

Data warehousing on Hadoop

Marek Grzenkowicz

Roche Polska



Agenda

Introduction

Case study: StraDa project

- Source data
- Data model
- Data flow and processing
- Reporting

Lessons learnt

Ideas for the future

Q&A


Some context



Roche

Roche
Pharmaceuticals
(Pharma)

Roche Diagnostics
(Dia)



Objectives of the project StraDa

A large, dark blue arrow pointing from left to right, containing text and a list of items.

Measure the
performance of the labs

- Workload
- Turnaround time (TAT)

A large, light blue arrow pointing from right to left, containing text and a list of items.

Discover and
understand the reasons

- Hardware configuration
- Tasks and work organization
- Other, unknown factors

There are **7 types** of the Workload KPI and **19** TATs.

Source data – files

22- PARC DES POTERIES	33 - BARR	010210124	LEN	10010210124	?	04	0	?	0	?	LiaisCapi	16	?	?	
1900	2015-02-10 10:43:54	M	ANA LIAISON XL1	31	LIASON XL	LXL	3					2015-02-11 22:04:37	2015-02-10	ANALYSEO	
3170008496	OUI	10245													
22- PARC DES POTERIES	33 - BARR	010210124	LEN	10010210124	?	00	0	?	0	?	Serum 03	?	?		
1900	2015-02-10 10:43:54	M	ANA MODULAR EE	91	MODULAR	MODULAR 3						2015-02-11 22:04:37	2015-02-10	ANALYSEO	31700084
96	OUI	10245													
22- PARC DES POTERIES	33 - BARR	010210124	LEN	10010210124	?	04	0	?	0	?	LiaisCapi	16	?	?	
2003	2015-02-10 13:40:54	M	ANA LIAISON XL1	31	LIASON XL	LXL	3					2015-02-11 22:04:37	2015-02-10	ANALYSEO	
3170008496	OUI	10245													
22- PARC DES POTERIES	33 - BARR	010210124	LEN	10010210124	?	04	0	?	0	?	LiaisCapi	16	?	?	
2004	2015-02-10 14:22:13	M	ANA LIAISON XL1	31	LIASON XL	LXL	3					2015-02-11 22:04:37	2015-02-10	ANALYSEO	
3170008496	OUI	10245													
22- PARC DES POTERIES	33 - BARR	010210124	LEN	10010210124	?	04	0	?	0	?	LiaisCapi	16	?	?	
2004	2015-02-10 14:42:10	T130	ANA LIAISON XL1	31	LIASON XL	LXL	3					2015-02-11 22:04:37	2015-02-10	ANALYSEO	
3170008496	OUI	10245													
22- PARC DES POTERIES	33 - BARR	010210124	LEN	10010210124	?	04	0	?	0	?	LiaisCapi	16	?	?	
2004	2015-02-10 14:42:12	T130	ANA LIAISON XL1	31	LIASON XL	LXL	3					2015-02-11 22:04:37	2015-02-10	ANALYSEO	
3170008496	OUI	10245													
22- PARC DES POTERIES	33 - BARR	010210124	LEN	10010210124	?	00	0	?	0	?	Serum 03	?	?		
2004	2015-02-10 14:42:13	T130	ANA MODULAR EE	91	MODULAR	MODULAR 3						2015-02-11 22:04:37	2015-02-10	ANALYSEO	31700084
96	OUI	10245													
22- PARC DES POTERIES	33 - BARR	010210124	LEN	10010210124	?	04	0	?	0	?	LiaisCapi	16	?	?	
2007	2015-02-10 14:42:14	T130	ANA LIAISON XL1	31	LIASON XL	LXL	3					2015-02-11 22:04:37	2015-02-10	ANALYSEO	
3170008496	OUI	10245													
22- PARC DES POTERIES	33 - BARR	010210124	LEN	10010210124	?	00	0	?	0	?	Serum 03	?	?		
2007	2015-02-10 14:42:14	T130	ANA MODULAR EE	91	MODULAR	MODULAR 3						2015-02-11 22:04:37	2015-02-10	ANALYSEO	31700084
96	OUI	10245													

Format: TSV files

Size: 1-100 MB, usually ~25 MB

Header: 1 line with column names + 1 line with file metadata

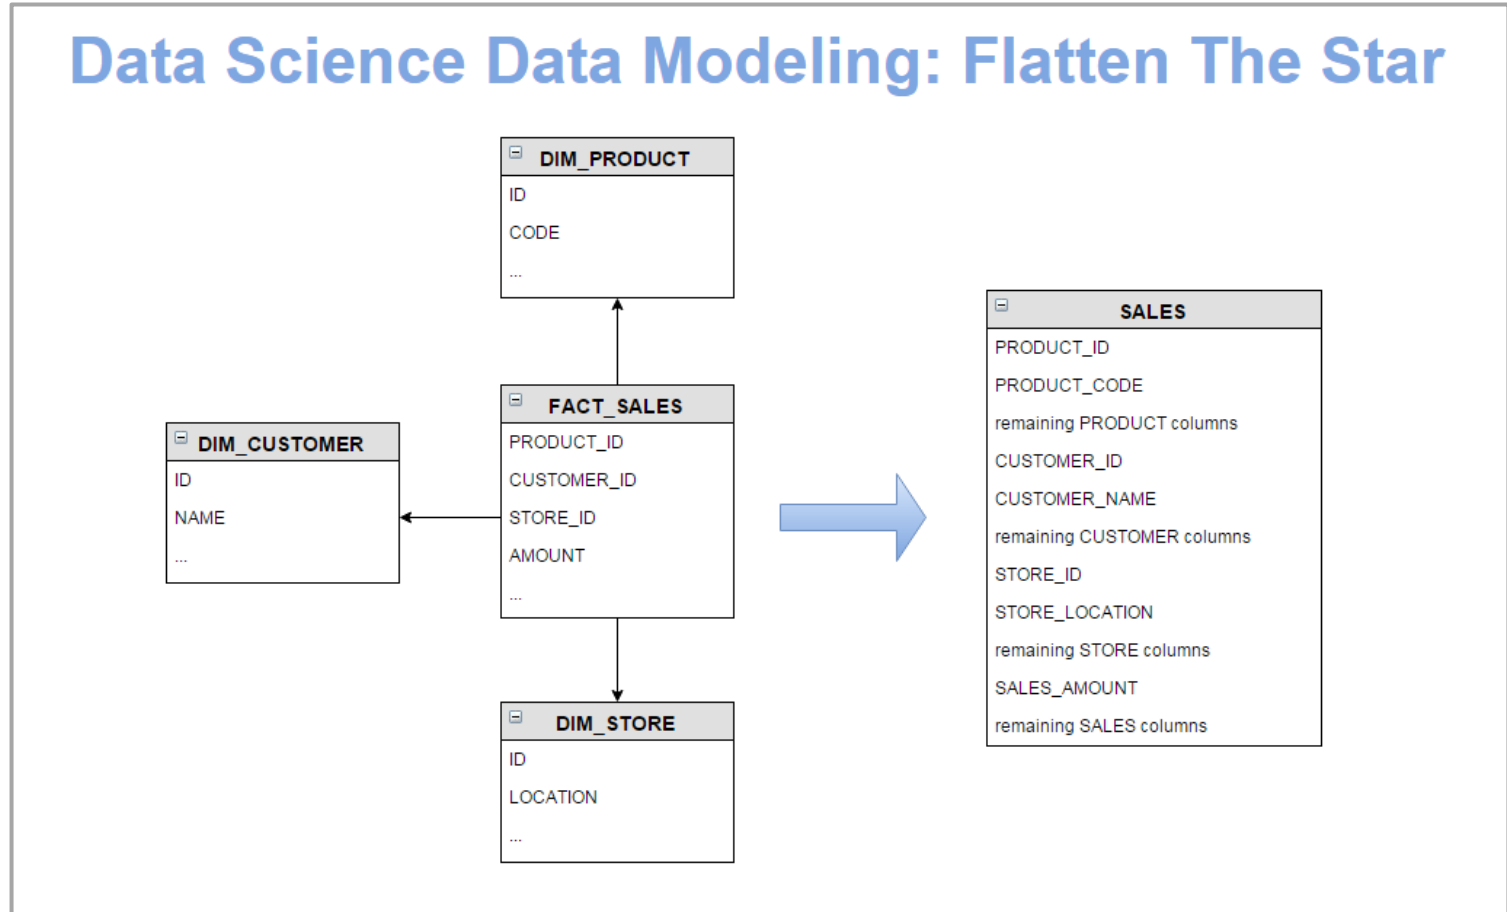
Content: events generated by the instruments and related IT systems

