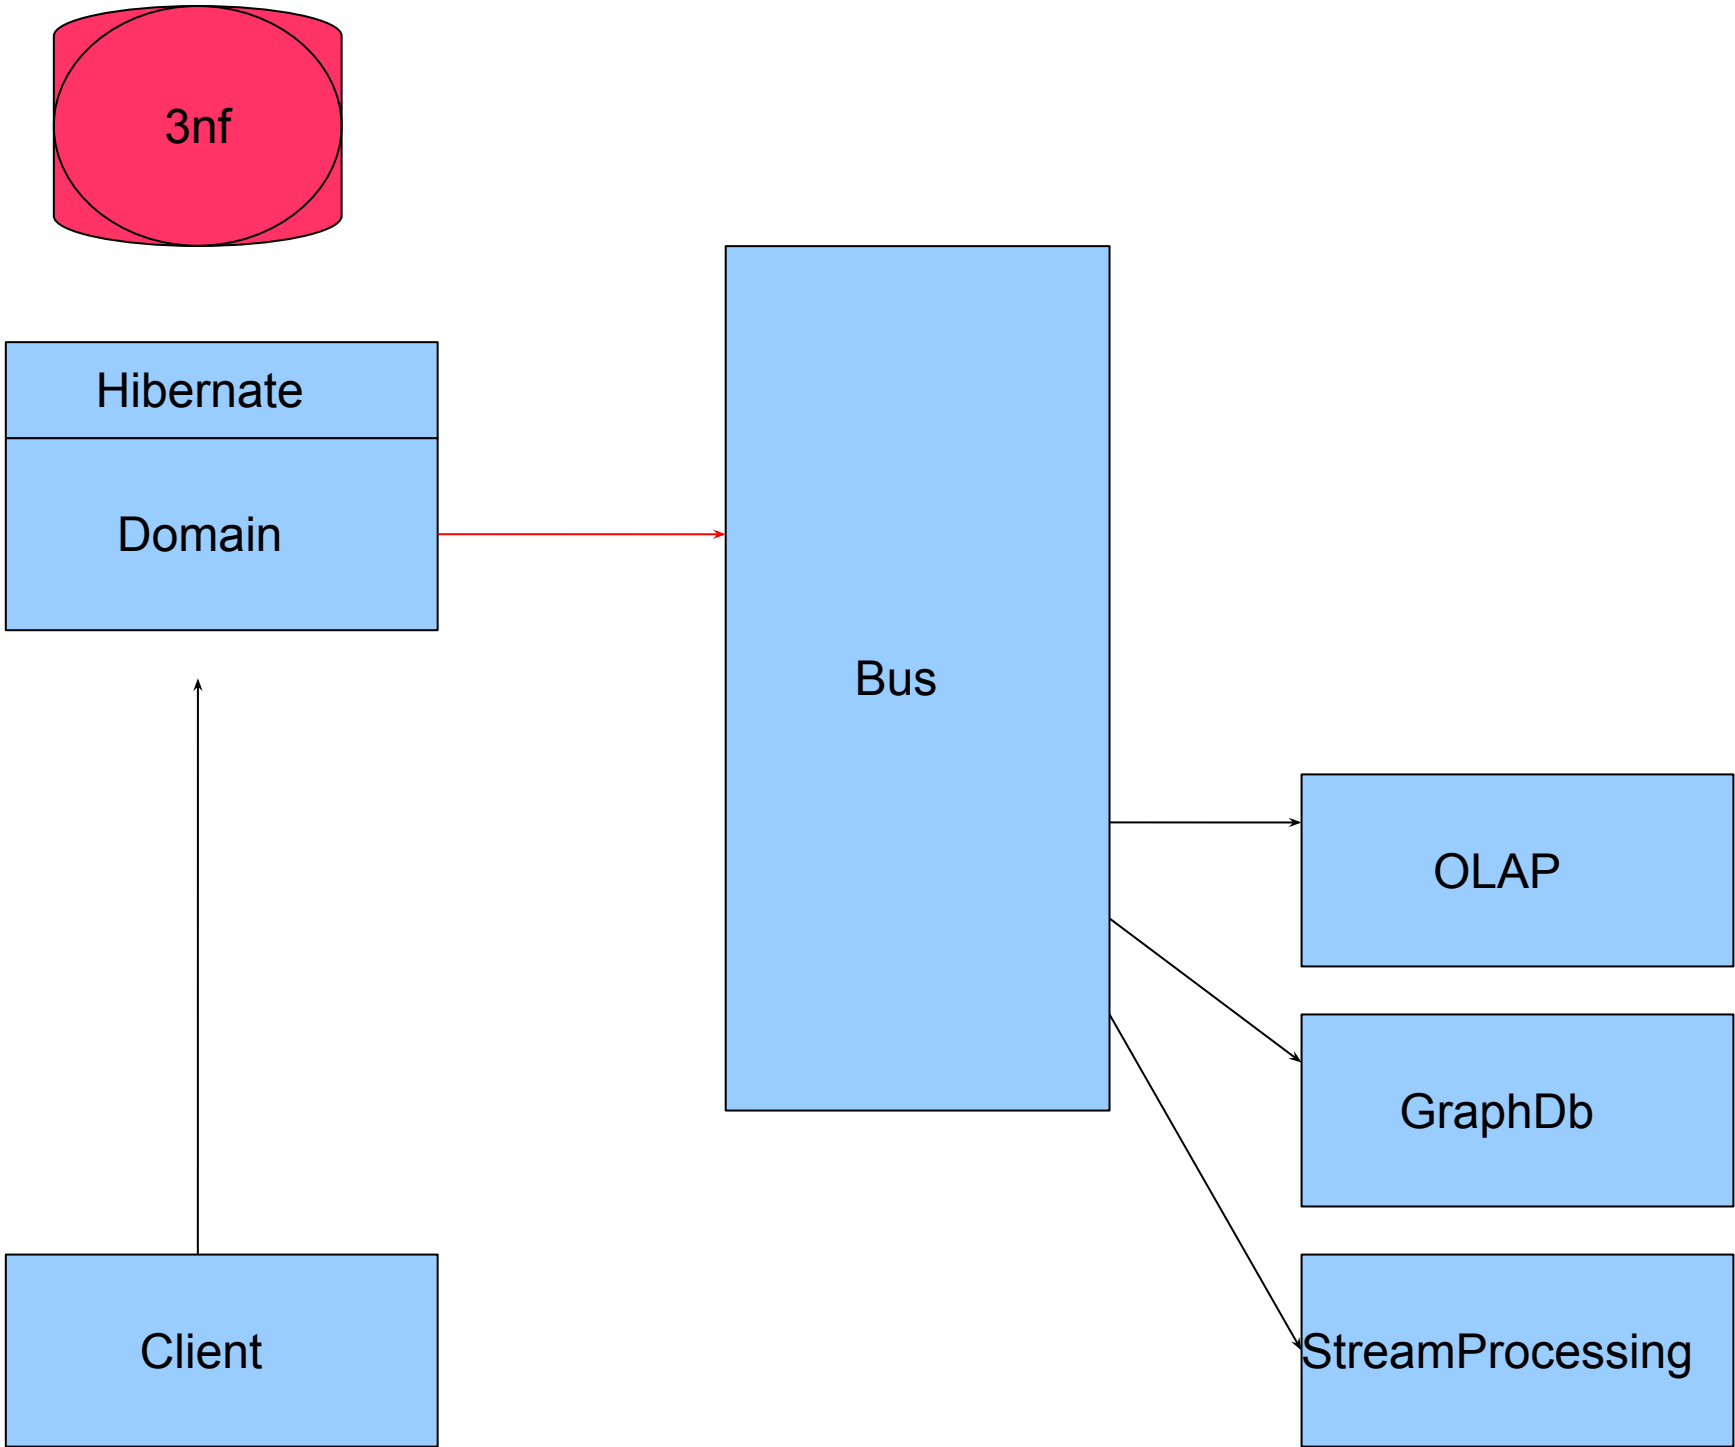





Id	ParentId	Data
1	0	Parent
2	1	Child1
