Measuring 2.0

ir. drs. Niels Basjes
Lead IT Architect Scalable Solutions

Agenda

- **Introduction**
- **Bol.com**
  - Personalization
- **Measuring 1.0**
  - Limitations
- **Measuring 2.0**
  - What do we really want
  - What is possible
- **Stream Processing**
  - Design
  - Implementation
  - Operations

Niels Basjes

nbasjes@bol.com
@nielsbasjes
https://github.com/nielsbasjes

TU-Delft Computer Science
Nyenrode Business School

Software developer
Research Scientist (NLR)
Infra Architect (NLR)
WebAnalytics Architect
Lead IT Architect (Bol.com)

Contributor Apache Hadoop, Pig, HBase, Storm, Flink, …

Committer Apache Avro

bol.com

Bol.com

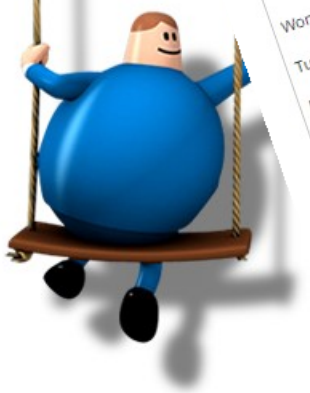

bol.com

> 11 million products for sale
> 40 million in catalog

> 6 million active customers
> 27 million visits per month

> 2500 million pageviews/year

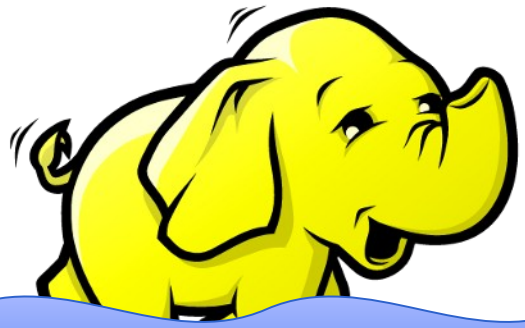

Hadoop in production
since 2010

bol.com

Happy customers

Change the website to fit the needs of the visitor

bol.com

Search Suggestions

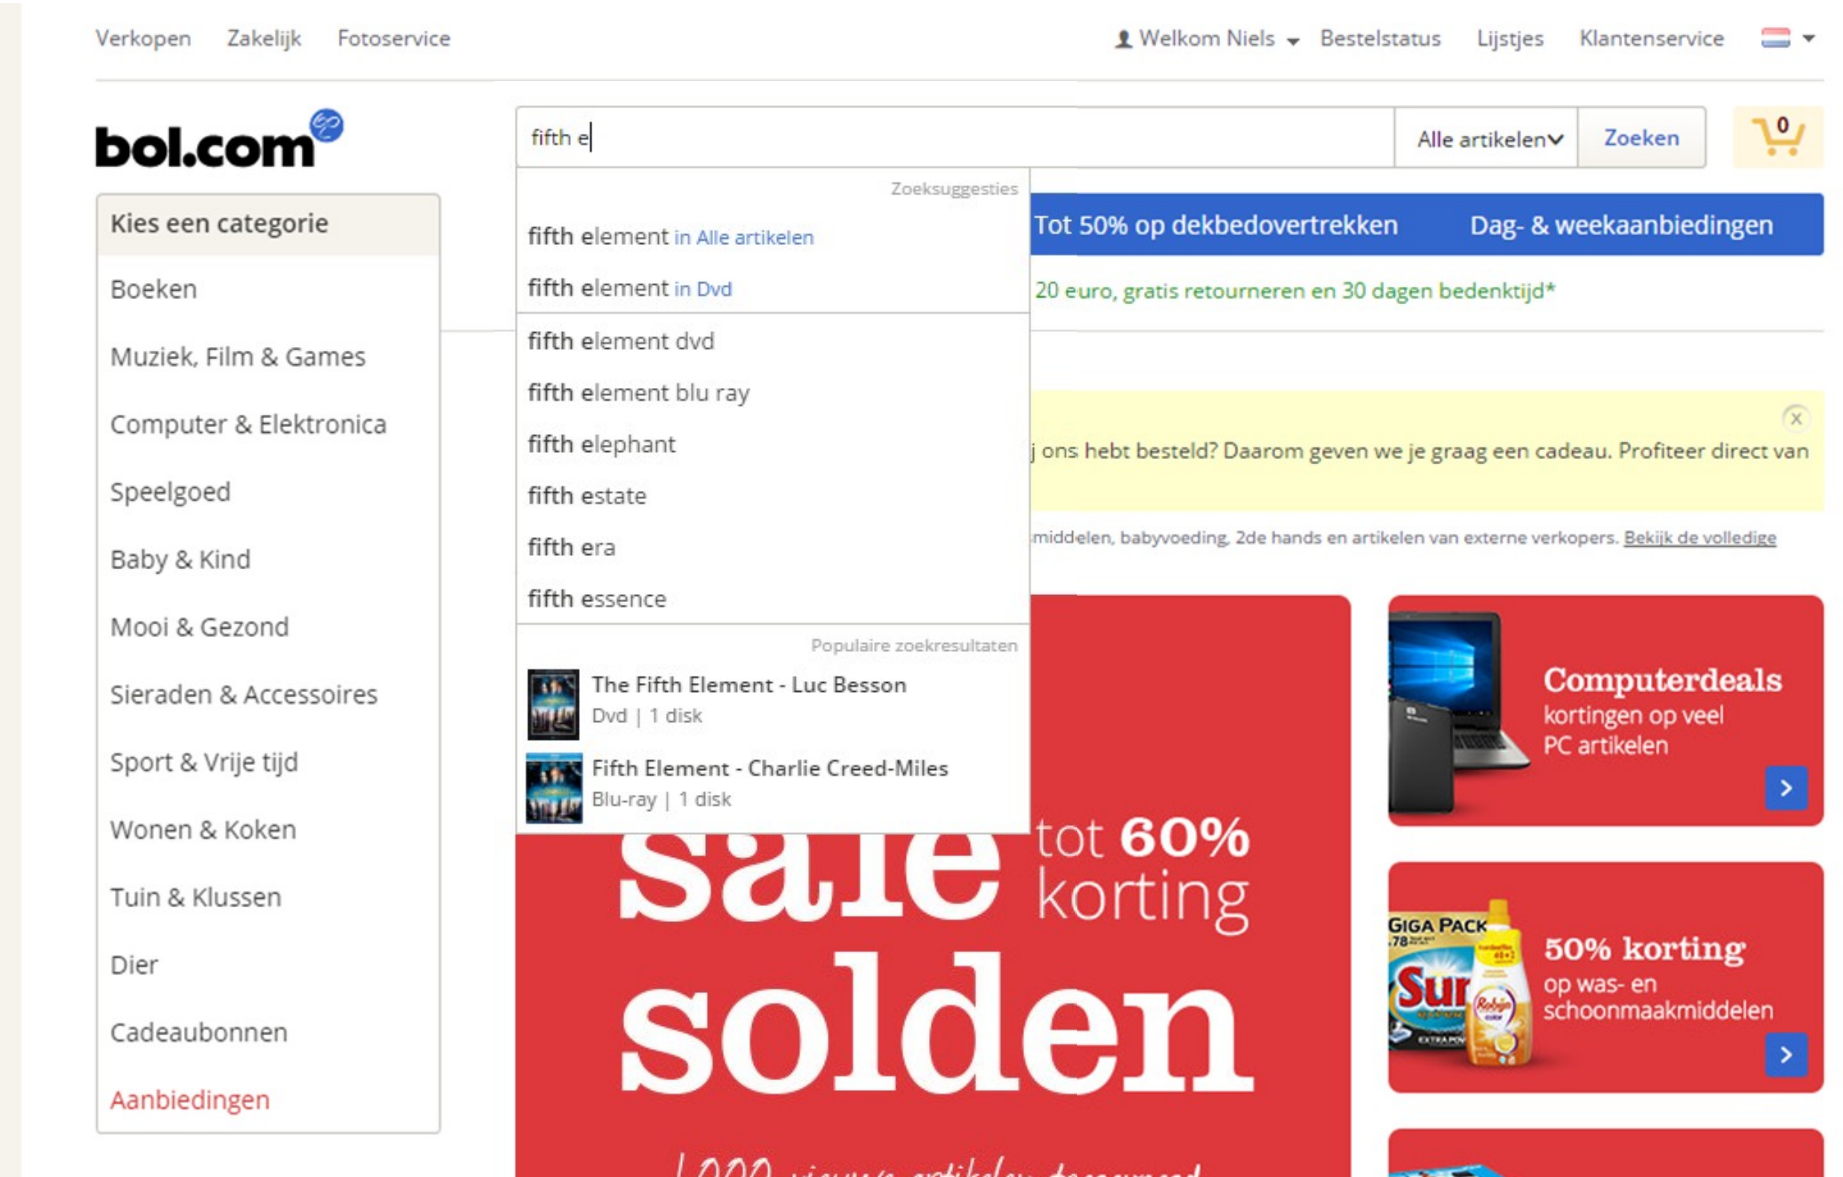


Personalization
Long term history
Browsing history
Purchase history
Wishlist
Browsing history
Verkopen
Zakelijk
Fotoservice
Welkom Niels
Bestelstatus
Lijstjes
Klantenservice
bol.com
Alle artikelen
Zoeken
Kies een categorie
Boeken
Muziek, Film & Games
Computer & Elektronica
Speelgoed
Baby & Kind
Sieraden & Accessoires
Sport & Vrije tijd
Wonen & Koken
Tuin & Klussen
Dier
Cadeaubonnen
Aanbiedingen
Frisse Start
Beste van 2015
Dag- & weekaanbiedingen
Gratis verzending vanaf 20 euro, gratis retourneren en
Beste Niels,
€ 7,50 korting bij bol.com.
sale
tot 60% korting
solden
Computerdeals
PC artikelen
50% korting
op was- en schoonmaakmiddelen

Based on Measuring 1.0

We have some issues here...

bol.com

Measurements are…

- **too old.**
  - Available once every 24 hours.
  - So personalization is a ‘day behind’.