Source data – events

Description	Code
Sample collected	2000 (first)
Order registered	1100
Test ordered	2900 (first)
Sample sorted manually	2002
Sample assigned to a transport box	2028 (first)
Sample sent to the lab	2026
Sample retrieved from a transport box	2029
Sample arrived in the lab	2027 (first)
Test request send to analytical instrument	3013 (first)
Test result produced	3003
Last result produced	3003 (last)
Result manually validated	3006
Order complete	2012 (last)

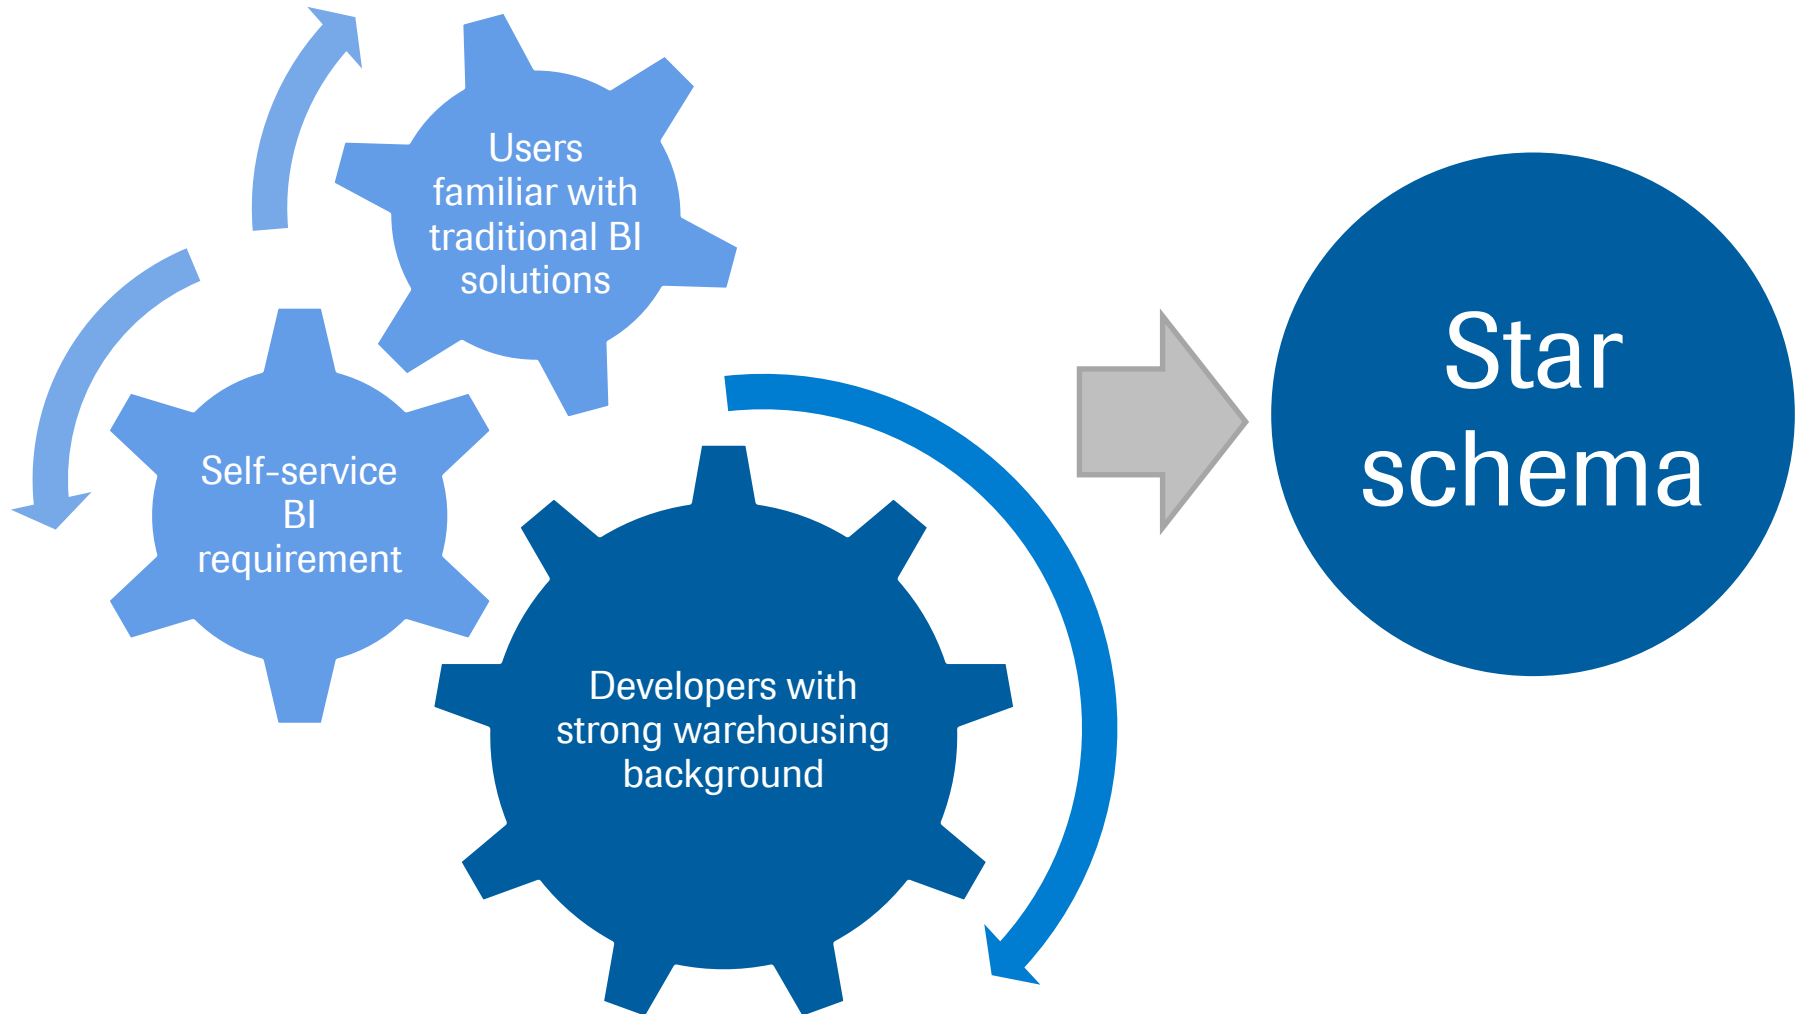


Dimensional modeling and Big Data modeling

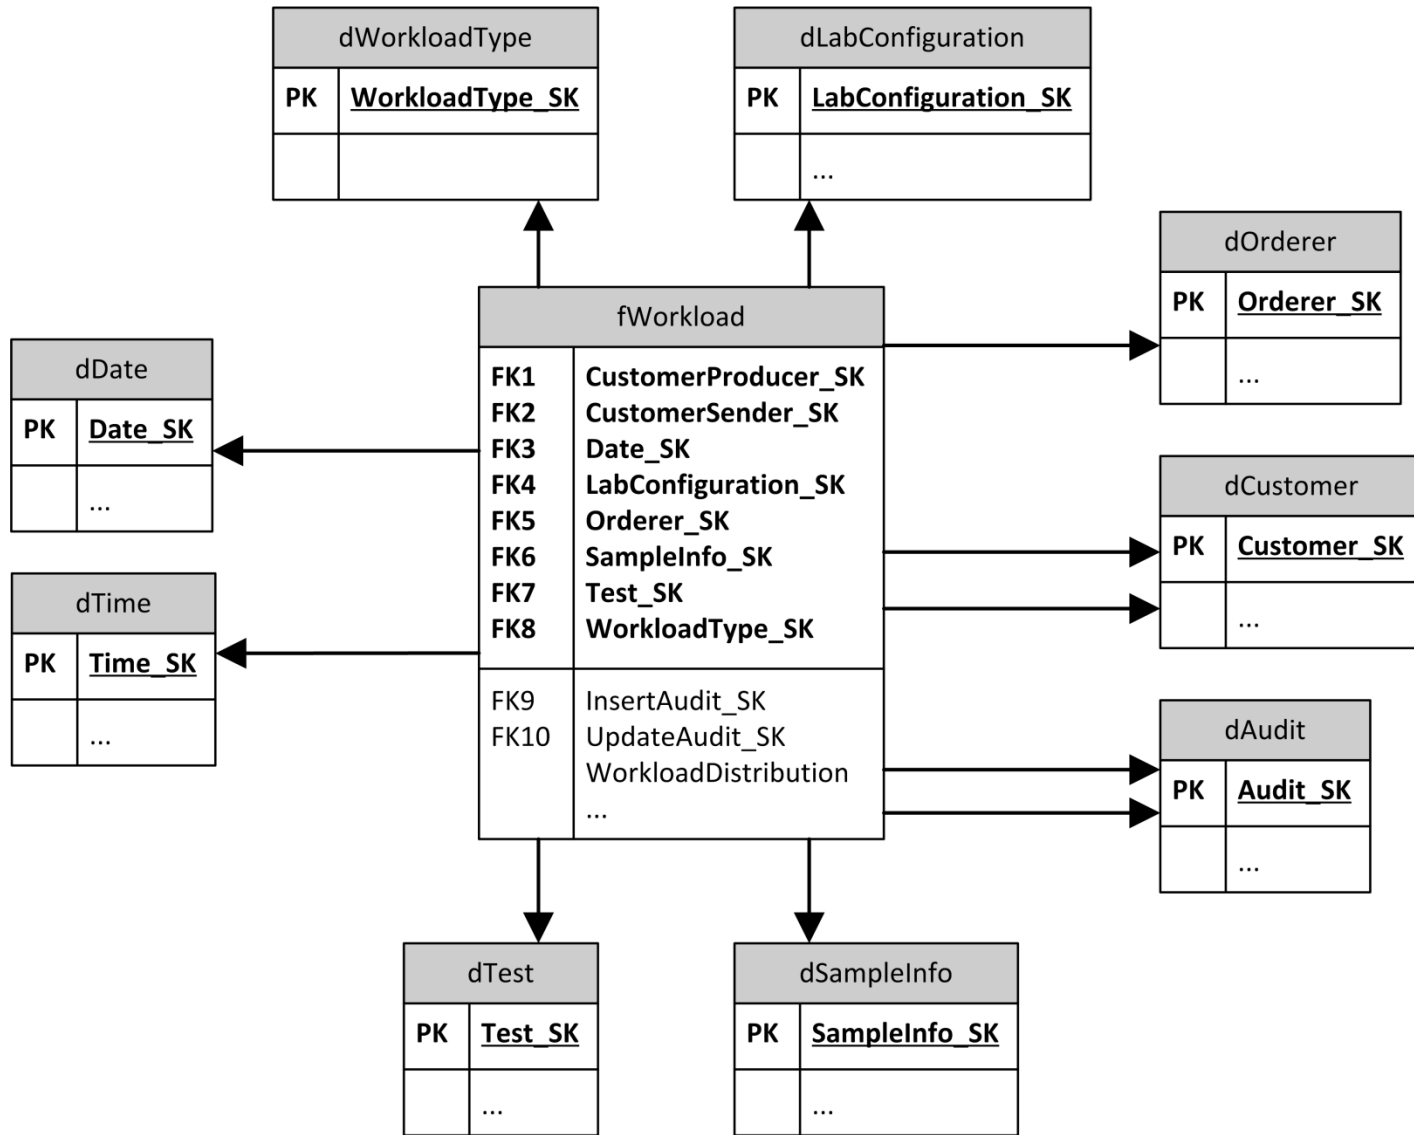


Bill Schmarzo, ECM – [Hadoop Data Modeling Lessons - by Vin Diesel](#)

