

The Rise of Multi-model Databases

Vitaliy Rudnytskiy & Abdel Dadouche





**Click 'Rate Session'
to rate session
and ask questions.**

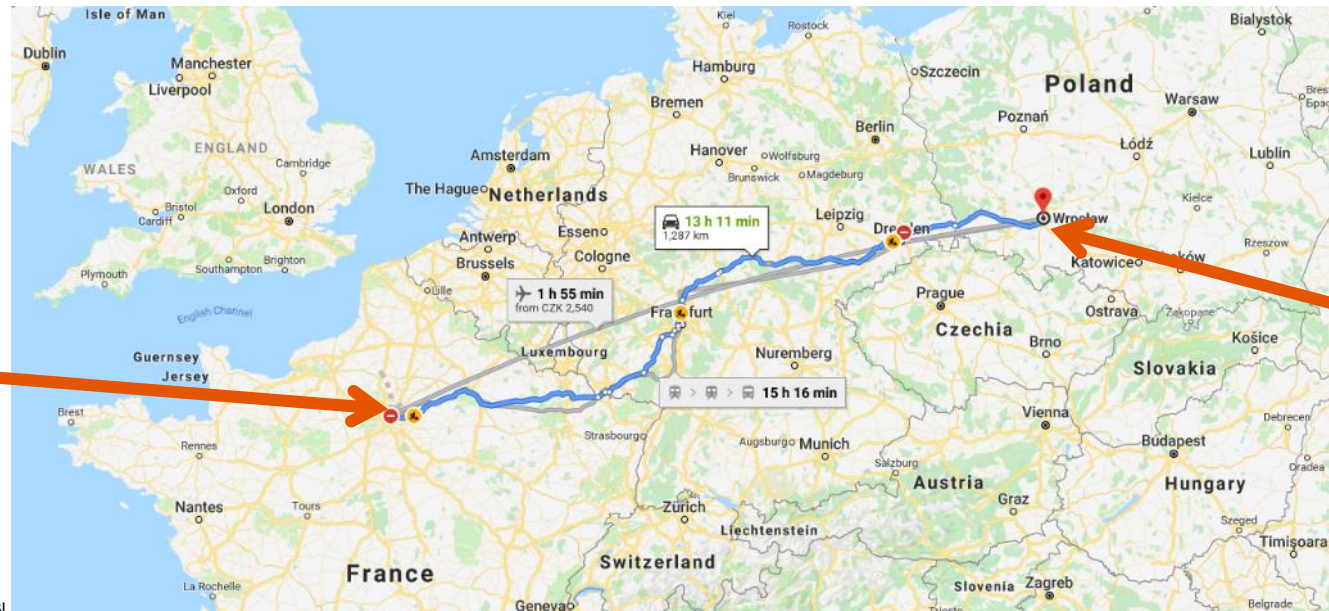
Let's start with your presenters today

Abdel(halim) Dadouche

- From: Paris, France
- Topics: Machine Learning and Data Engineering
- Hobbies: Hackathons and Home Improvement

Vitaliy Rudnytskiy (a.k.a. Witalij Rudnicki)

- From: Wrocław, Poland
- Topics: Analytics, Big Data & IoT
- Hobbies: Local organizer of SAP Community meetups & Beer tasting!



Ever heard about **SAP**?