Useless inspiration:
I was interested in this YESTERDAY

bol.com

Data relevance decay

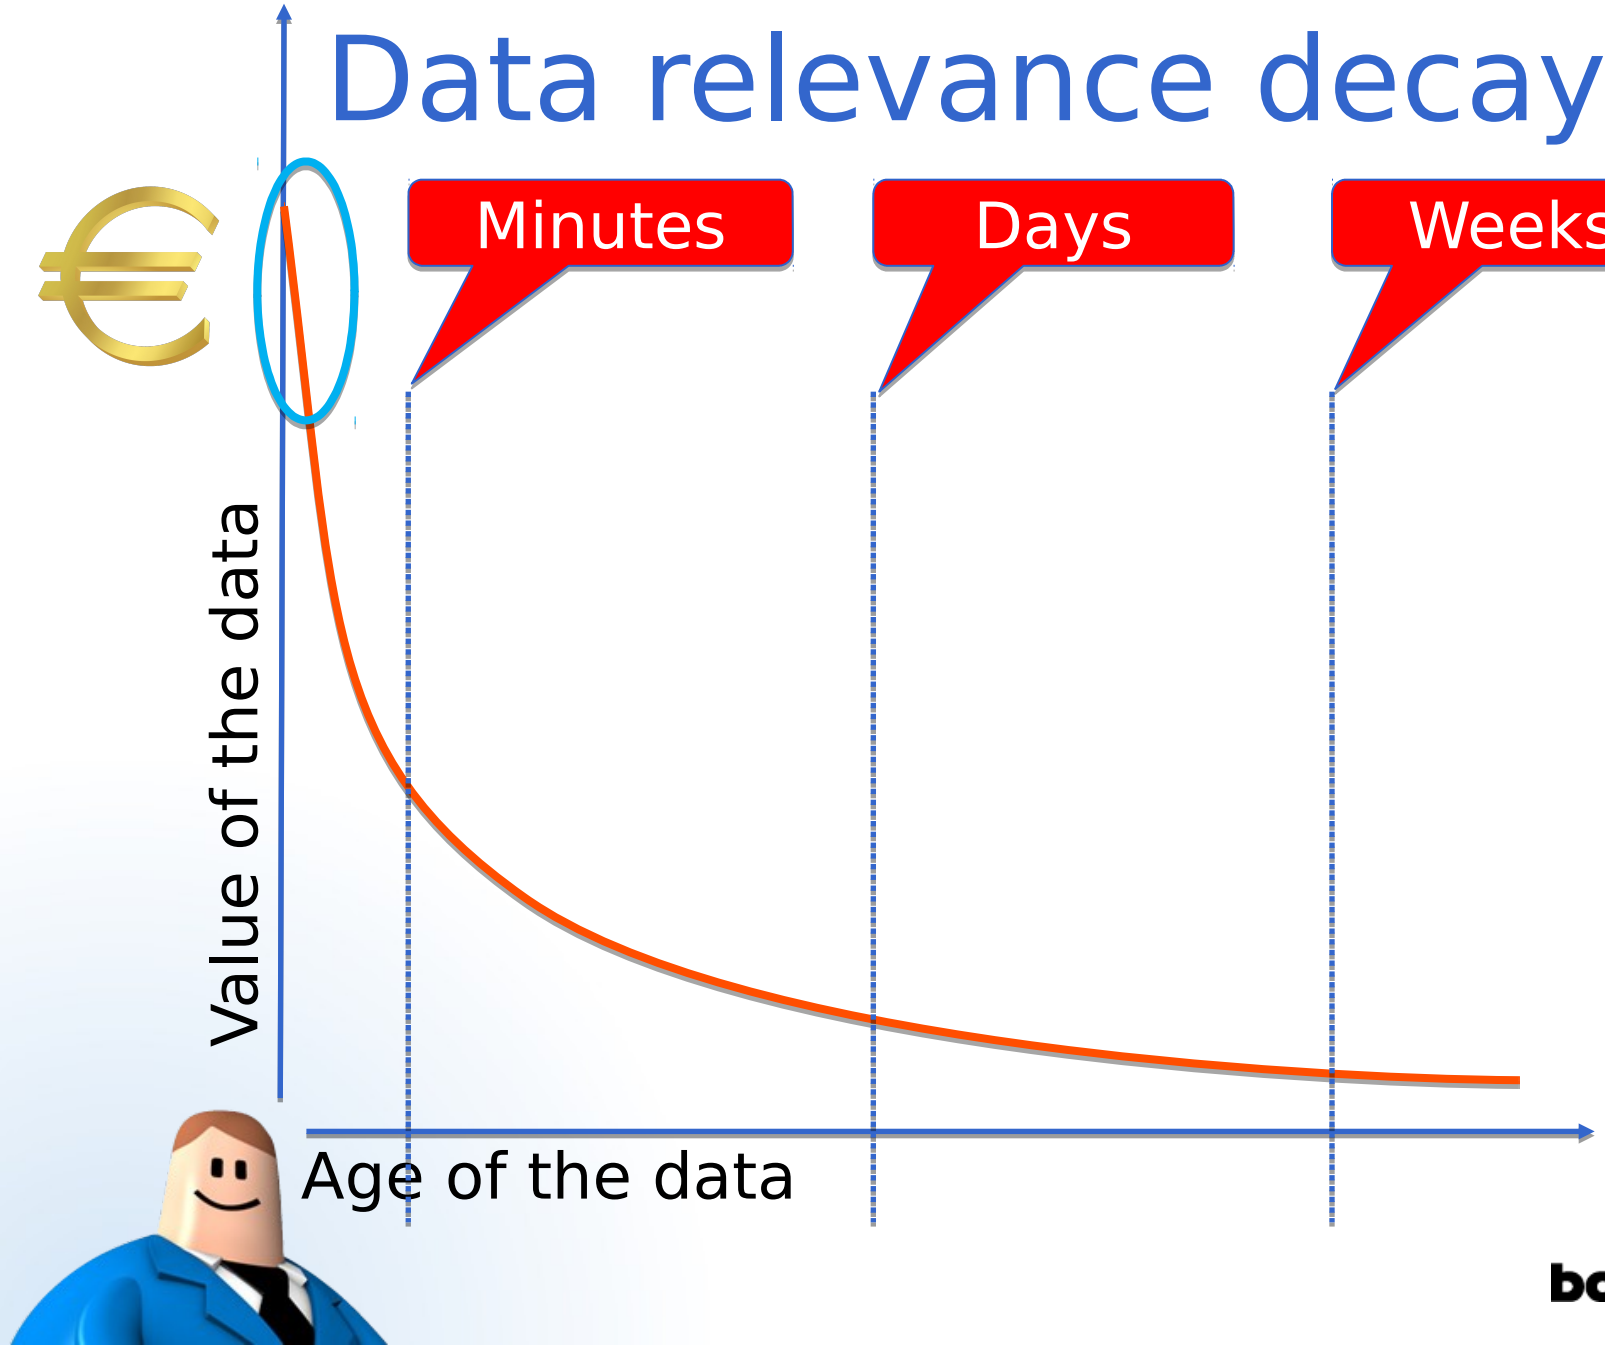

bol.com

Measurements are...

done using JavaScript ...

```
<script type="text/javascript">
   if(typeof(s)=='undefined'){s = {};s.t = function(){};}

var link = document.querySelector && document.querySelector('link[rel=
link && link.href; if (canonical) { s.prop3 = canonical; s.eVar26 = '[
s.eVar4=unescape('ps43');
s.pageName=unescape('\/nl\/index.html');
s.eVar3=unescape('main');
s.prop2=unescape('Home');
s.prop18=unescape('0');
s.prop1=unescape('main:algemeen');
s.prop29=unescape('VqeYkQpid0IAAELxnY8AAAb4');
s.eVar12=unescape('not logged in');
s.prop25=unescape('niet ssl');
s.prop50=unescape('\/nl\/');
s.eVar51=unescape('www.bol.com,DESKTOP');
```

Measurements are…

- **Heavy on clients**
  - Everything is javascript
- **We cannot measure “Everything”.**
  - Sometimes we run out of eVars.
  - No personalized banners.
- **Unclear/unspecified/complicated.**
  - Errors because eVars get ‘reused’ over time.

bol.com

Measurements are...

- **Incomplete/Inaccurate.**
  - JavaScript is sometimes slow
  - Everything goes on the URL

s.products=unescape(';1001004007516421;;;;evar3=books|evar7=1001004007516421_Pluk van de Petteflet|evar35= 171|evar41=KI1200|evar47=B|evar9=New');

bol.com

All special deals page

http://www.bol.com/nl/aanbiedingen/index.html

De allerbeste aanbiedingen!

Vind ik leuk 553 Tweeten 18 Volgen

Dagaanbieding!

› Ontvang elke dag de dagaanbieding in je mail

Je hebt nog 13 uur, 36 minuten, 43 seconden

Aquaplay draagbare Waterbaan AquaLock 616 - Waterbaan

Prachtige draagbare waterbaan van Aquaplay. Dirigeer jouw bootjes door de kanalen en sluizen. Leer spelenderwijs hoe water beweegt.

~~€ 49,99~~

€ 29,99

Je bespaart: 40%!

Vandaag voor 23:00 uur besteld, morgen in huis.

+ In winkelwagentje

(7 reviews) › Lees reviews

Waterbaan | 88 cm | 3 - 6 jaar | 3 - 6 jaar | Waterbaan

De draagbare AquaLock 616 is een van de bestsellers van AquaPlay. Neem de draagbare AquaLock 616 waterwereld

Weekaanbiedingen!

Je hebt nog 4 dagen, 13 uur, 36 minuten, 43 seconden

34 weekaanbiedingen

Dead Or Alive (3DS)

(3 reviews)

Nintendo 3DS

Vandaag voor 23:00 uur besteld, morgen in huis.

~~€ 29,99~~

€ 9,99

Je bespaart: 67%!

+

Electrolux UltraCaptic Stofzuiger

Zakloos | Hepa | 1400 W | Stofzuiger | 1400 W

Vandaag voor 23:00 uur besteld, morgen in huis.

~~€ 379,00~~

€ 179,00

Je bespaart: 53%!

+

Sam Smith - In The Lonely

~~€ 17,99~~

All special deals page

s.products=unescape(';1004004011338840;;;;evar3=|evar8=1004004011338840_Dead or Alive|evar35=170|evar47=EGAME|evar9=New |evar27=not scarcity,;9200000018376017;;;;evar3=|evar8=9200000018376017_Electrolux UltraCaptic UCALLFLOOR Stofzuiger|evar35=160|evar47=HH|evar9=New|evar27=not scarcity,;9200000023965194;;;;evar3=|evar8=9200000023965194_In The Lonely Hour (Deluxe Edition)|evar35=174|evar47=POP|evar9=New|evar27=not scarcity,;9200000030746171;;;;evar3=|evar8=9200000030746171_Billendoekjes Lotion - 24x72 stuks|evar35=160|evar47=BA evar9=New|evar27=not shown,;9200000010313916;;;;evar3=|evar8=9200000010313916_Philips SB7200 - Bluetooth-speaker - Zwart|evar35= evar47=PERAU|evar9=New|evar27=not shown,;9200000010695 cinema se 35=183|evar47=HNT|evar9=New|evar27=not scarcity,;9200 ogy|evar3 var47=DVDMO|evar9=New|evar27=not scarcity,;920000002 r47=EBOOK =New|evar27=not scarcity,;9200000009380504;;;;evar3 skoolbarb 0D8 47cm - Groen|evar35=160|evar47=KTH|evar9=New| 8319_Acer Duo - Dual Sim - Zwart| 35=160|evar47=MOBPHON r8=920000 illette ProGlide Manual Scheerapparaat + 9 stuk , 07_Nautica Yacht Cl a r8=9200000024802093 r 3=|evar8=9200000 or|evar35=160|evar47=BPH|evar9=New|evar27=not shown ;9200000026165 63 47=B|evar9=New|evar27=not scarcity,;9200000020338 ;;;;evar3=|evar8 Bronzer|evar35=170|evar47=BPH|evar9=New|evar27=not shown,;9200000008 All-in-One Printer|evar35=160|evar47=PCACCE|evar9=New|evar27=not sca s Senseo Viva Caf\u00E9 HD7825\/70 - Blauw|evar35=160|evar47=KIT|eva 00000020099724_Just One Evil Act|evar35=170|evar47=UK|evar9=New|eva 394_Katy Perry Killer Queen for Women - 30 ml - Eau de Parf ;evar3=|evar8=9200000010152983_Pelgrim PKV085WIT Tafelmodel 002756012;;;;evar3=|evar8=9200000002756012_Curver Decobin P t shown,;9200000010562035;;;;evar3=|evar8=9200000010562035_ evar9=New|evar27=not scarcity,;9200000026525667;;;;evar3=|evar8=9200 r47 lie - 60 Ca psu - Zwart|ev ar3 wk - Backpa ck 9063_D-Link DCS-5222L - Cloud Camera 5000 - Netwerkcamera|evar35=160|evar47=PCACCE|evar9=New|evar27=not scarcity,;9000026073540;;;;evar3=|evar8=92 40_Berthoud Green power- Elektrische onkruidbestrijder|evar35=160|evar47=PGT|evar9=N ar27=not scarcity,;9200 0000107 3=|evar8=9200000010740892_Nutribird P15 Tropical Onderhoudsvoeder Papegaaienvoer - ar35=160|evar47=PGT|e var9=Ne scarcity,;9200000020478011;;;;evar3=|evar8=9200000020478011_Iiyama ProLite GB2773 nitor|evar35=173|ev ar47=NTB evar27=not scarcity,;9200000021950688;;;;evar3=|evar8=9200000021950688_Packard Be 69KB-1145NL8.1nO - Laptop ar35=160|evar47=NTB|evar9=New|evar27=not scarcity,;920000 kup Plus Slim 1T USB 3.0 Mobiele harde schijf (zwart)|evar35=170|evar47=DAT evar3=|e var8=92 000027266545_Packard Bell iMedia S A3621 - Desktop|evar35=1 550908;; ;;evar3=|evar8=1004004013550908_Aquaplay draagbare Waterbaan AquaLock 616 - Waterbaan|evar35=160|evar47=TOYS|evar9=New|evar27=not shown')

s.products=unesca pe('

');

bolpx01.gif?AQB=1&ndh=1&t=7%2F7%2F2014%20...
/tracking/s45251220294740

GET

414
Request-URI Too Long

Page is invisible for stats!

So...

things need to change

bol.com

Business goal

We want to be able to know and process everything about our webshop

NOW

Because we want to be able to assist our visitors

NOW

bol.com

Business goal

- **While the visitor is still on the site**
  - Respond to what they have done so far.
    - Last months/year
    - Within seconds after the ‘click’.
- **Batches multiple times per day**
  - Search suggestions
  - Search rank
  - Recommendations

bol.com

Measuring 2.0

bol.com

It’s really all about…

- **Measuring**
  - Better
- **Processing**
  - Faster
- **Applications**
  - More relevant

bol.com

Goals of “Measuring 2.0”

- **Measure everything of our website**
  - All interactions (also AJAX)
  - All channels (also mobile, email, ..
  - All countries
  - All details
  - All visitors (also Googlebot)
- **More reliable data**
- **Lowest possible load on the client**
- **Lowest possible latency (< 1 second)**

bol.com

Goals of “Measuring 2.0”

- **Developer**
  - Easy to build
  - Easy to validate
  - Test automation
- **Business**
  - Always measure everything
  - Data is “independent”
    - New questions are allowed

bol.com

Goals of “Measuring 2.0”

- **Privacy**
  - Personal Data Protection Act
    - *Wet bescherming persoonsgegevens*
  - No profiling beyond 2 years
- **Security**
  - Avoid storing “login” info.
- **Business**
  - Do “profiling” on everything
  - For many years (>2)

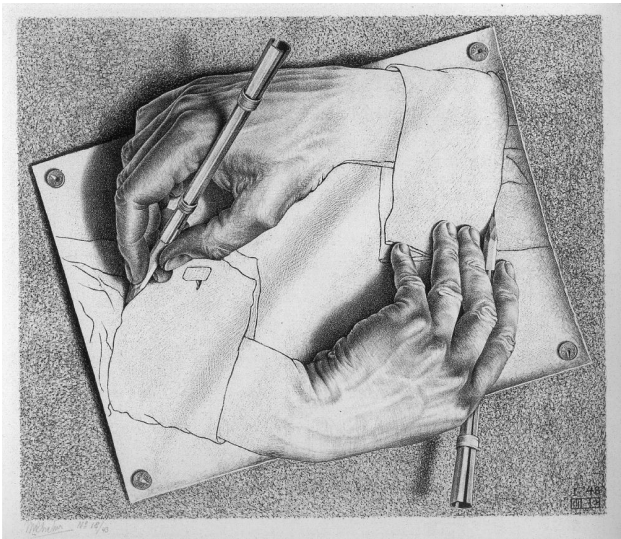

Measure everything

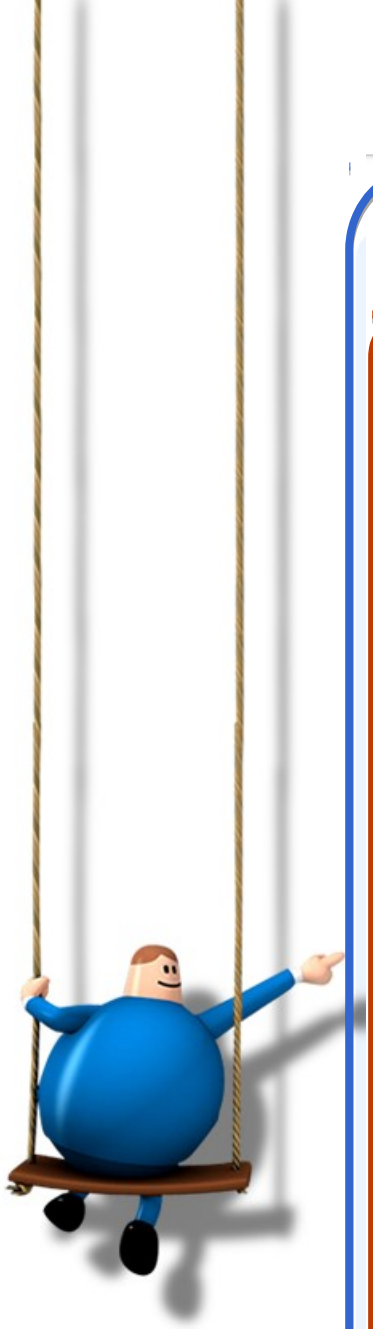

bol.com

Boeken > Nederlandse boeken > Kind & Jeugd > Voorlees- & Sprookjesboeken > Pluk van de Petteflet

Pluk van de Petteflet

Herziene Editie.Nu Met Nieuwe Illustraties

Auteur: Annie M.G. Schmidt | 39 reviews | Stel een vraag

Favoriet 8

Inkijkexemplaar

Bindwijze: **Hardcover**

~~€ 19,95~~ **16,96**

Je bespaart 15%

Vandaag voor 23:00 uur besteld, morgen in huis

Prijs en levertijd

Verkoop door **bol.com**

- Gratis verzending vanaf 20 euro
- 30 dagen bedenktijd en gratis retourneren
- Ophalen bij een bol.com afhaalpunt mogelijk
- Dag en nacht klantenservice
- Cadeautje? Laat het voor je inpakken en bezorgen

+ In winkelwagentje Zet op verlanglijstje

Andere verkopers (11)

Tweedehands
vanaf € 11,95

› Bekijk en vergelijk alle verkopers
vanaf € 16,96

Samenvatting

Pluk had een klein rood kraanwagentje. Hij reed ermee door de hele stad en zocht naar een huis om in te wonen. Af en toe stopte hij. En dan vroeg hij aan de mensen: 'Weet u niet een huis voor me?'