StraDa data model



Data model



Initial vision

Data model

- **Star schema**
 - Alternative: a flattened table for each business process

Master data

- **SQL Server**
 - Alternative: Hadoop as a single storage system

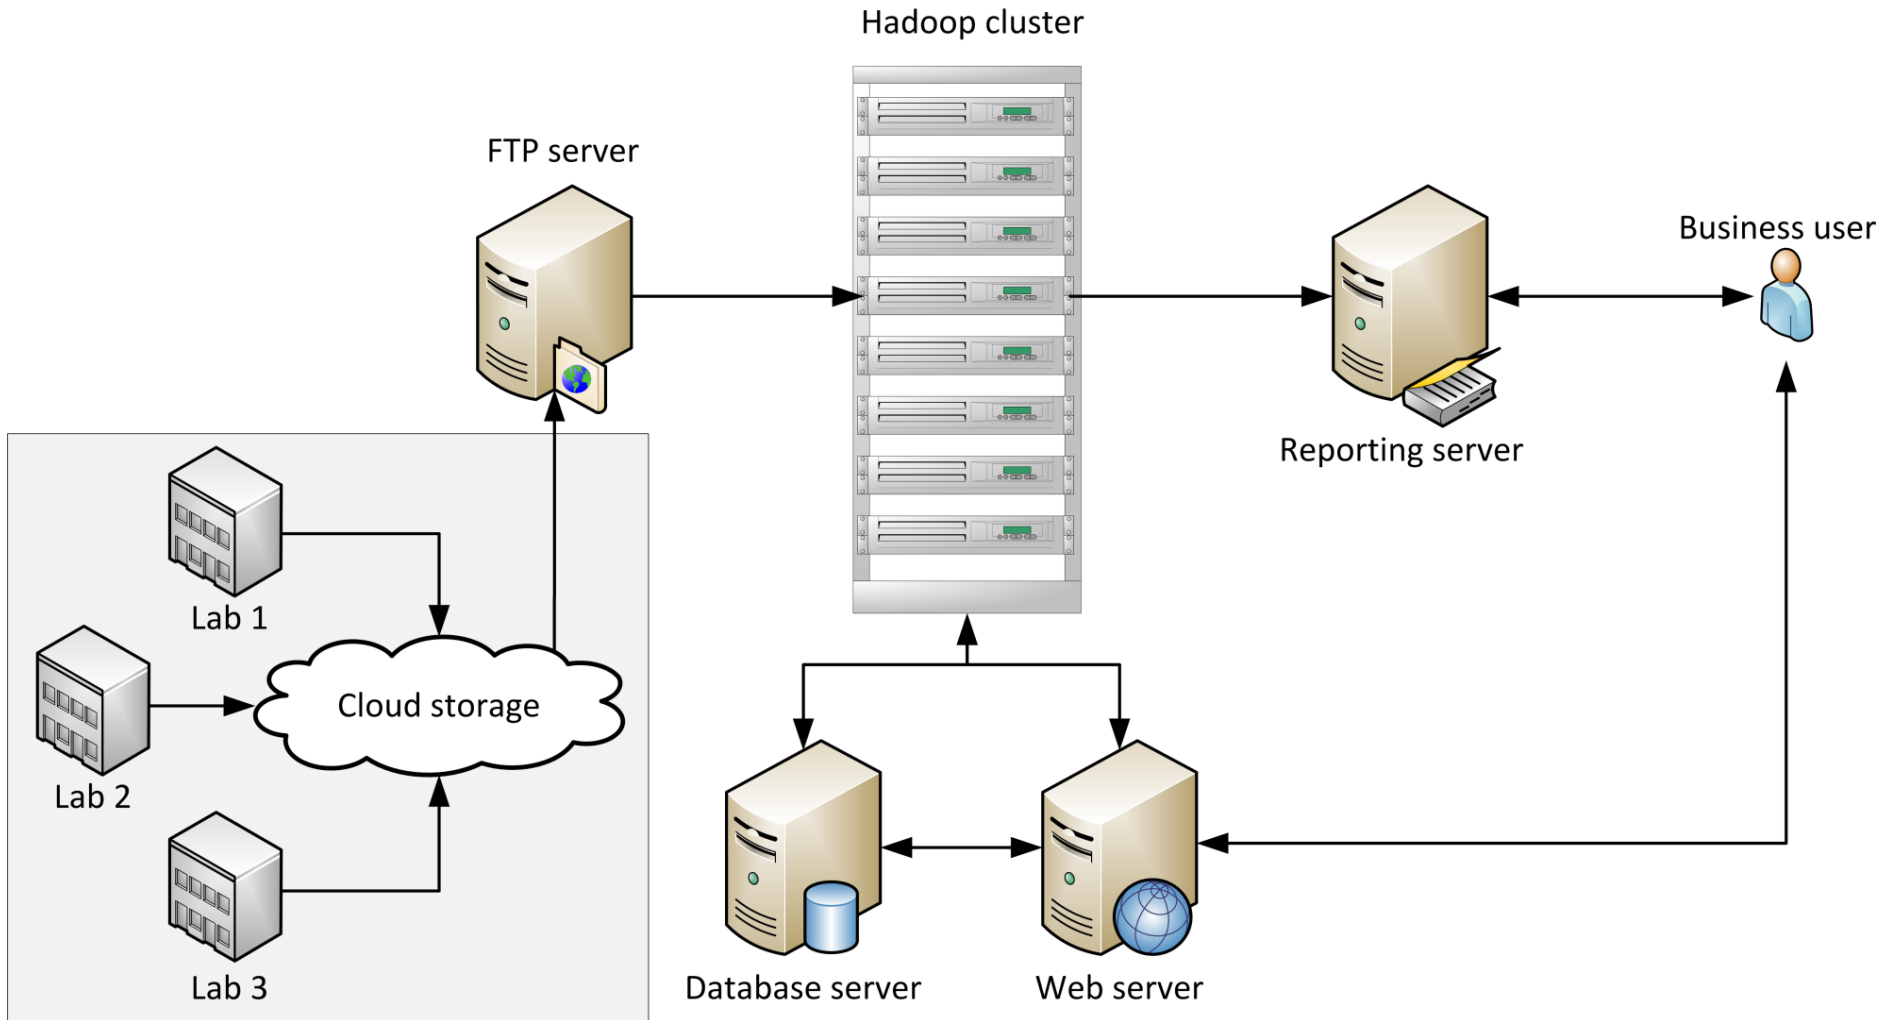
ETL

- **Hive, Pig**
 - Alternative: M/R jobs developed in Java

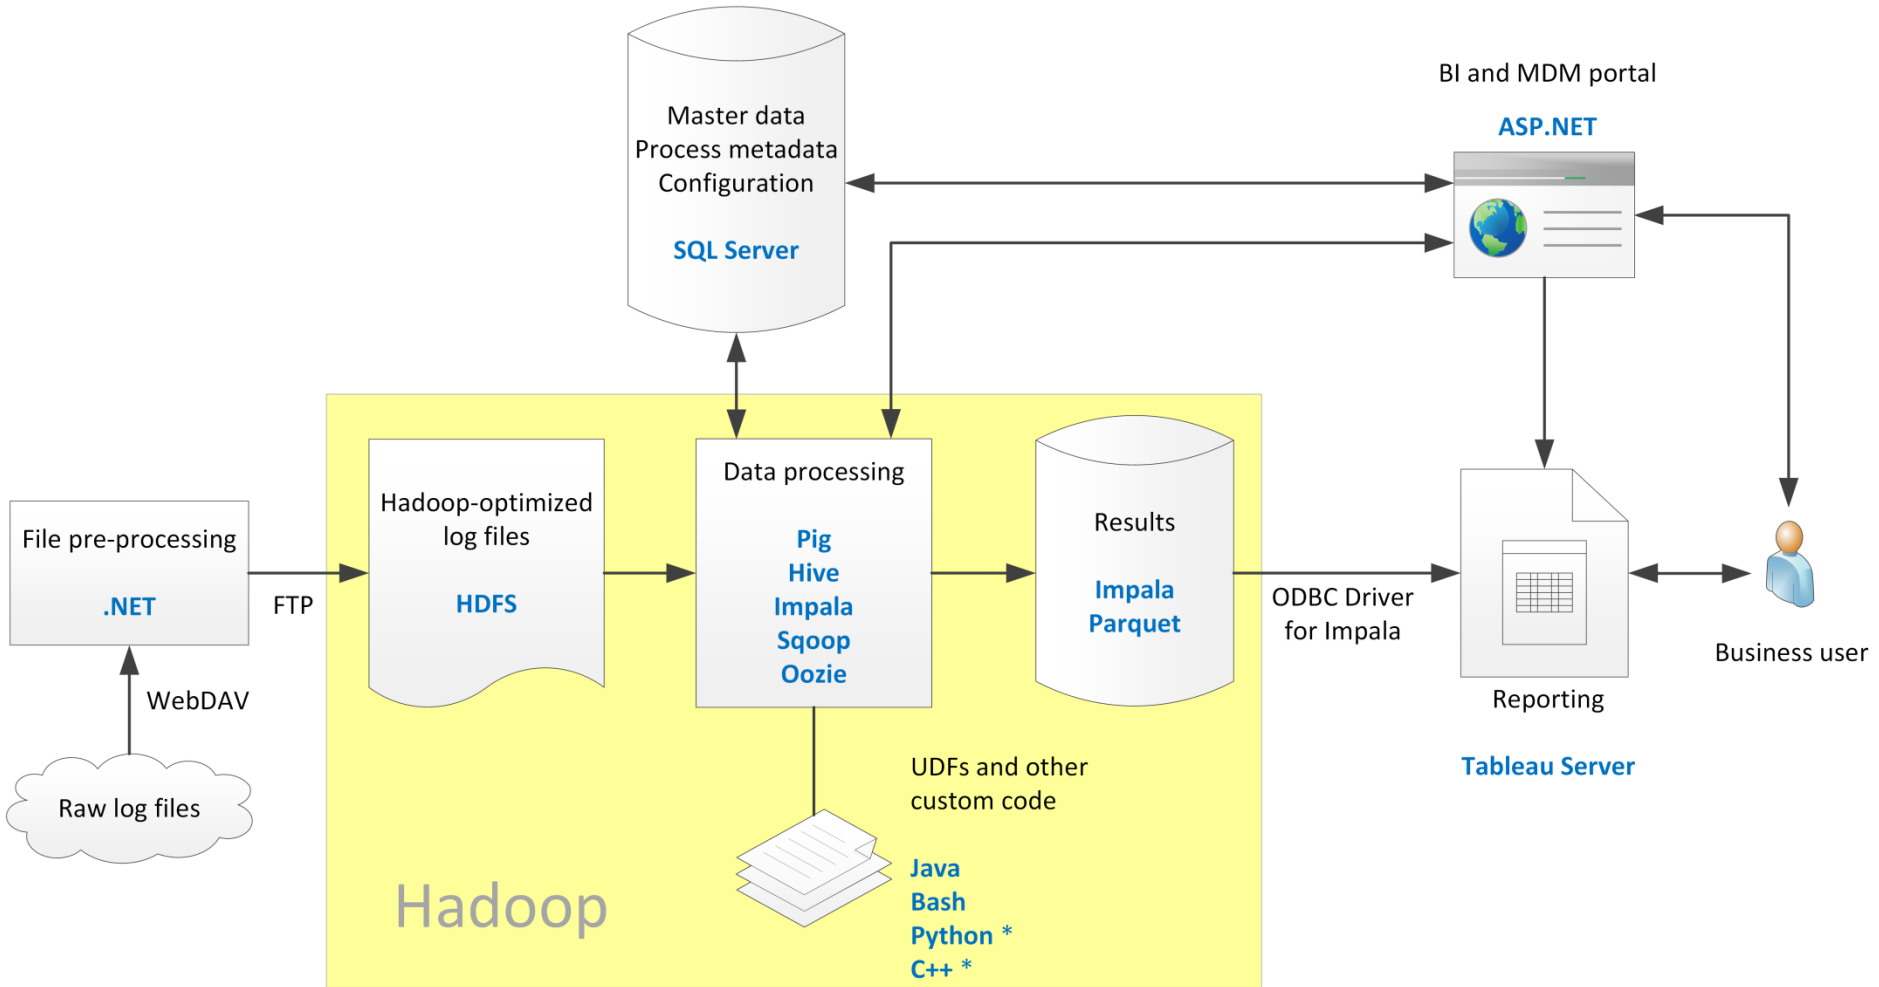
Reporting

- **Tableau**
 - Alternative: a dedicated Big Data reporting tool; R and Shiny

Overview of the data flow

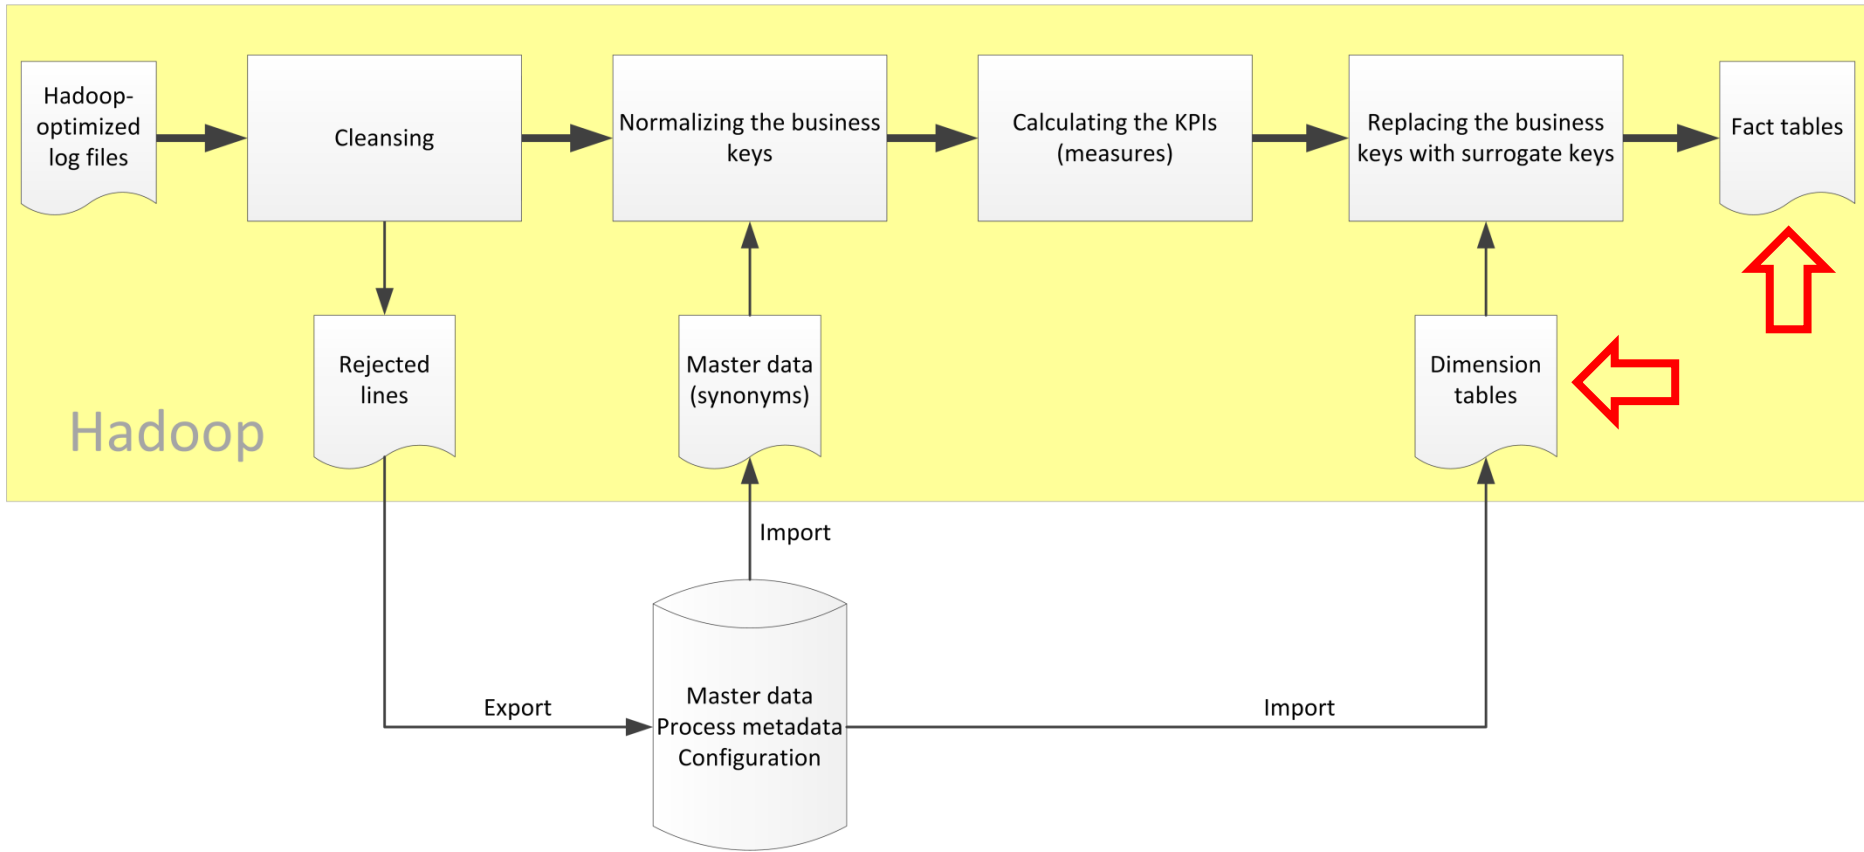


Tools and technologies



* Not included in the production release

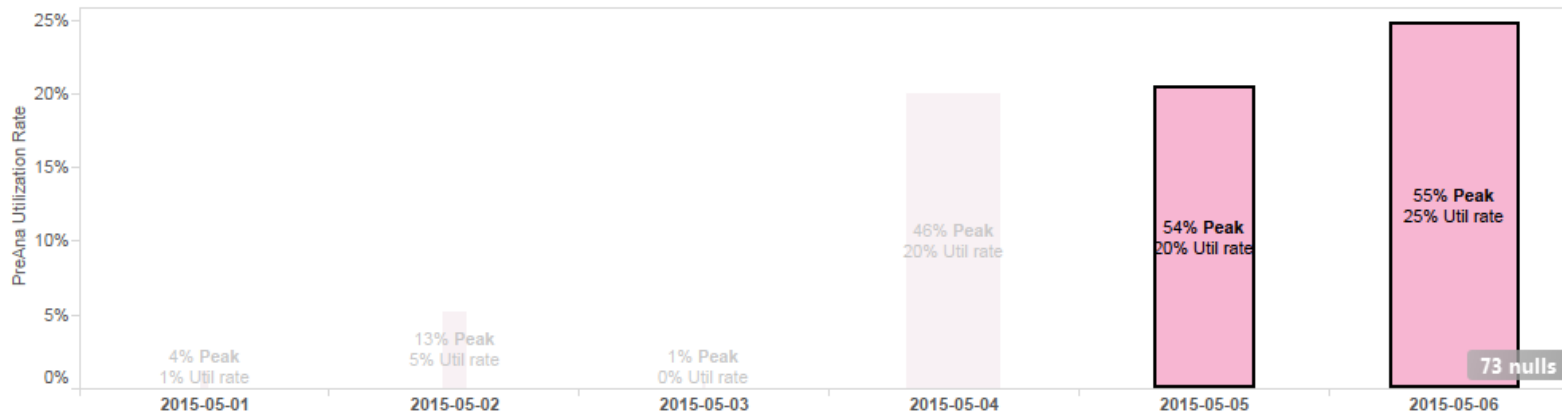
ETL



Reporting



Daily PreAna util rate



Lab Configuration

3170004327 - LABO D... ▾

MinDate

1-5-2015