Enterprise software
company founded in
1972 in **Germany**

96k+ employees across
140+ countries

400k customers WW

77% of the world's
transaction revenue

25 industries

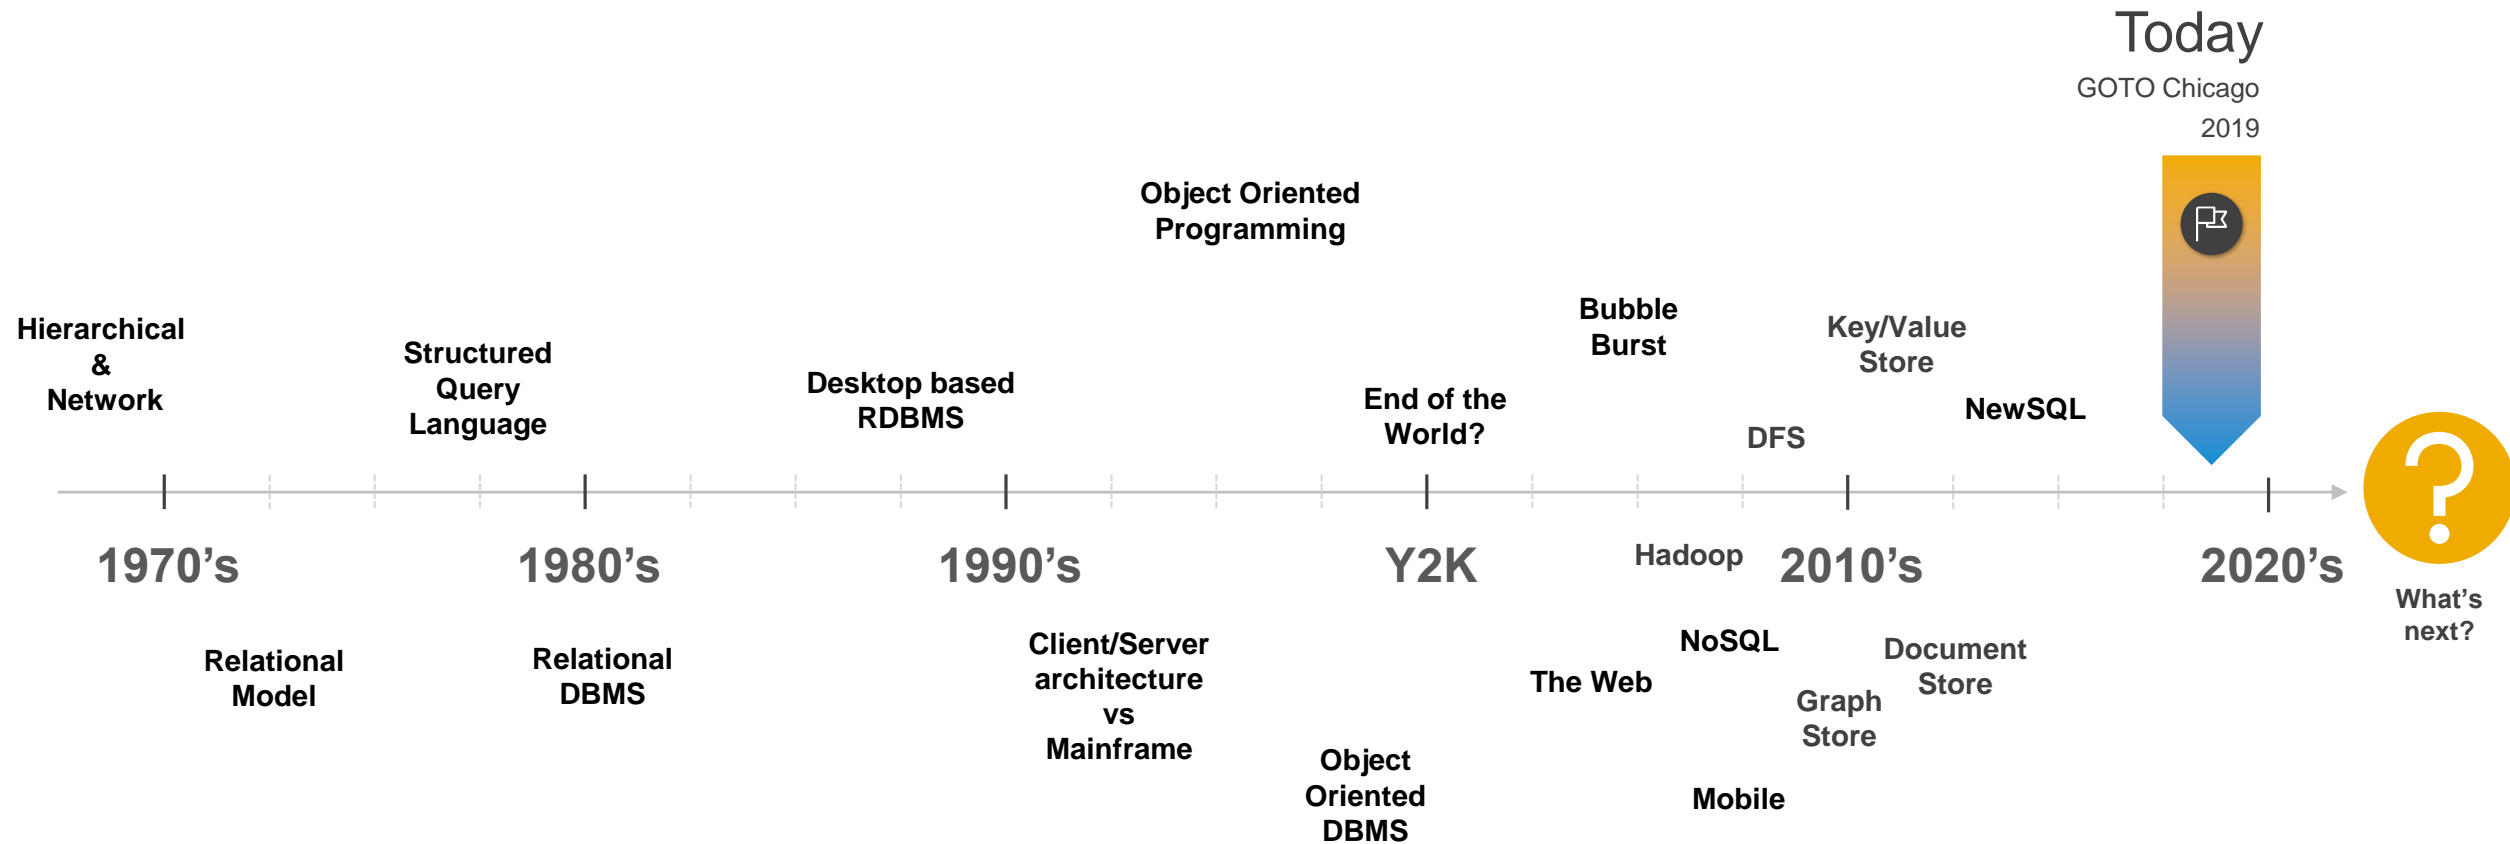
12 lines of business

The world's largest
business network