Uiteindelijk vindt hij wat: het torentje van de Petteflet. Daar maakt hij kennis met de Stampertjes, mevrouw Helderder, Aagje en Zaza, maar ook met andere dieren als Dollie, Langhors en de Krullevaar kan hij het goed

Annie M.G. Schmidt
Auteur

Blijf op de hoogte

Co-auteur: Fiep Westendorp

Uitgever: Querido Kinderboek

- Nederlands
- 200 pagina's
- Querido Kinderboek
- april 2010

› Alle productspecificaties

Anderen bekeken ook:

Jip en Janneke
Annie M.G. Schmidt
€ 25,95

Floddertje
Fiep Westendorp
€ 15,95

Pluk redt de dieren
Fiep Westendorp
€ 15,95

Otje
Querido Kinderboek
€ 17,50

Cause and Effect

is what really interests us

bol.com

Our usecases

- **Banner optimization**
  - Look / Search Next page better banner
- **A/B testing**
  - Show feature Use Buy product
- **Search Suggestions**
  - Search Find Choose Buy product
- **Attribution modeling**
  - Show ad Click Buy product
- **…**

bol.com

Behavioral analytics

- **Cause and effect**
  - **Action:** We show something
  - **Reaction:** To click or not to click

bol.com

The “measurement”

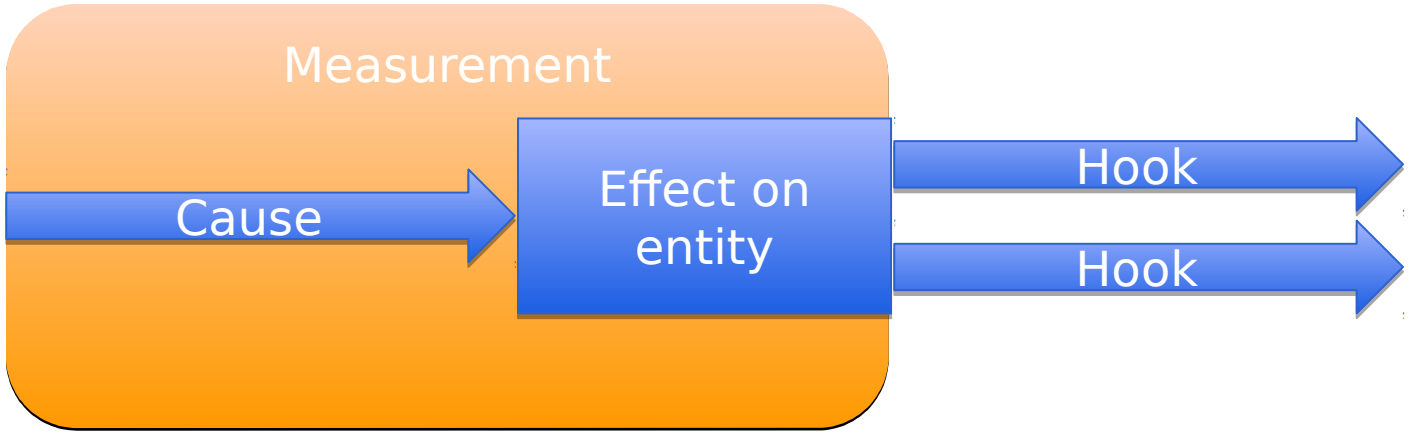

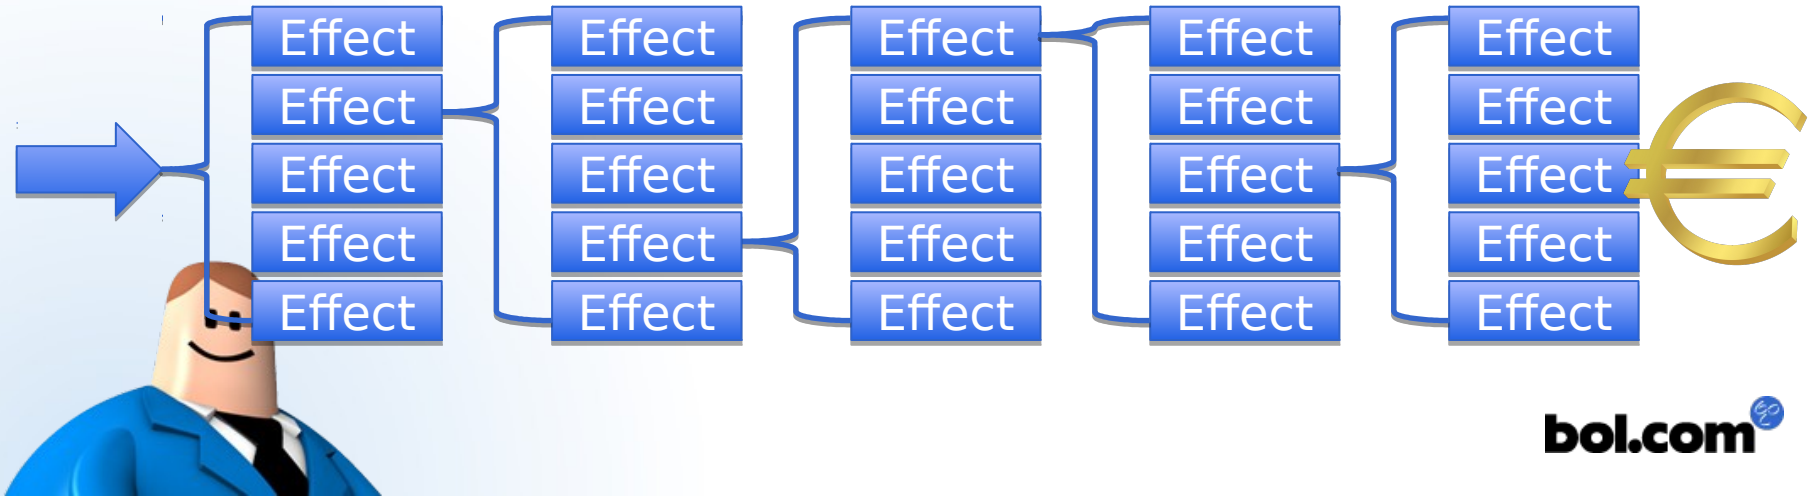

bol.com

How do we measure?

How can we measure?

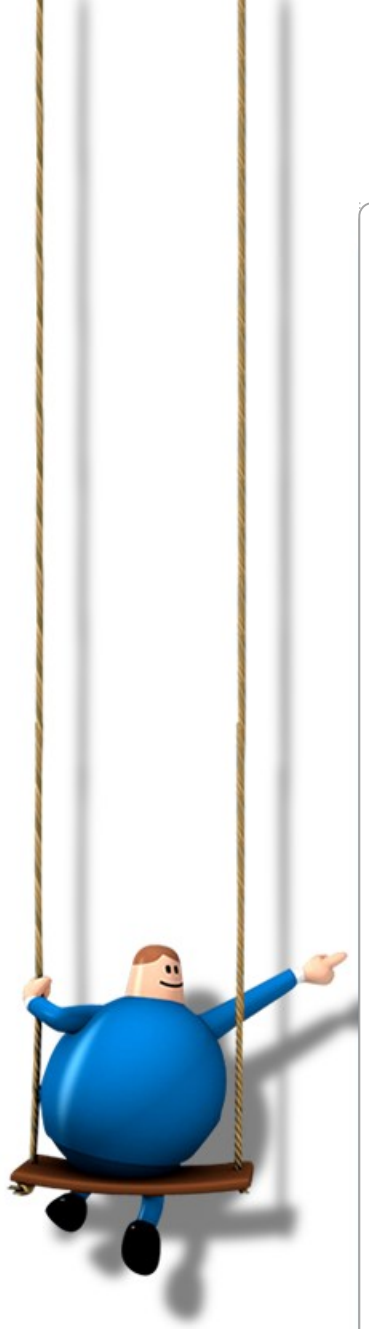

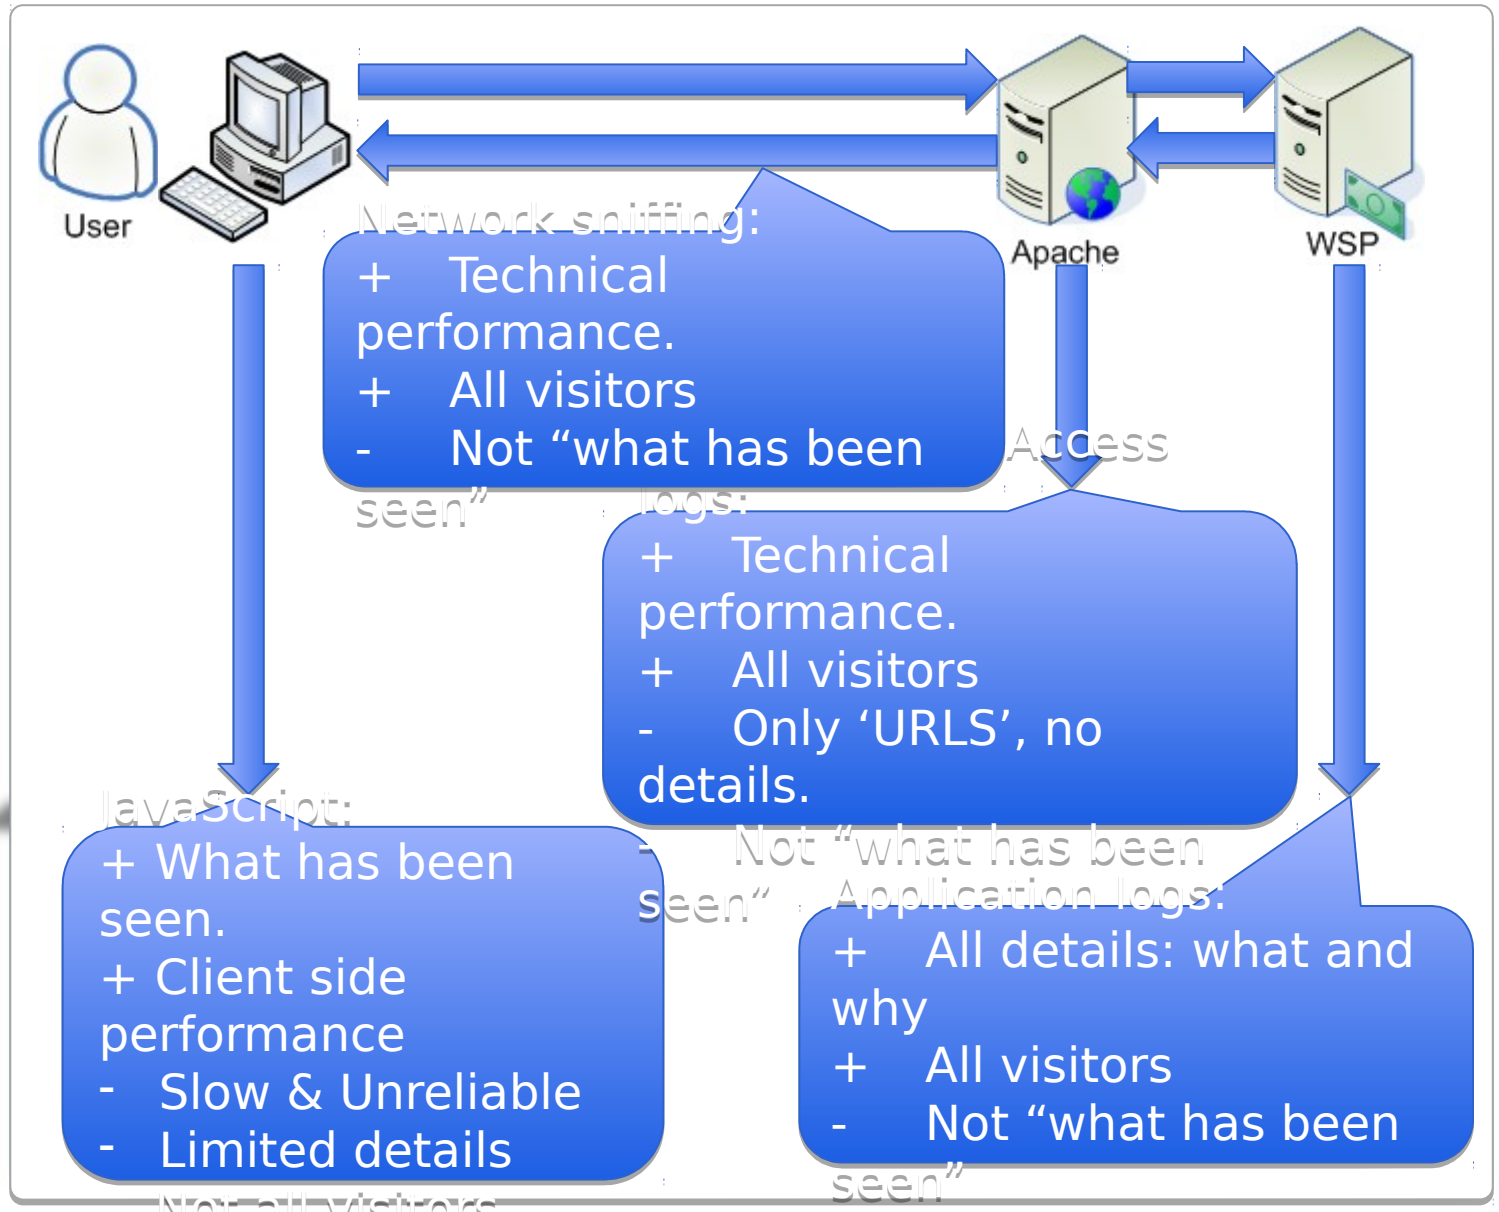

bol.com

Hybrid approach

We measure everything serverside

(unless it really cannot be measured that way)

If we need to measure clientside then it must be as lightweight as possible.