MaxDate

7-5-2015

Capacity Calc

Automatic ▾

Immuno Capacity per hour


360

System

(All) ▾

System

■ C8000_Cce cobas 8000..




Utilization rate

[actual workload : capacity ratio]

Reporting

Customer Name: (All) | MinDate: 1-6-2015 | MaxDate: 8-6-2015 | Hour: (All) | Workflow Span: Total TAT Collection | Module Manufacturer: (All) | Parameter: (All)

Sample is urgent: (All) | Sample Type: (All) | Orderer Code: (All) | TAT range bin (minutes): 0 | Percentile X: 50% | Percentile Y: 31%



TAT distribution (continuous)

[vertical lines – different percentiles]

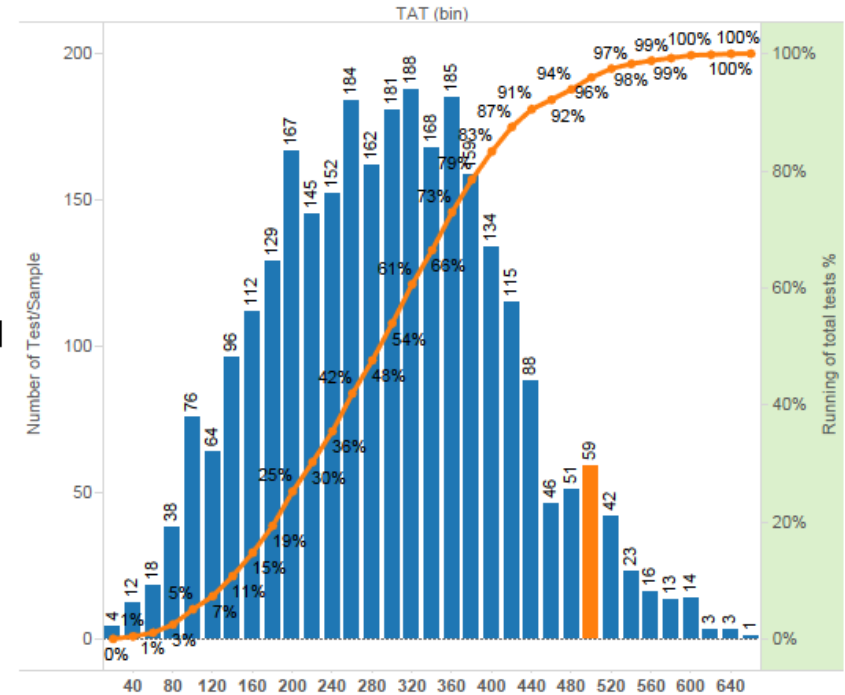
Reporting

Customer Name: (All) | MinDate: 1-6-2015 | MaxDate: 8-6-2015 | Hour: (All) | Workflow Span: Total TAT Collection | Module Manufacturer: (All) | Parameter: (All)