3	1	Child2

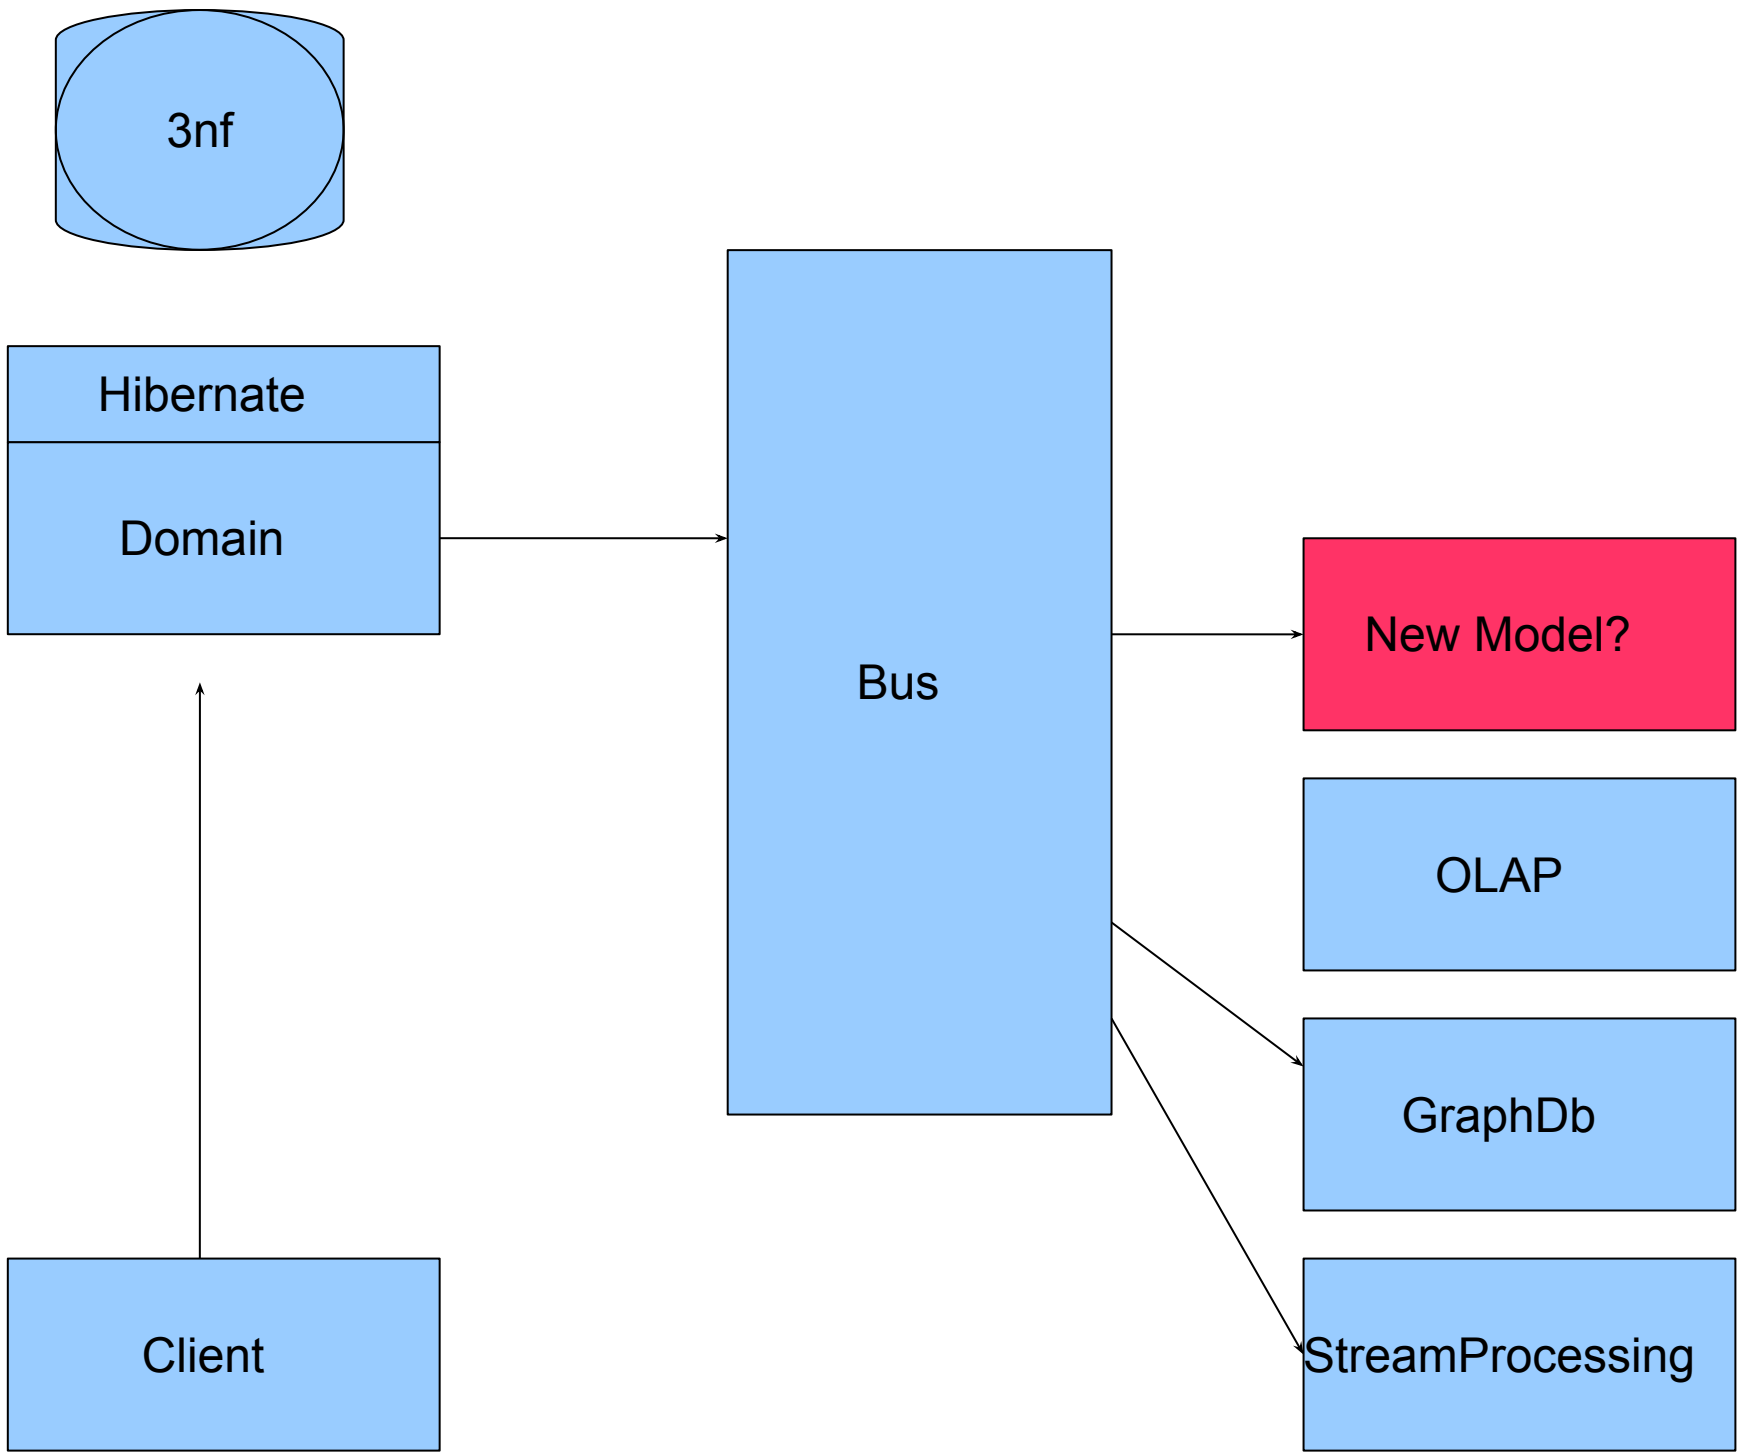
Id	ParentId0	ParentId1	ParentId2	ParentId3	ParentId4	Data
1	0	0	0	0	0	Parent
2	1	0	0	0	0	Child1
3	1	2	0	0	0	Child2