JavaScript:
+ What has been seen.
+ Client side performance
- Slow & Unreliable
- Limited details
- Not all visitors

Application logs:
+ All details: what and why
+ All visitors
- Not “what has been seen”

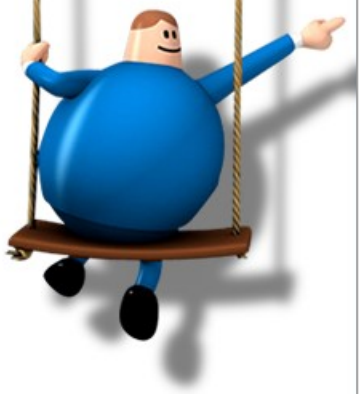

bol.com

Processing the data

bol.com

Requirements

- **Online:**
  - Near real-time (< 1 second)
  - Have long history available
  - Seamless integration history + real-time
- **Offline:**
  - Incremental batch jobs during the day
  - Have long history available

bol.com

Lambda Architecture

- **Basic idea:**
  - Batch tools are good at ‘a lot of data’
  - Realtime tools are good at ‘low latency’
- **Lambda Architecture:**
  - Combine the two at query time.

Cut boundary at a millisecond.

bol.com

Many batches per day

Joining the pieces becomes very hard

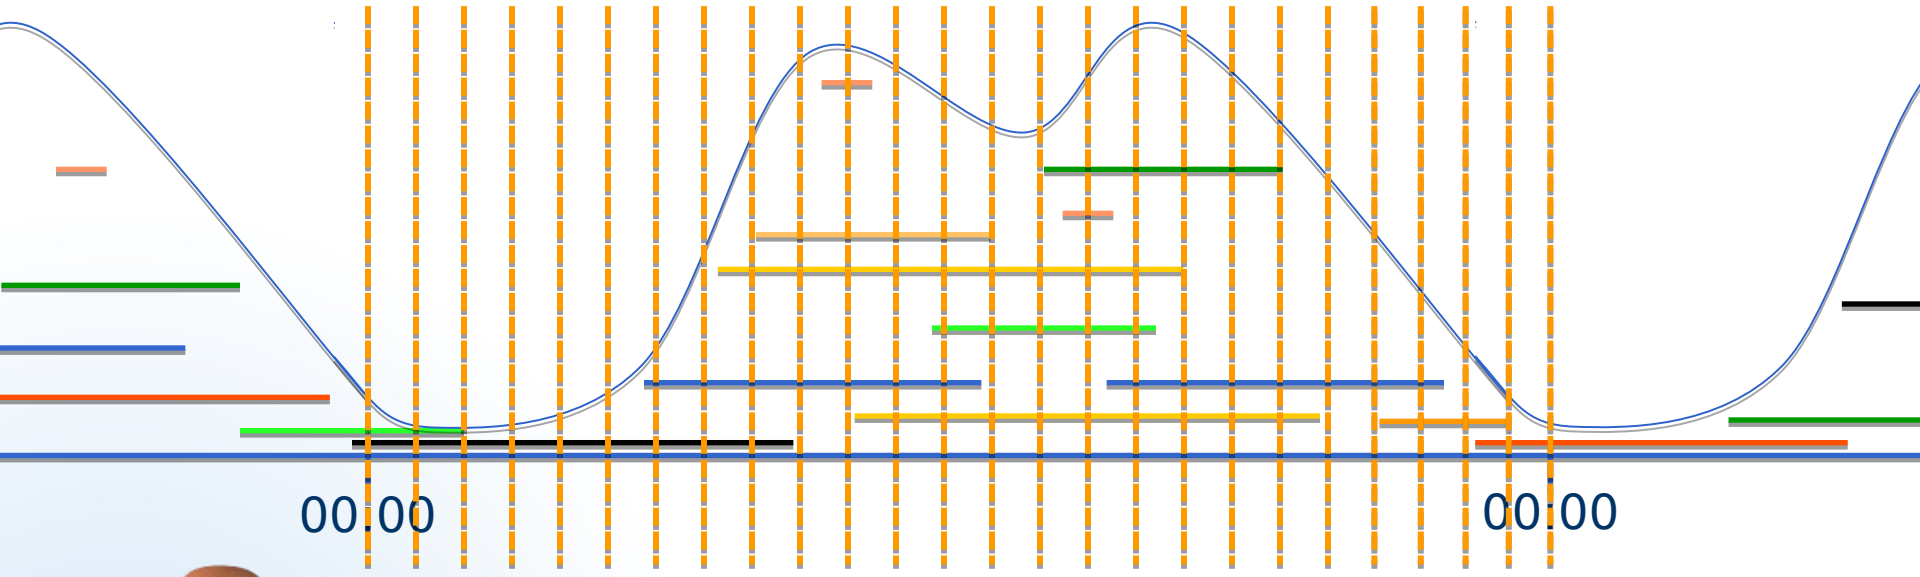

bol.com

Sessionized Lambda Architecture

by Niels Basjes (Bol.com)

bol.com

Sessionized Lambda Architecture

Extension of the Lambda Architecture

- **Bounded event streams**
  - that stay together
- **Queries take time**
  - and multiple can overlap
- **Service orientation**
  - because it is part of a bigger thing

bol.com

Focus on Visits

- **Browsers**
  - The software installed on a computer
- **Sessions**
  - Start: Visit website
  - End: Close browser
  - Can last for days … weeks … !
    - Device is ‘suspended’
- **Visits**
  - Start: Visit website
  - End:
    - 30 minutes idle
    - max 12 hours active