Sample is urgent: (All) | Sample Type: (All) | Order Code: (All) | TAT bin range (minutes): 0 | TAT bin size (minutes): 20

Lab Configuration	Site Configuration	System	Module	Submodule
3170007035 - * CTRE BIOLOGIQUE CHEMIN VERT	* CTRE BIOLOGIQUE CHEMIN VERT	C6000-2	C6000-2 ..	Unknown
		cobas 6000	C6000-2 E1	C6000-2 E11 cobas 6000 e601 Cellule 1 1
		cee 3	e601 2	C6000-2 E12 cobas 6000 e601 Cellule 2 2
			e601 3	C6000-2 E21 cobas 6000 e601 Cellule 1 1
			e601 3	C6000-2 E22 cobas 6000 e601 Cellule 2 2
			Unknown	Unknown
		CAPILLA..	Unknown	Unknown
		Cobas 8000 Ccee	Cobas 80..	Unknown
		cobas 8000	Cobas 8000 C-1	cobas 8000 c701 rotor A 1
		iCcee 2	Cobas 8000 C-2	cobas 8000 c701 rotor B 2
			Cobas 80..	Cobas 8000 ISE1 cobas 8000 ISE1 1
			Cobas 8000 E11	cobas 8000 e602 cellule 1 1
			Cobas 8000 E12	cobas 8000 e602 Cellule 2 2
			Cobas 8000 E21	cobas 8000 e602 cellule 1 1
			Cobas 8000 E22	cobas 8000 e602 Cellule 2 2
	Unknown	Unknown		
	Unknown	Unknown		
Liaison XL..	Unknown	Unknown		
Liaison XL..	Unknown	Unknown		
Remisol N..	Unknown	Unknown		
Vidas VID..	Unknown	Unknown		
3170007094 - LAM	LAM TRENZS	AIS ND_in..	Unknown	Unknown
TRENZS BP40068	BP40068	BR E411 c..	Unknown	Unknown
		BR Int400..	Unknown	Unknown

TAT Distribution



TAT distribution (histogram)

[orange line – cumulative distribution; orange bar – 95th percentile]

Some geeky numbers

Input files	Fact tables	
300 GB	350 million (10^6) rows	2.5 GB
	1.5 billion (10^9) rows	16 GB