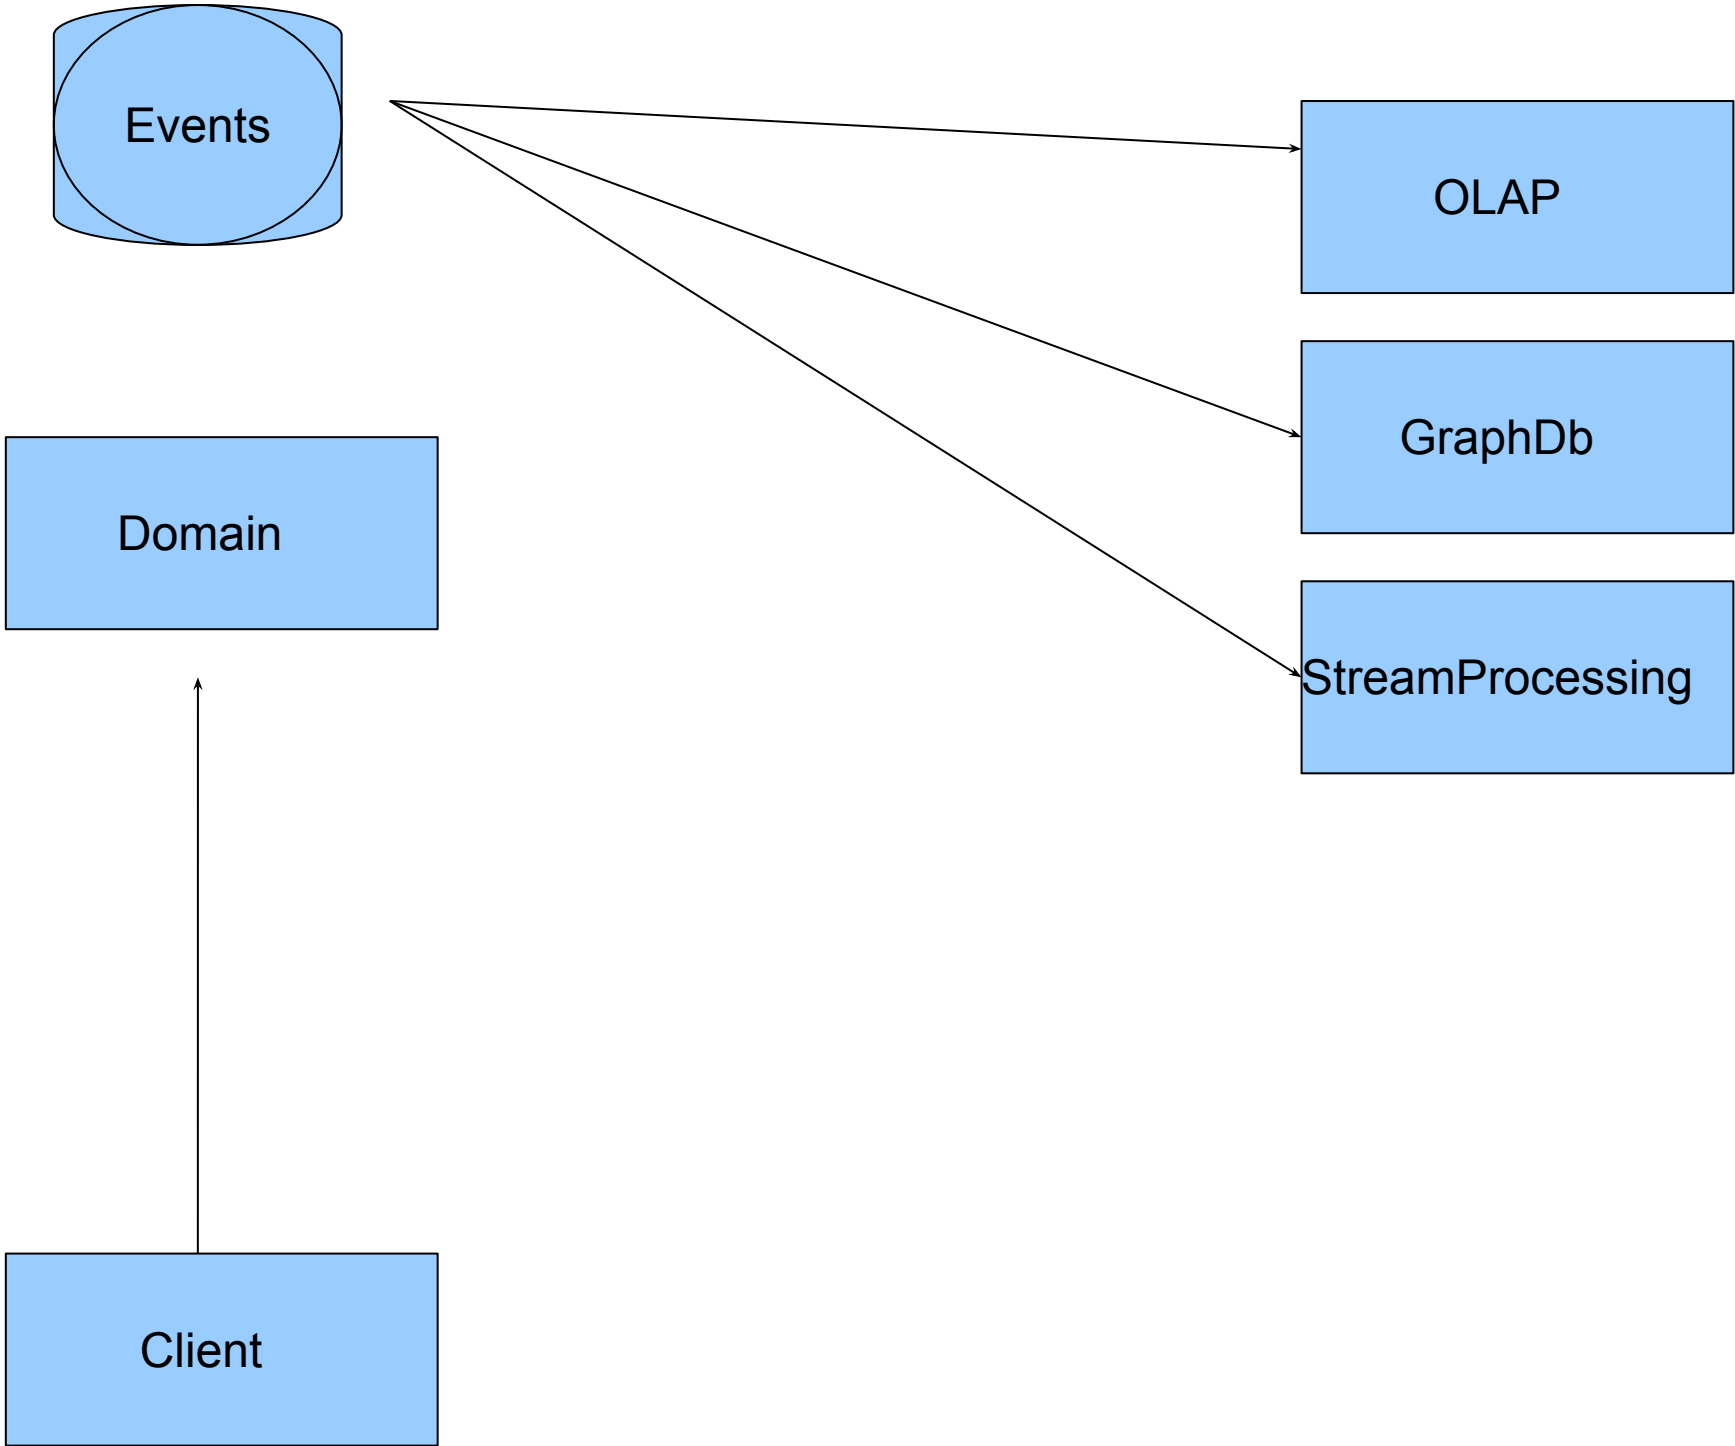
Wrong models cause massive accidental complexity

A queue with 'infinite history' allows new models to easily be brought online and to sync