bol.com

Processing Architecture

bol.com

Measuring 2.0

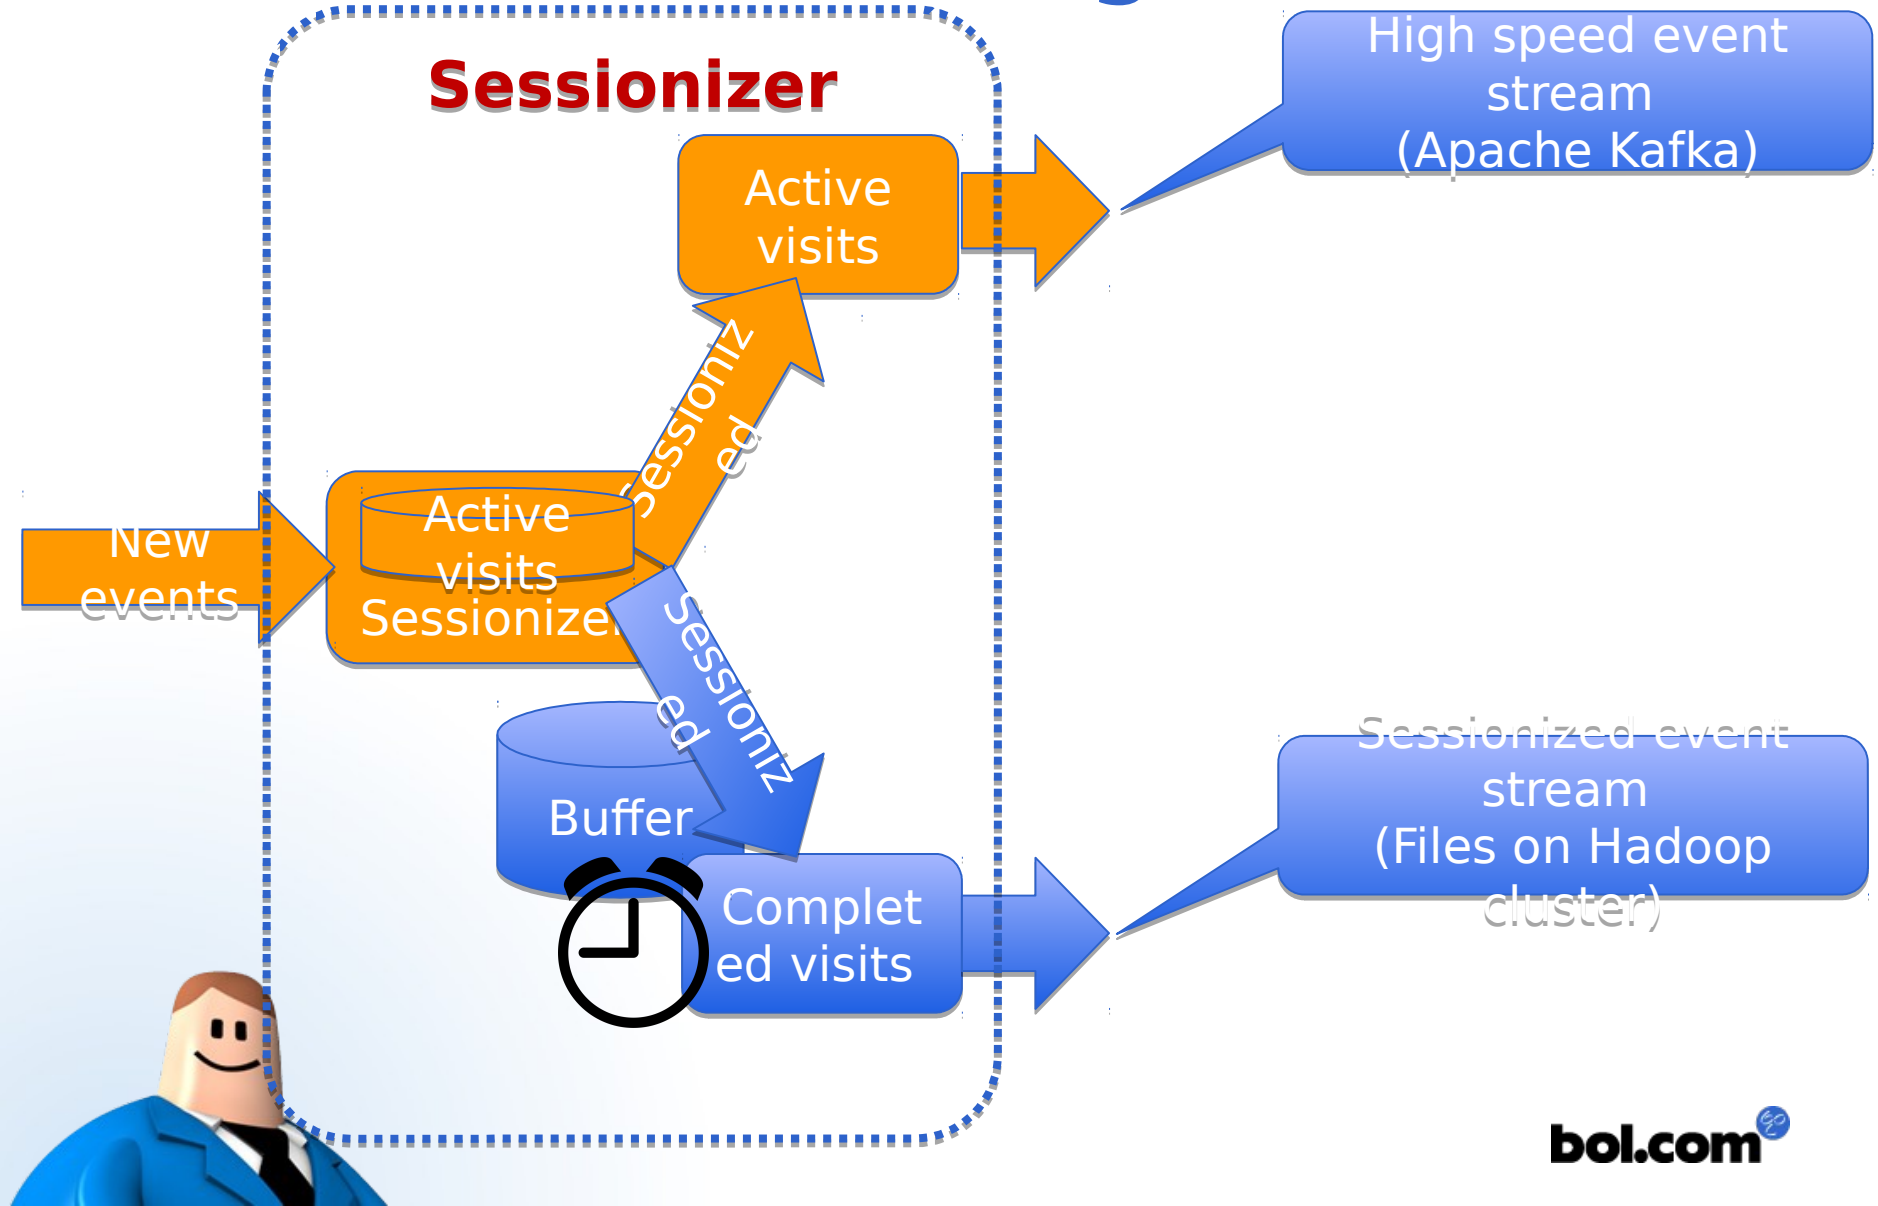

bol.com

“Full” Sessionized Lambda

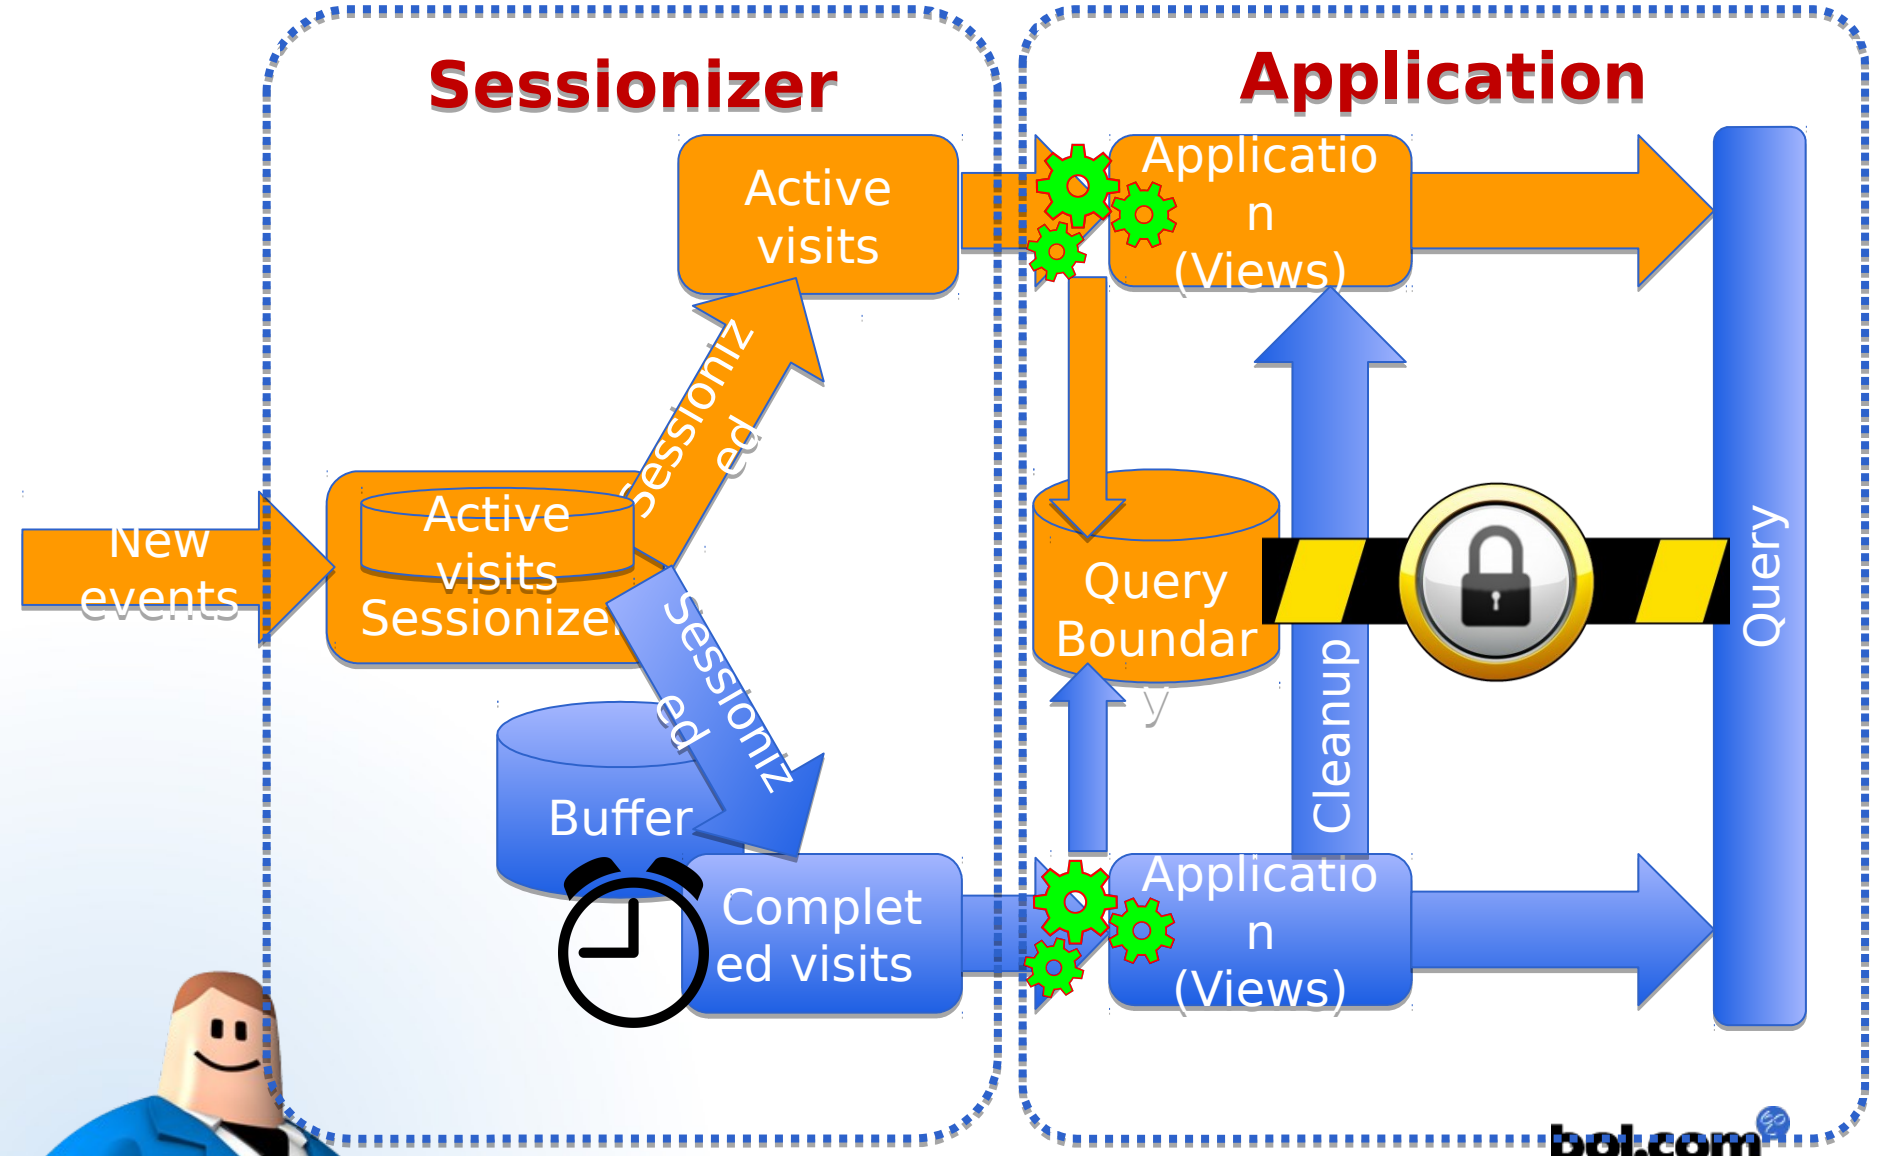

bol.com

“Speed only” application

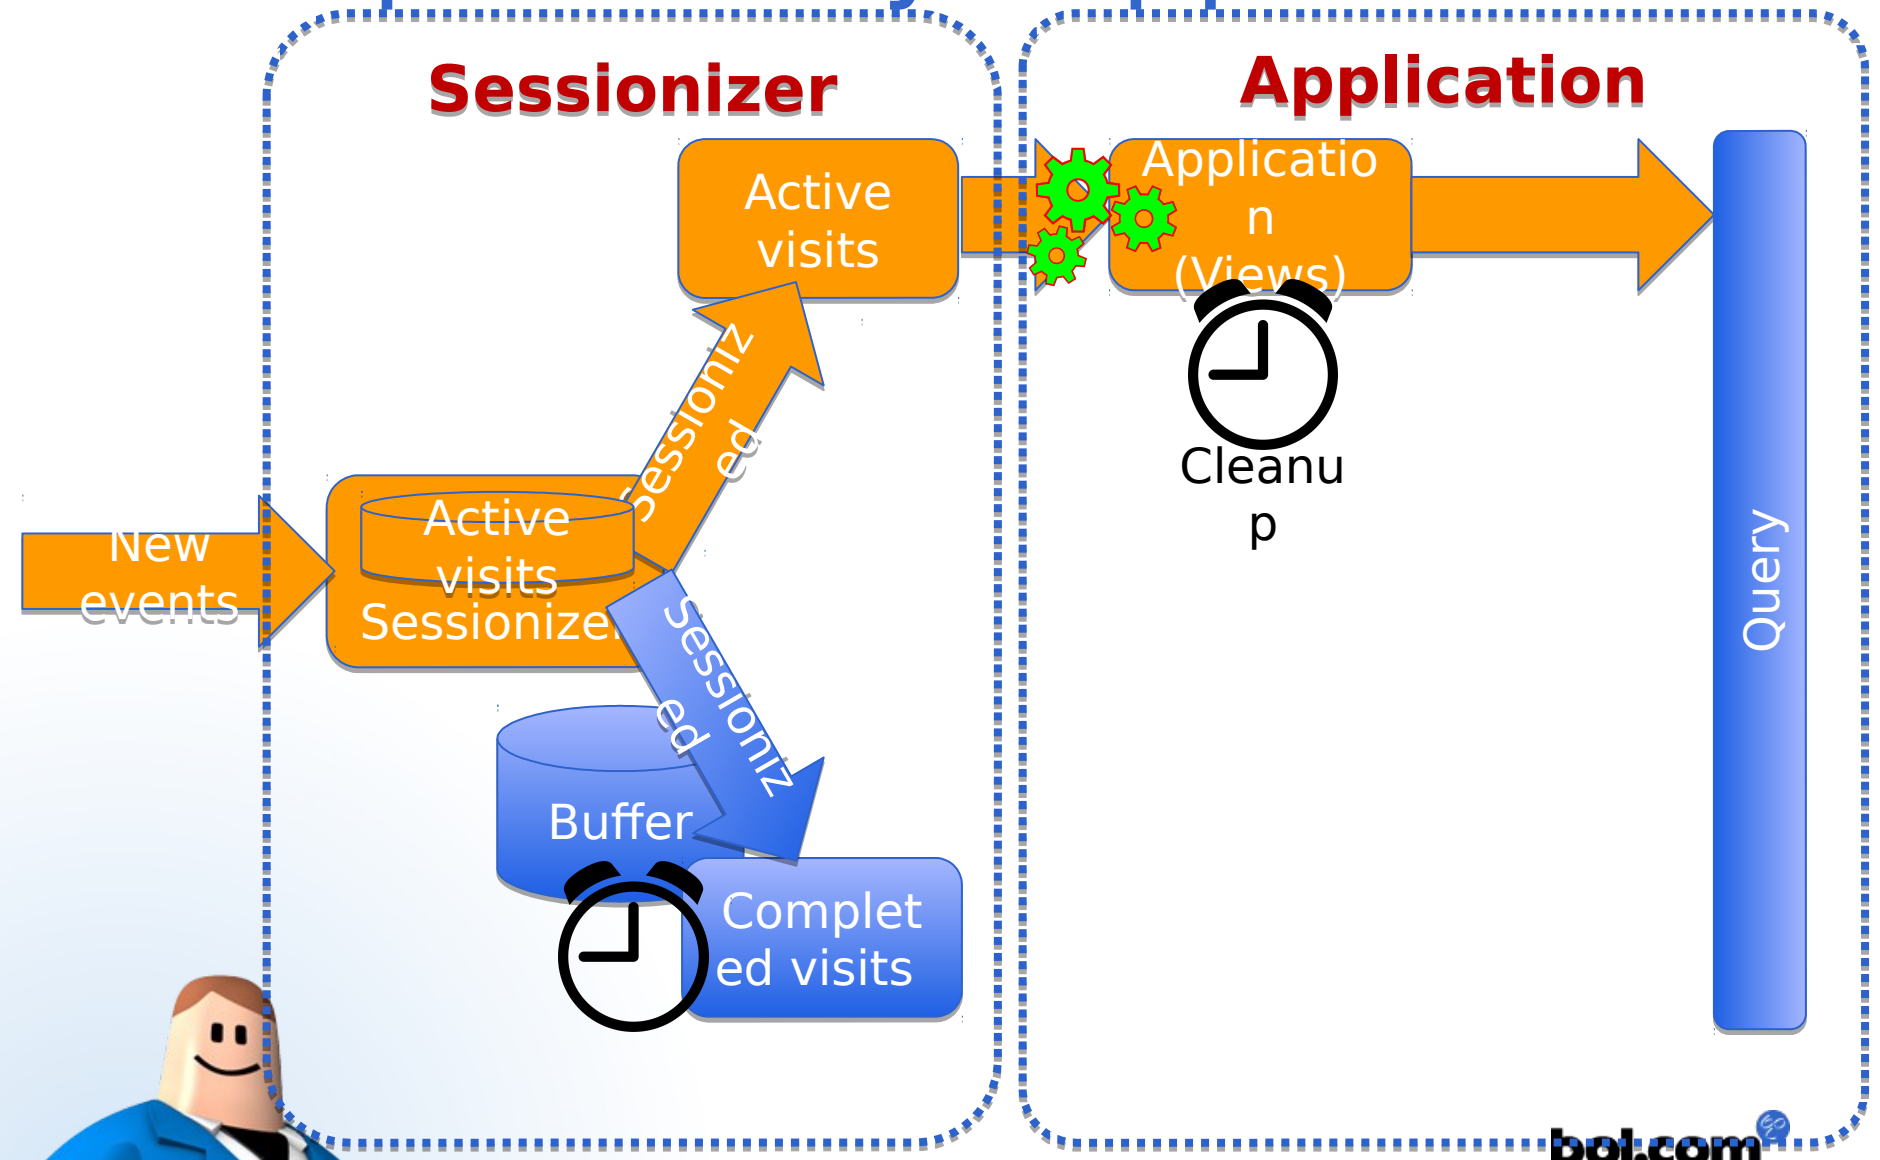

“Simple hybrid” application

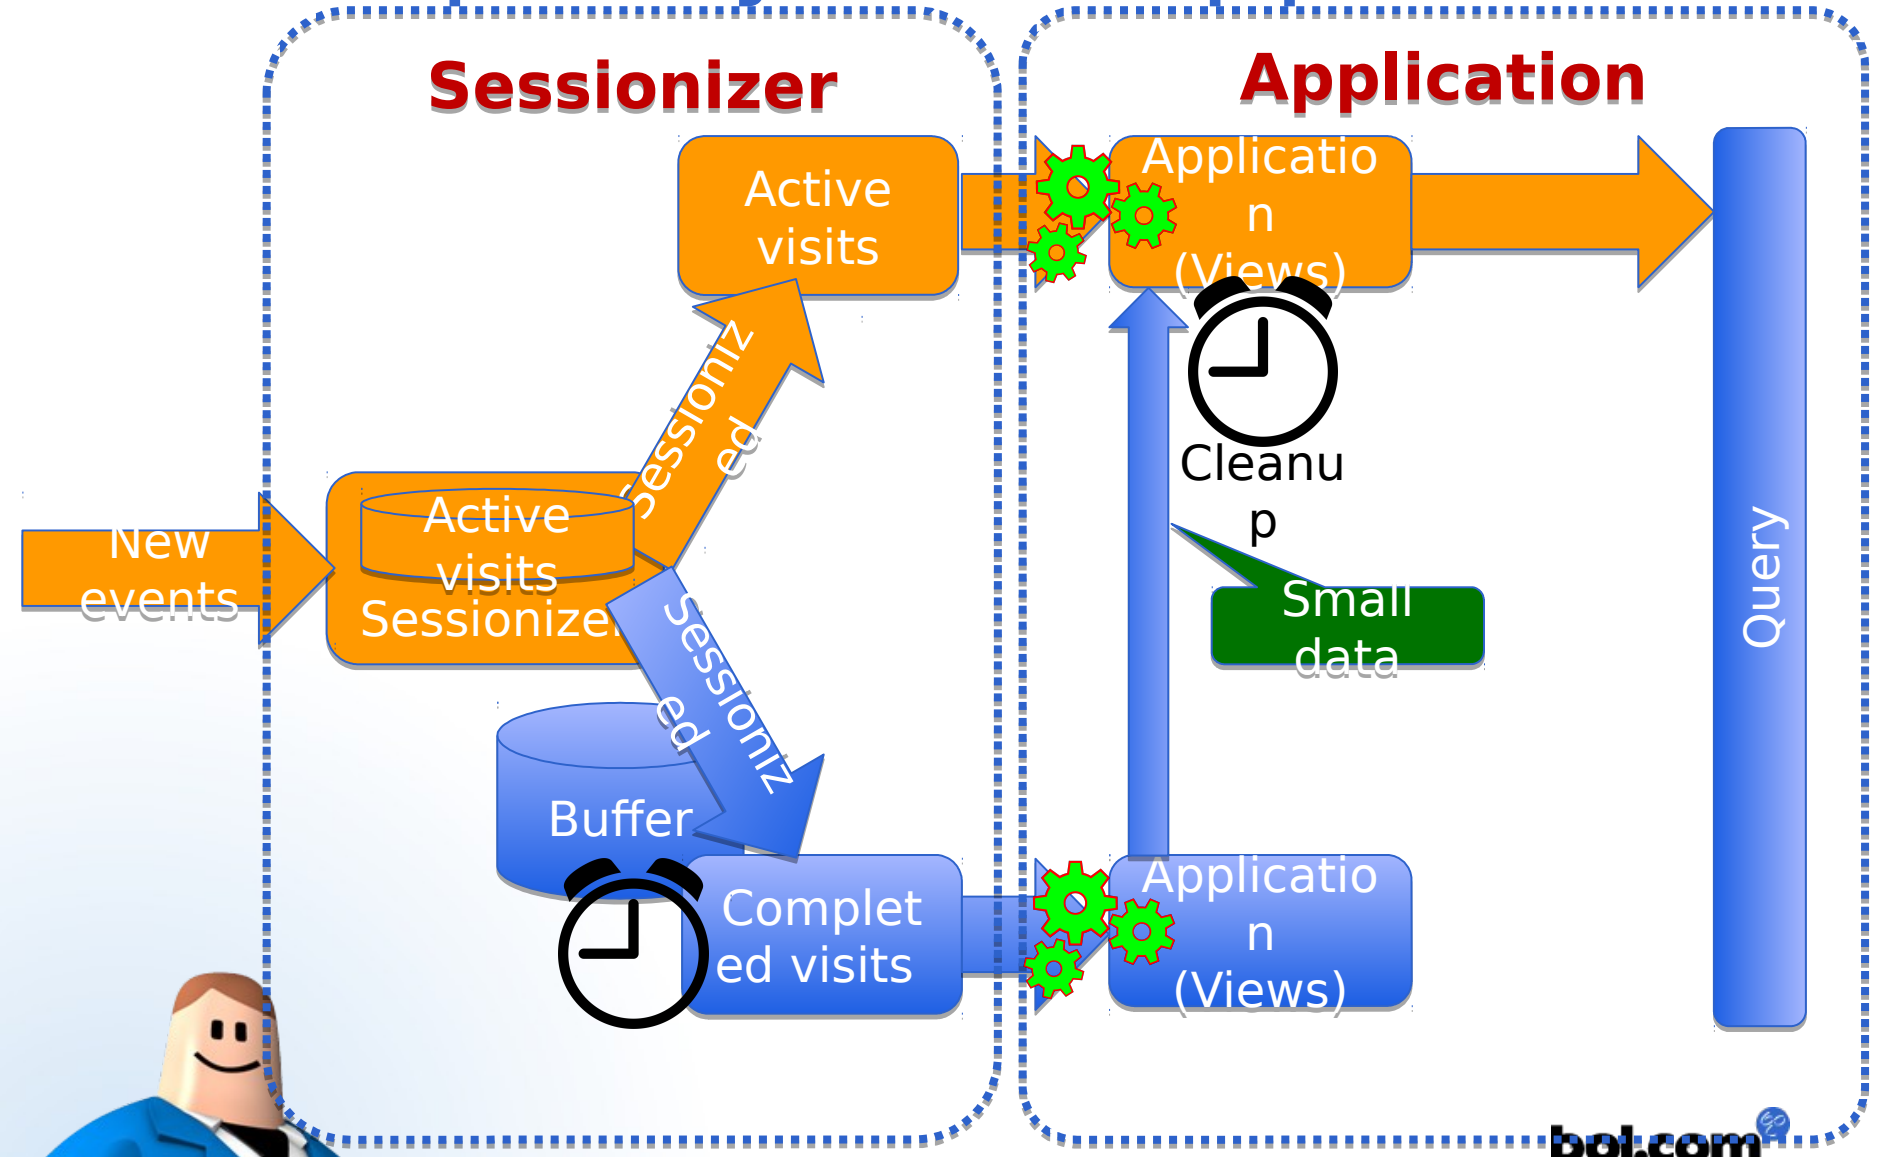

bol.com

“Batch only” application

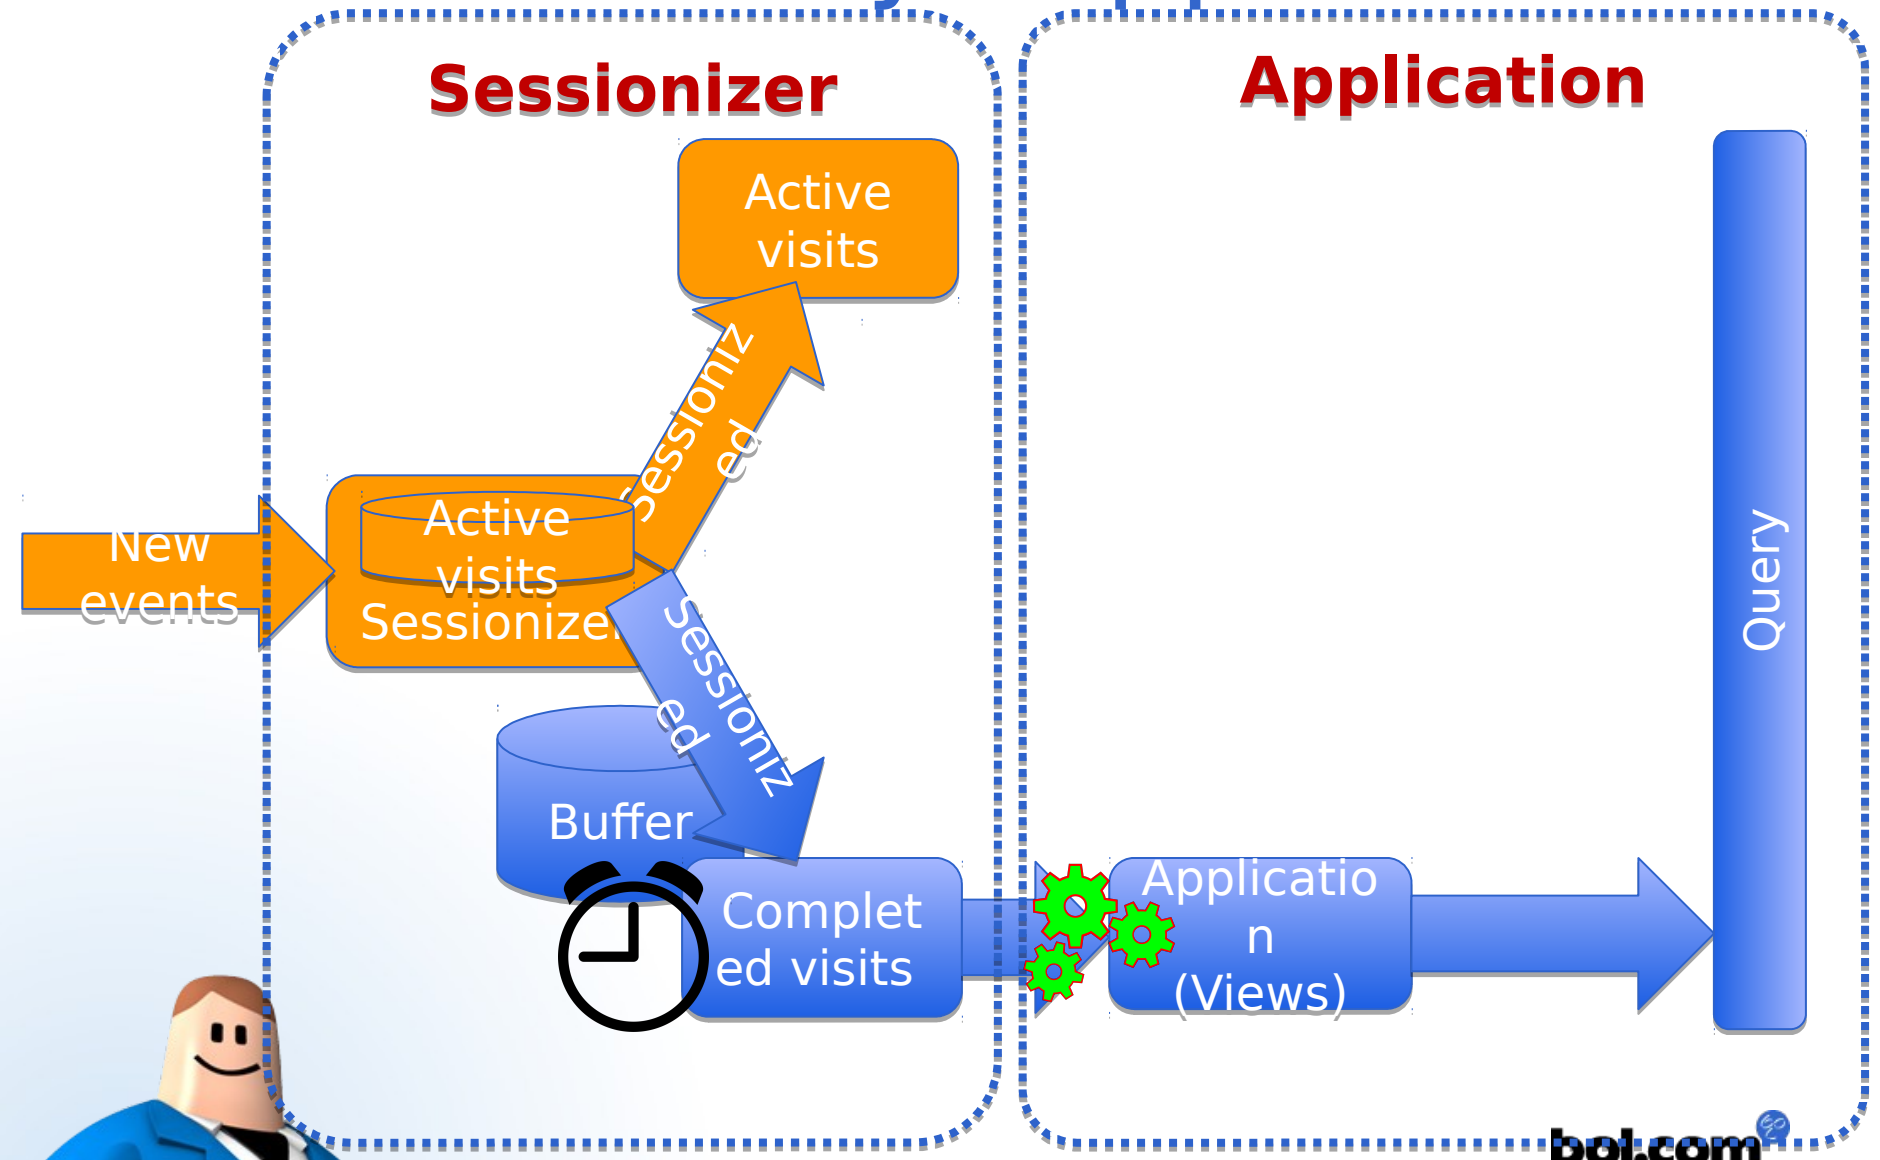

bol.com

Privacy vs Profiling

bol.com

Privacy protection

Identifiable data may not be used for profiling for more than 2 years *)

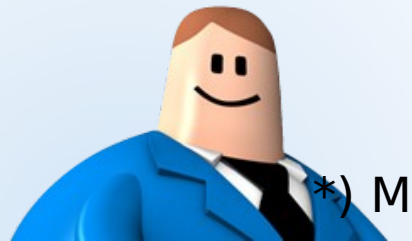

*) My wording & understanding, I’m no lawyer.

bol.com

Identifiable data

(Groups of) data elements that can identify a single person

In the clickstream:

- **IP Address**
- **Customer number**
  - *After logging in*
- **Unique browser ID**
  - Random value in cookie
    - *Visitor can clean this*

bol.com

Long term “profiling”?

“Looking for behavior patterns”

- **Query pattern:**
  - GROUP BY ${CustomerId}
- **Same result**
  - GROUP BY **ENCRYPT**(${CustomerId})
  - GROUP BY **HASH**(${CustomerId})
  - GROUP BY **HASH**(${CustomerId}-**SALTED**)

Hash collissions are not important

bol.com

Our solution

- **Data in motion (streaming)**
  - Identifiable & Anonymous
    - Kafka expires after 4 weeks
- **Data in rest (files on disk)**
  - Identifiable
    - Delete files after 2 years
  - Anonymous
    - Hash with *yearly* secret salt.
      - *So only ‘group by’ within a specific year*

bol.com

What can we analyze?