Input files: uncompressed text files; 1 country, 16 months

Fact tables: compressed Parquet files

Lessons learnt

Hadoop is not fully mature yet

You need a Hadoop administrator in the team

Broad skillset is necessary

Lack of proven best practices and literature

GitHub is essential for a Hadoop developer



site:github.com parquet convertINT96



Bugs / Unwanted features / Surprises



- Hive supports only equality comparisons in the JOIN predicate [[by design](#)]
- **The beeline client may fail when executing a HiveQL script that contains comments** [[HIVE-8396](#)]



- Impala does not fully support non-ASCII characters – they can be stored and retrieved but not manipulated [[by design](#), pending future release]

```
UPPER('Viscérale') -> 'VISCÉRALE'
```
- **Non-standard encoding of date and time values that is incompatible with Parquet** [[IMPALA-2111](#)]

Bugs / Unwanted features / Surprises



- **Lack of collaboration and productivity features (IDE) necessary for teams bigger than 2-3 developers**
- No way to migrate solutions between different environments



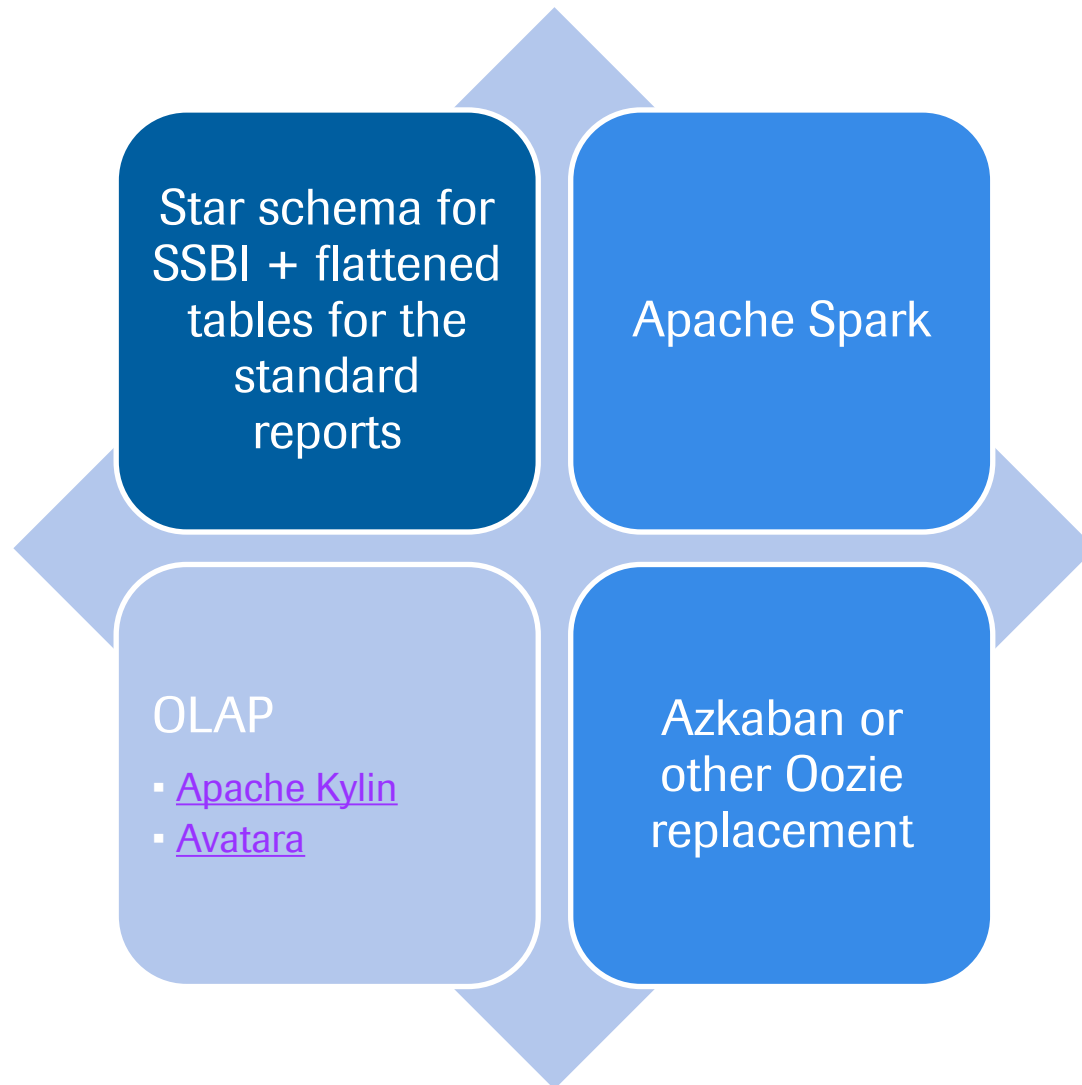
- **There is no Impala action** [[OOZIE-1591](#)]
- Workflows sometimes get corrupted and stop loading in the editor and there is no easy way to fix them
- Some Oozie features are not supported by the Hue editor
- When a workflow is shared and then edited or run, the non-owner can no longer access its deployment folder [[HUE-2376](#)]

Bugs / Unwanted features / Surprises



- The hive client cannot be used in a kerberized cluster, because it was not designed to follow the Sentry security rules [by design]
- **Workflows that contain credentials cannot be exported** [[HUE-1900](#)]
- Additional, rather complex configuration is needed to make all the log (workflows, M/R jobs, etc.) available for all the team members [by design]

Ideas for the future



Conclusion

Can I build a data warehouse on Hadoop?

- Yes.

Can the star schema be used?

- Yes, but it is a usability/performance tradeoff.
- YMMV, so test it carefully.

Can I just put Hadoop in place of my RDBMS?

- No, it is way more complex than that.

Is it worth it?

- It depends, so don't follow the Big Data hype blindly.

Questions and (hopefully) answers



marek.grzenkowicz@roche.com

<http://it.roche.pl/>

Recommended materials

1. Ralph Kimball and Matt Brandwein – [Hadoop 101 for EDW Professionals](#)
 - [Hadoop 101 for EDW Professionals – Dr. Ralph Kimball Answers Your Questions](#)
2. Ralph Kimball and Eli Collins – [EDW 101 for Hadoop Professionals](#)
3. Ralph Kimball – [Newly Emerging Best Practices for Big Data](#)
4. Ralph Kimball – [The Evolving Role of the Enterprise Data Warehouse in the Era of Big Data Analytics](#)
5. Josh Wills – [What Comes After The Star Schema? \(slide deck\)](#)

Doing now what patients need next