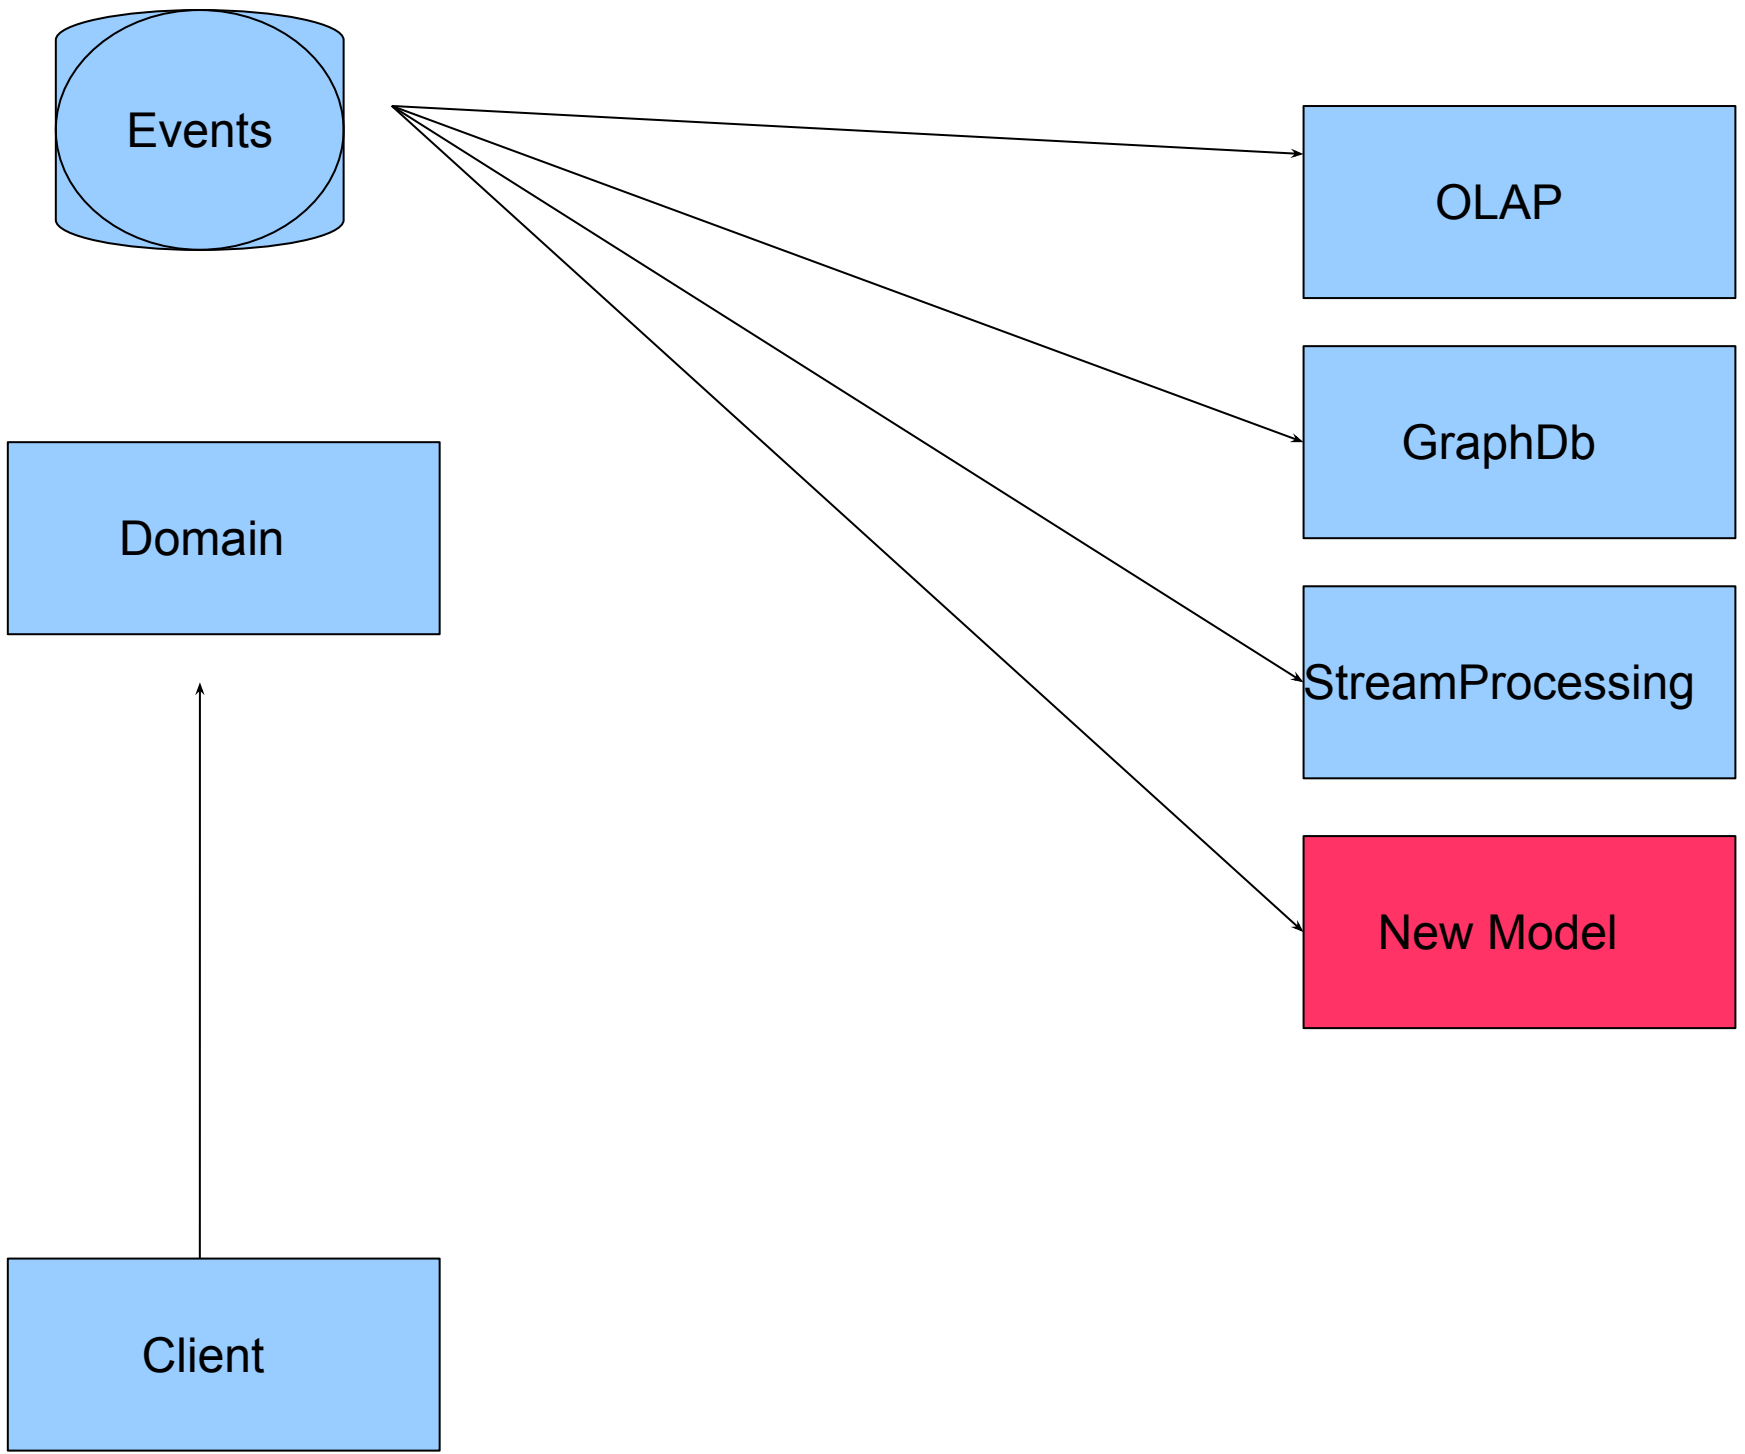




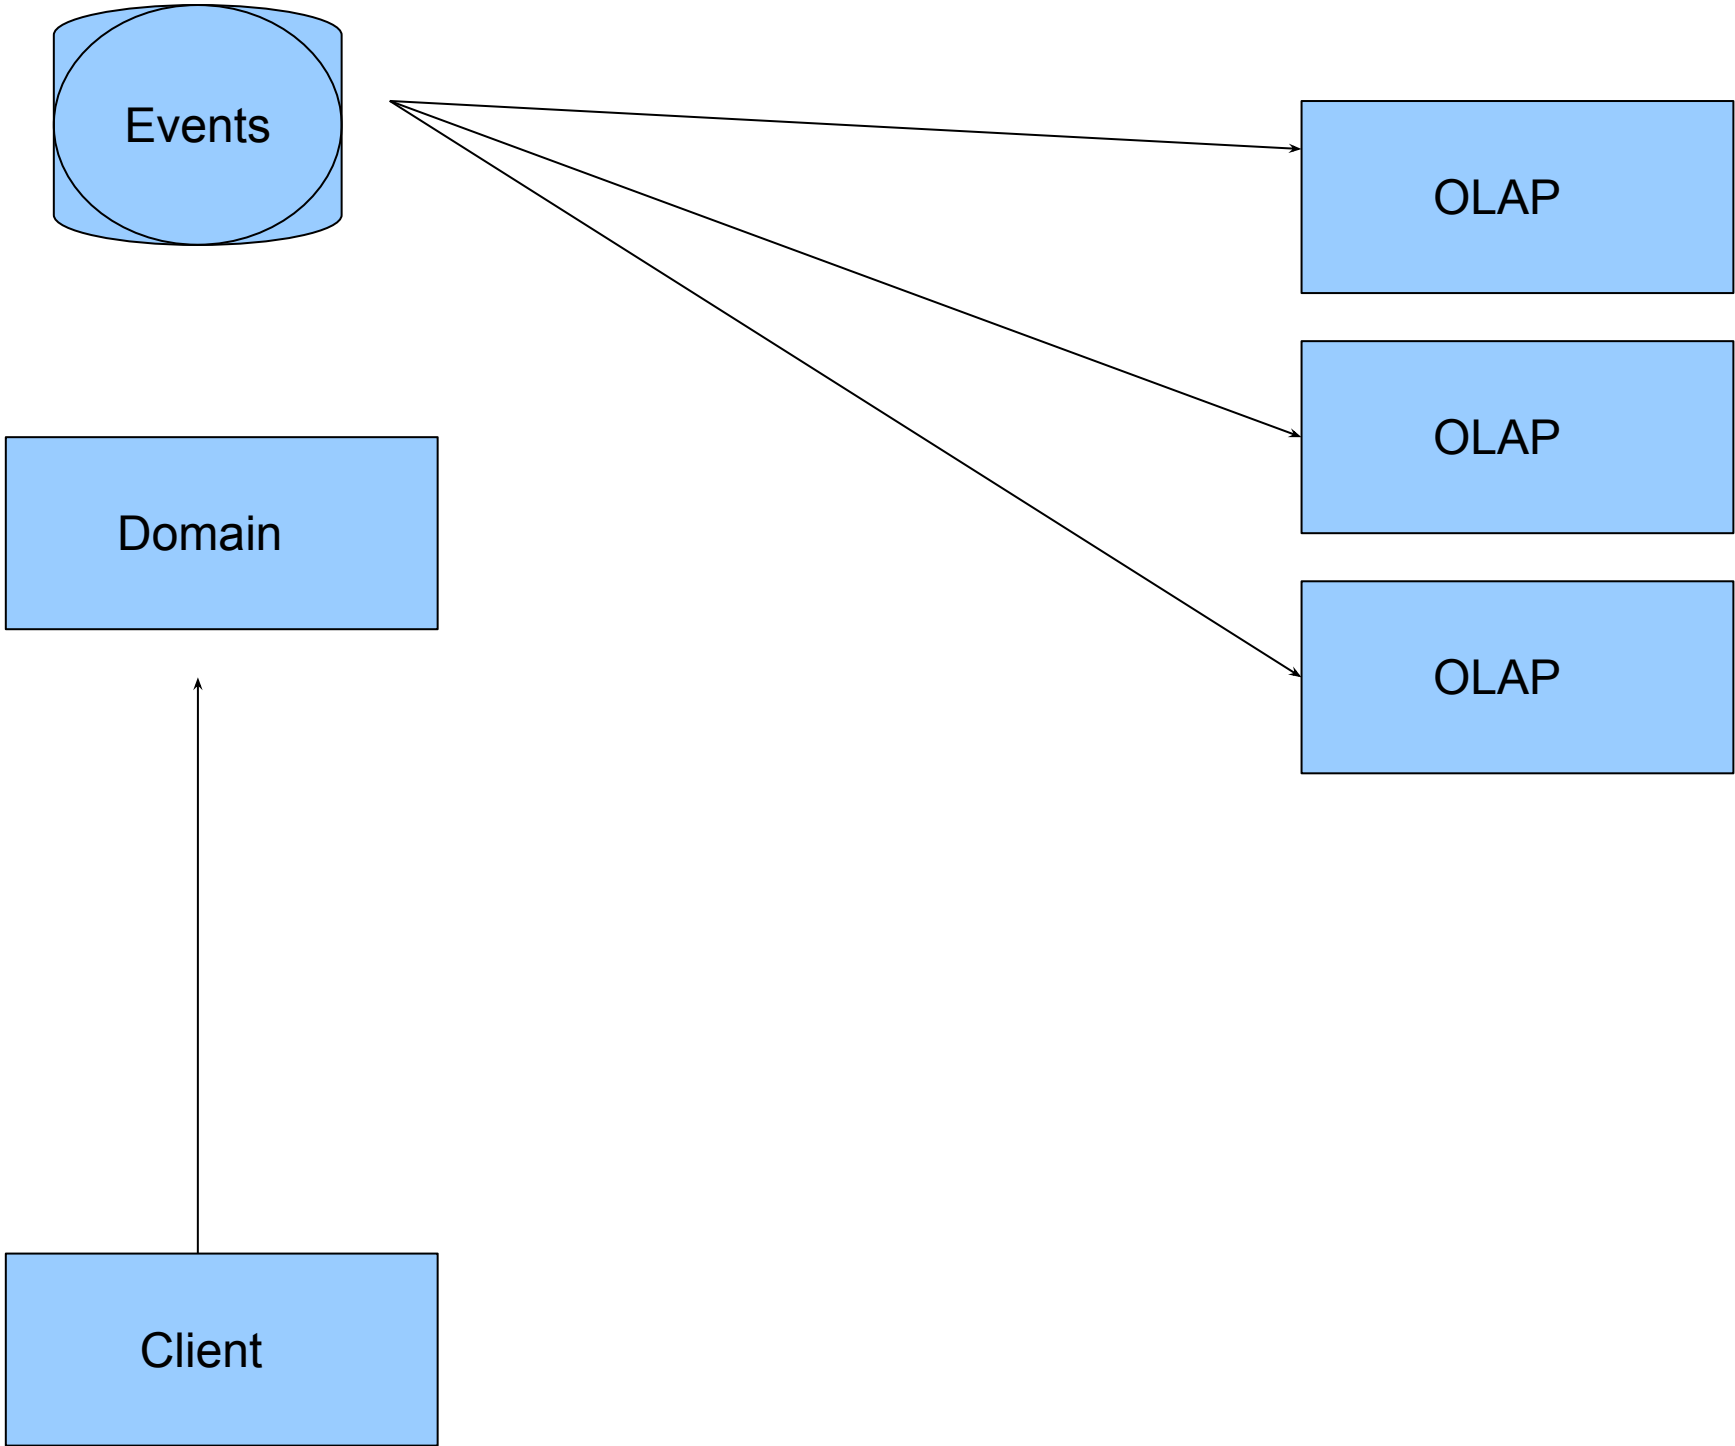


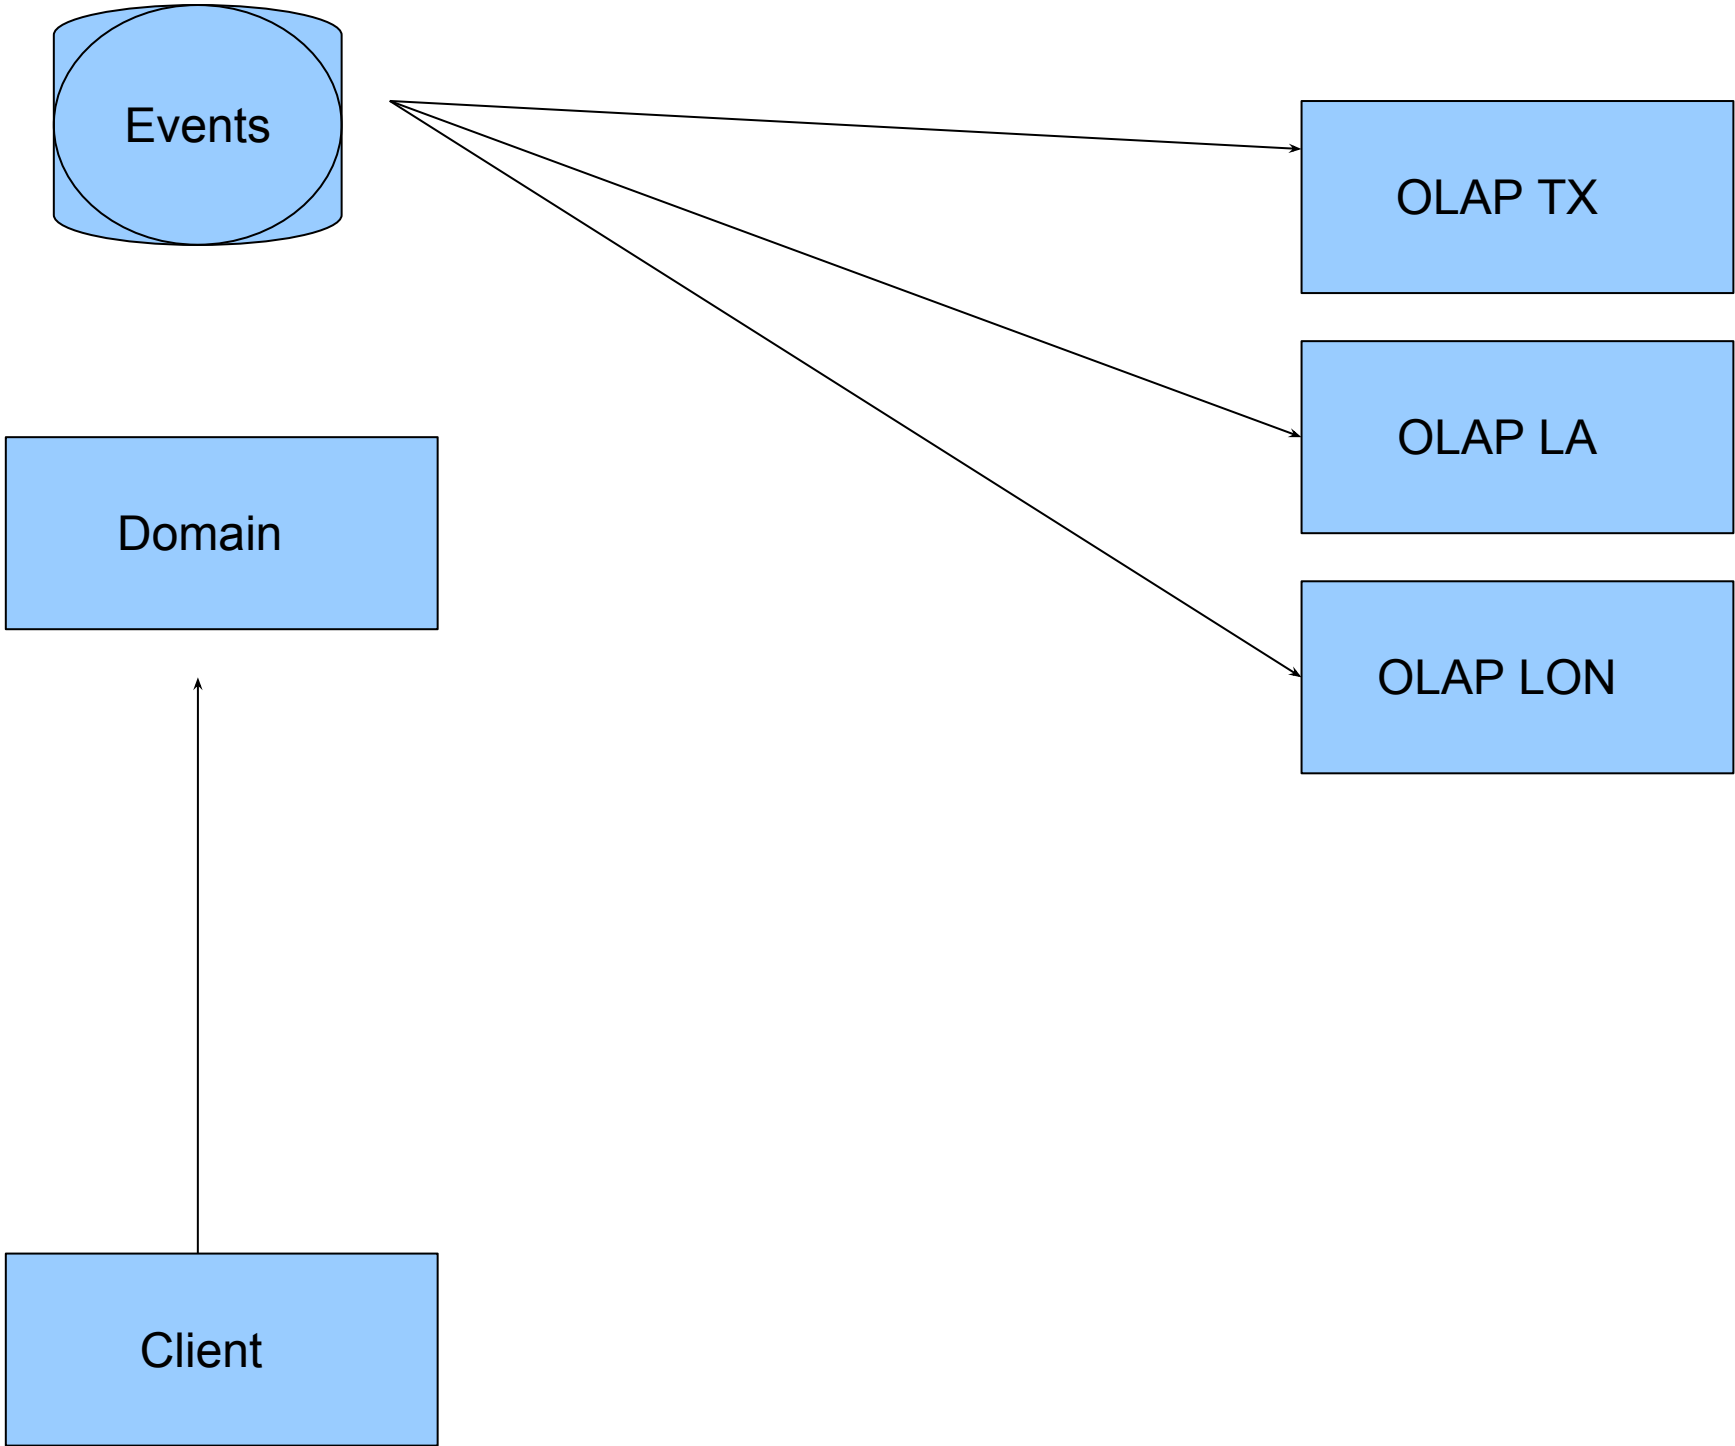






Multiple models simplify many things!





Projection : Handles<Message1>,  
Handles<Message2>

The framework/tool space is confusing.

Kafka

EventStore (GetEventStore)



# Axon event stores

Eventuate

# Axon framework

Akka.Persistence

# Axon framework

Concurrency?

Reactive Streams?

Prediction: More to come



Lots of choices!

Questions?