We can analyze the exact behavior of

- **Named individuals**
  - in the last 24 months.
- **Anonymous individuals**
  - per year for many years.

bol.com

Implementation

Implementation design

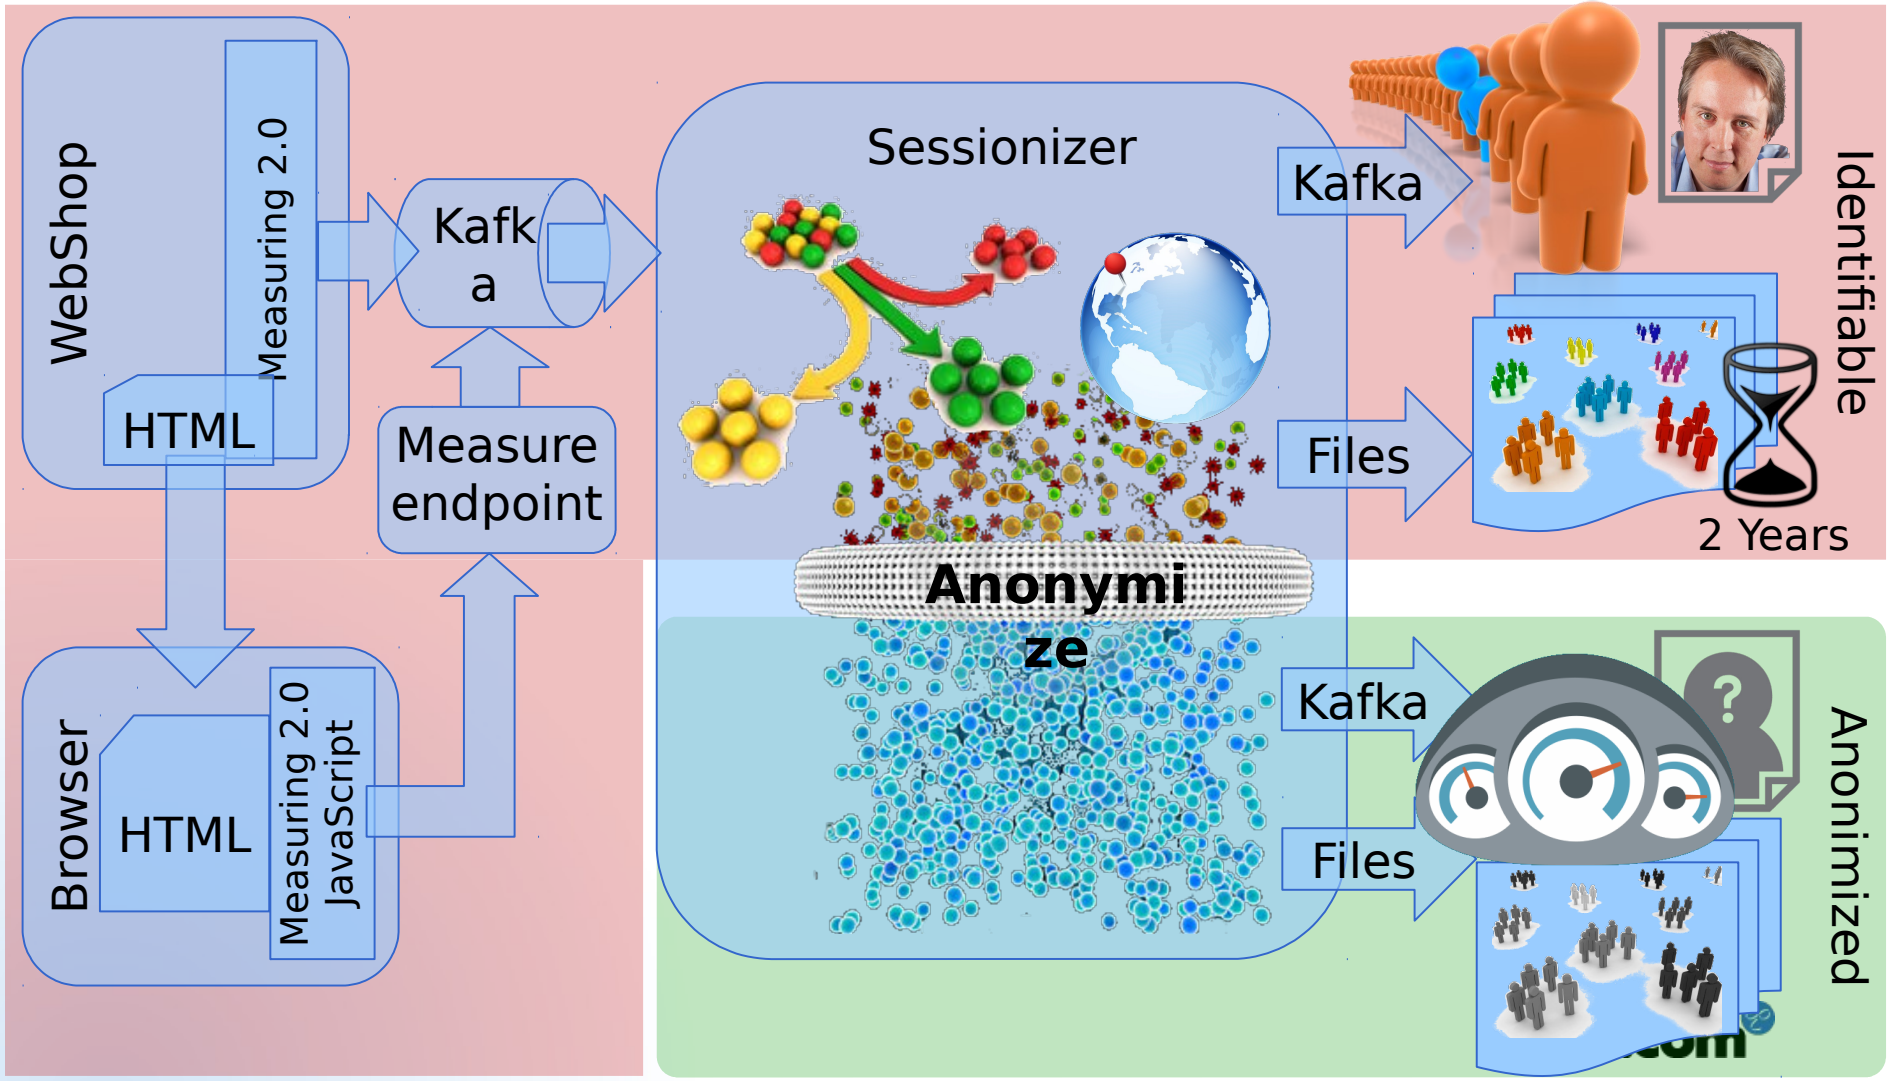

2014: First prototype

Apache Storm

- **Stream processing**
- **Low latency**
- **Many users**

STORM

bol.com

Challenges

```
TopologyBuilder builder = new TopologyBuilder();
builder.setSpout("words", new TestWordSpout(), 10);
builder.setBolt("exclaim1", new ExclamationBolt(), 3)
        .shuffleGrouping("words");
builder.setBolt("exclaim2", new ExclamationBolt(), 2)
        .shuffleGrouping("exclaim1");
```

Make sure these names match!

- **Hard to program**
  - Unfriendly API
- **“At least once”**
- **Need “Exactly once”?**
  - Trident Micro batches
    - No more low latency
- **Storm is stateless**
  - State is an application problem…
  - No recovery of state after failure
- **Run on an existing cluster?**
  - I never got it to run on Yarn (in 2014!).

bol.com

Spark Streaming ?

- **What about Spark Streaming?**
  - Micro batches
    - Too much latency

bol.com

Apache Flink

bol.com

Apache Flink

- **Low latency**
- **Exactly once**
  - With recovery after failure
- **Runs on Yarn**

bol.com

Flink supports our goals

- **Manages state in the framework**
  - and saves it in case of failure
    - CheckPoints & SavePoints
- **Windows**
  - The basic component for ‘visits’
- **“Event time”**
  - A ‘time out’ is based on the events.

bol.com

Challenges

- **Running on Kerberos secured cluster**
  - I fixed that for HBase (FLINK-2977)
  - It dies after 173.5 hours
    - common problem on secured Yarn
      - Delegation Tokens problem (HDFS-9276 ?)
- **My “Windows” are too big to fit in memory.**
  - Only keep the visit ‘state’ in memory
  - Persist the events in HBase
    - TODO: Evaluate RocksDB
- **Exactly once**
  - Not on Kafka output!

bol.com

Assign VisitId?

- **Window**
  - Only releases the records after “PURGE”
    - Realtime stream?
- **Custom operator**
  - Must handle everything manually

bol.com

Assign VisitId Operator

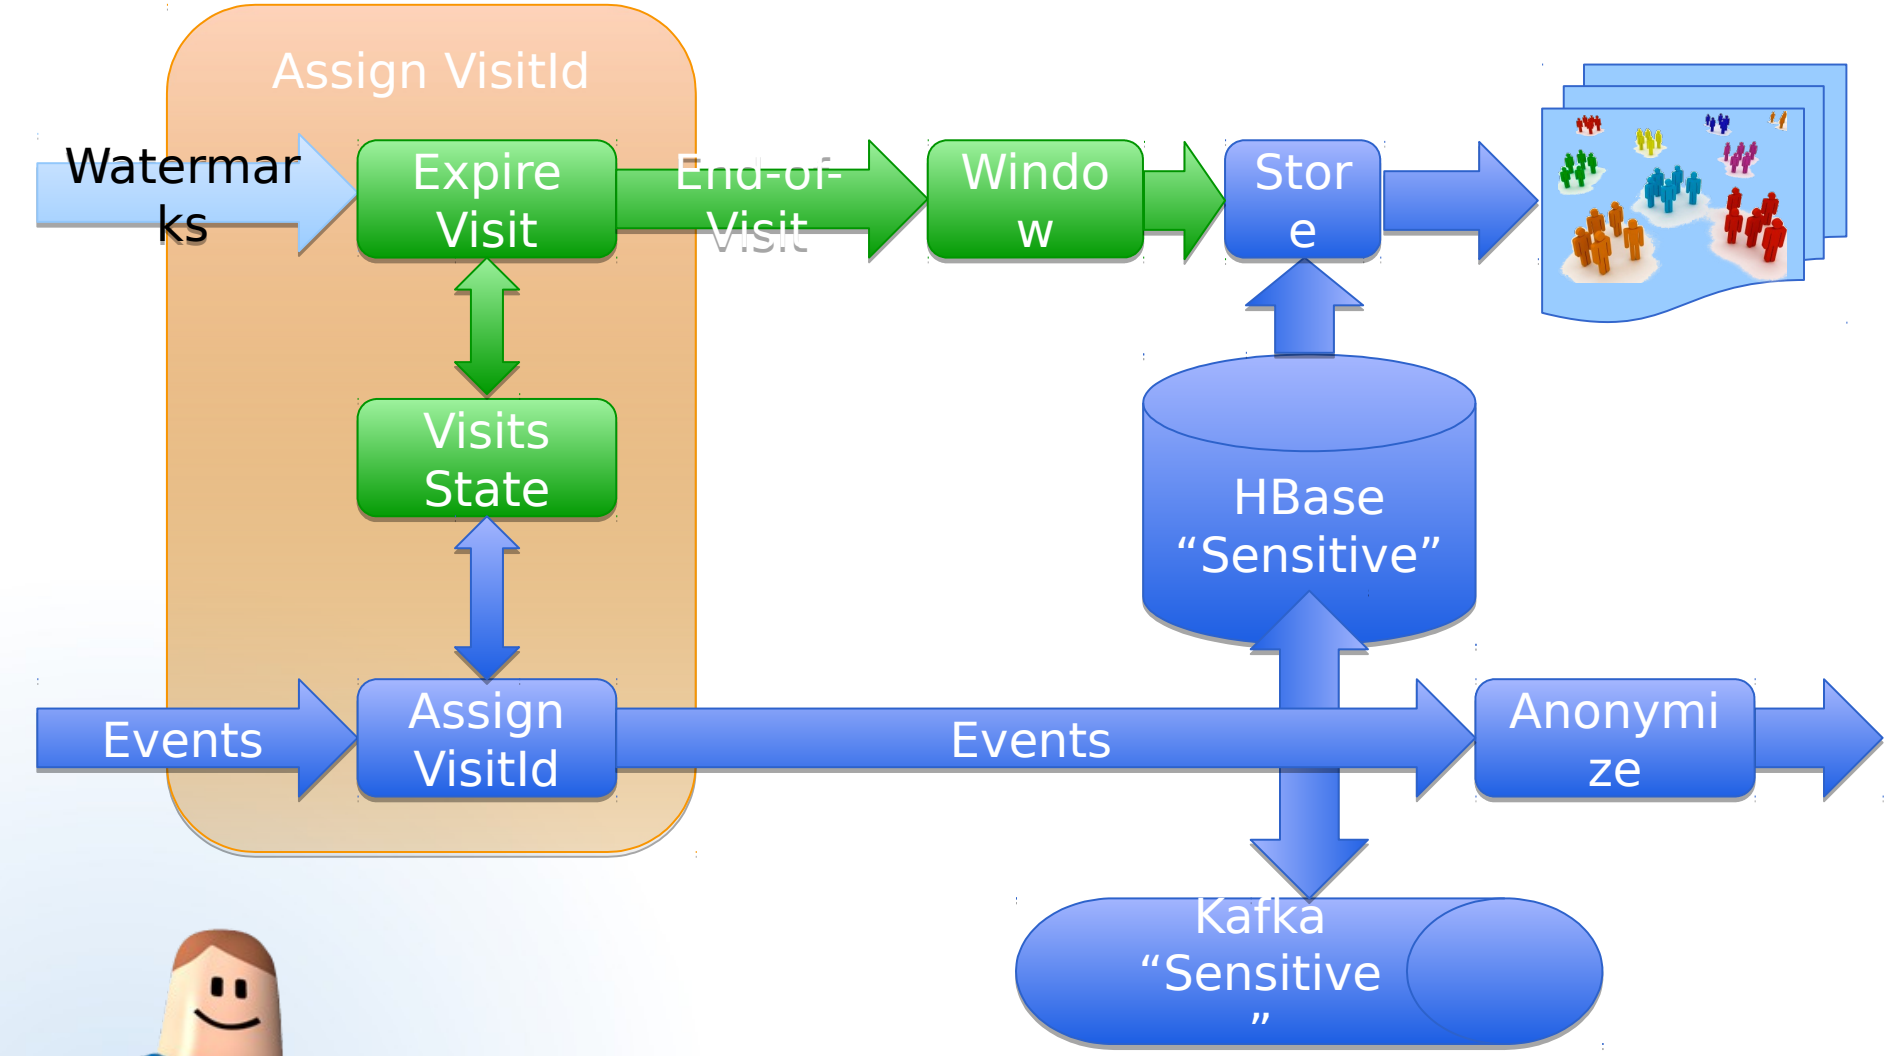

bol.com

Implementation

- **Caveats**
  - An operator has only 1 type of output
  - Watermarks are essential for Windows

Watermarks trigger window evaluation.

Final plumbing

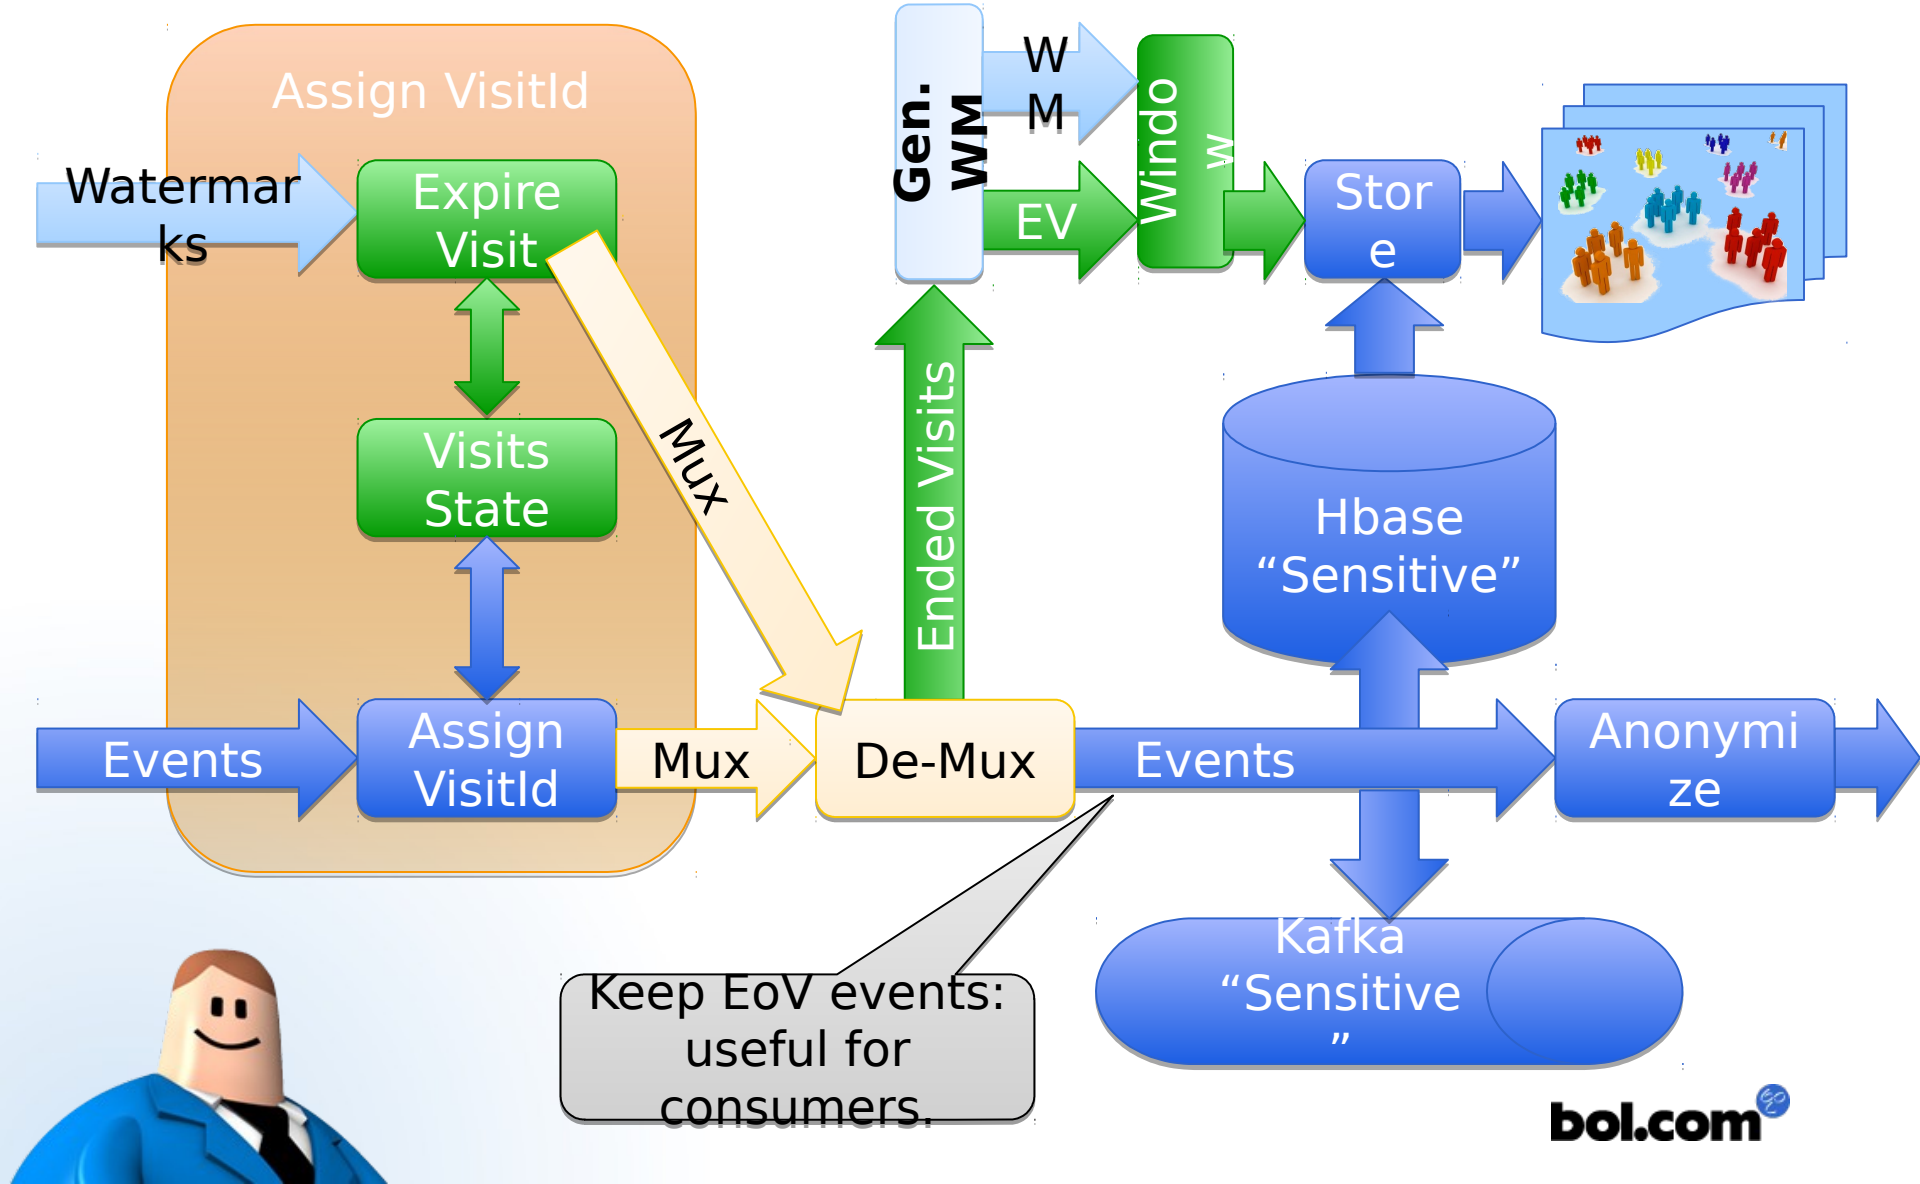

Debugging watermarks

- **Watermarks are**
  - The core of building Windows
  - Tricky to get right
    - Input data is ‘messy’
    - You can get a second ‘same window’
- **So**
  - Need debugging for the combination of watermarks & events.

bol.com

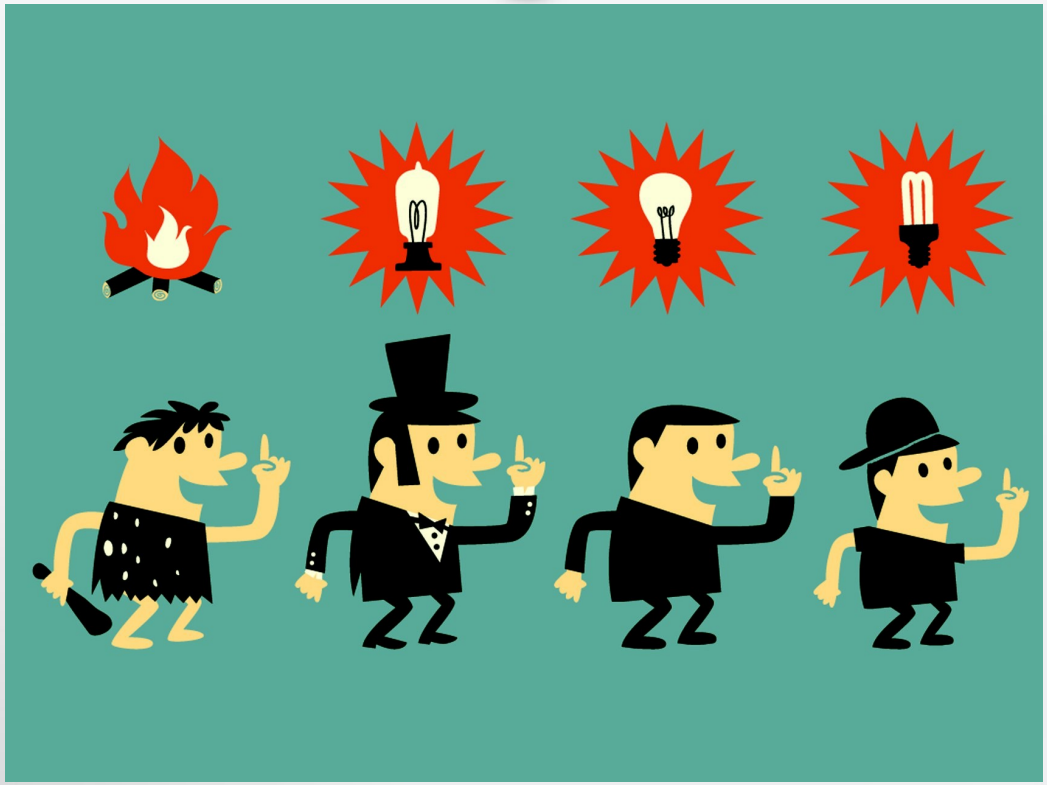

Evolution in streaming data

Because requirements change over time.

bol.com

Persisting data

- **Define Schema in: Apache AVRO**
  - Nice schema language (IDL)
  - Generates Java classes
  - Supports schema evolution

- **Persist in files: Apache Parquet**
  - Compresses well
  - Great on read use cases
  - Write straight from Avro classes.
    - See o.a.parquet.avro.AvroParquetWriter

bol.com

Persisting data

- **Stream: Apache Kafka**
  - High throughput low latency event pipe
  - Persists records for a few weeks
  - Requires the record to be a byte[]

- **Question:**
  - How to serialize record into byte[]?

bol.com

Persisting data

- **Apache AVRO**
  - Standard byte[] for the record.
  - Needs the original Schema on read !!
- **What if the Schema changes ???**
  - Kafka keeps the events for weeks
- **AVRO-1704 (Work in progress)**
  - Serialize record to a “Message” byte[]
    - schema fingerprint (hash)
    - pluggable schema ‘database’

bol.com

This is no hype

Thursday March 27th 2014

Thuiswinkel awards

Bol.com was elected as the 'Best webshop of the Netherlands'

During the last year bol.com made great steps forwards **by successfully applying customer profiling and BigData.**

bol.com

Join us

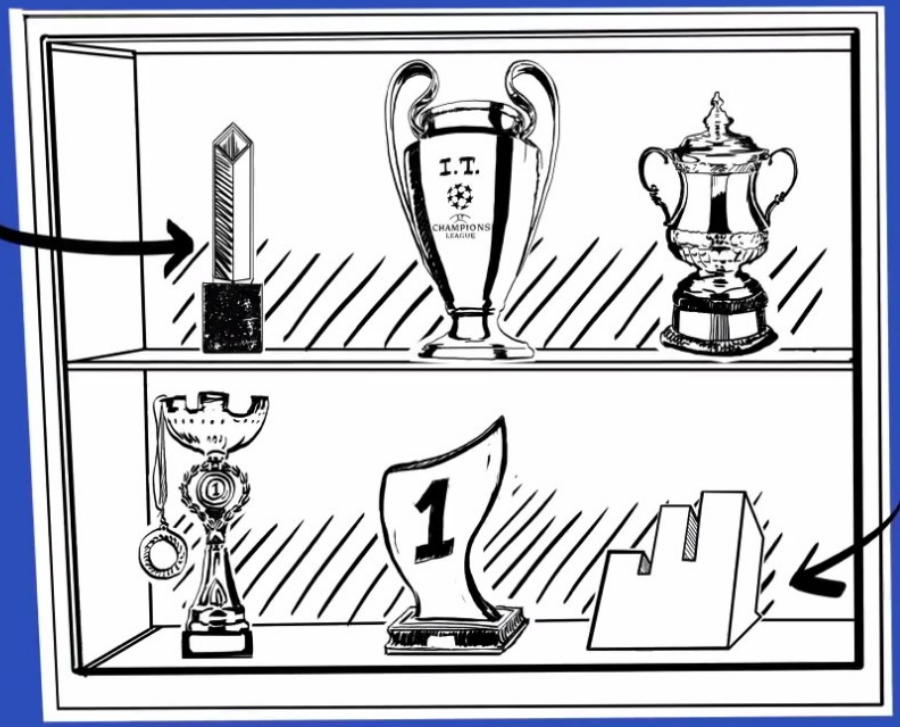

Thuiswinkel awards

effie awards

UITDAGEND
IT-TRAININGSPROGRAMMA

ON THE JOB COACHING

PERSONAL TRAININGS PROGRAMMA

40 SCRUMTEAMS

www.bol.com/banen


Q&A
You have
Questions
We have
Answers