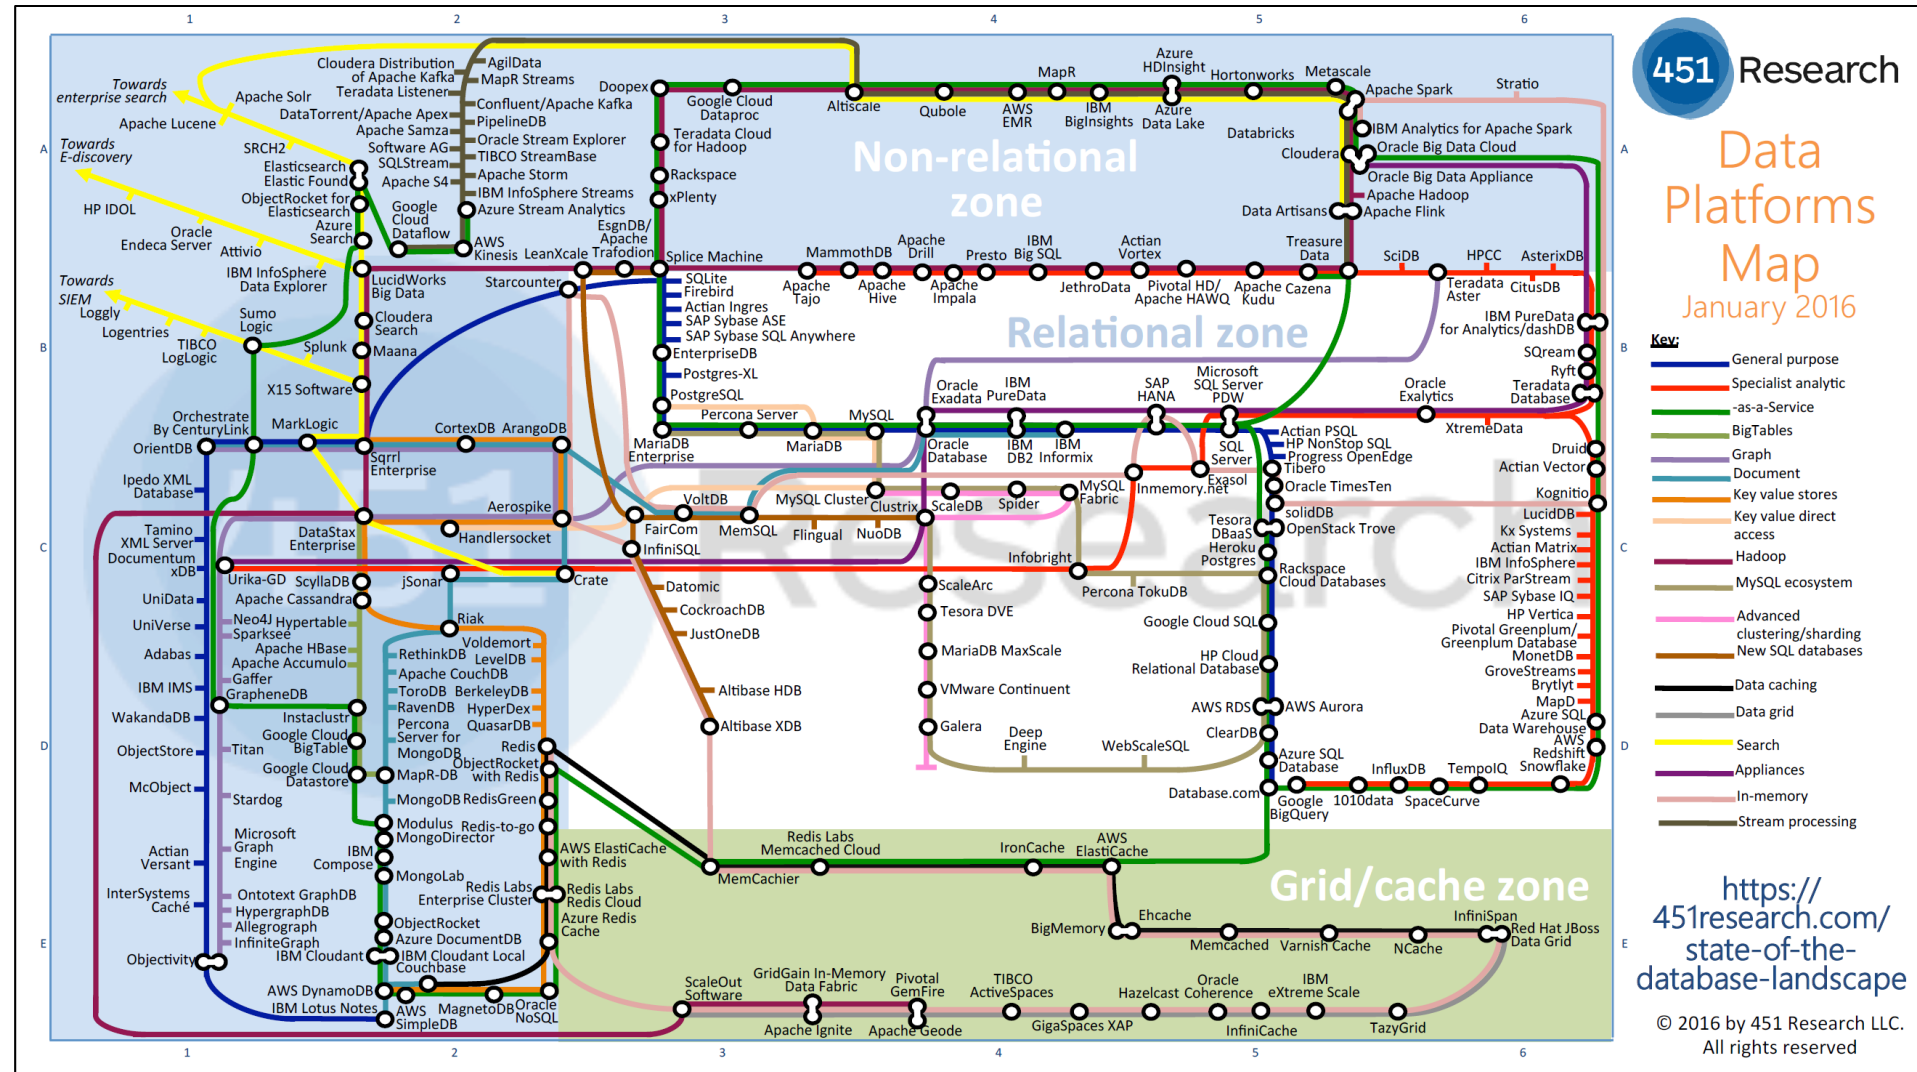


SAP customers produce about
77% of the world's beer, burgers
and deep dish pizzas!

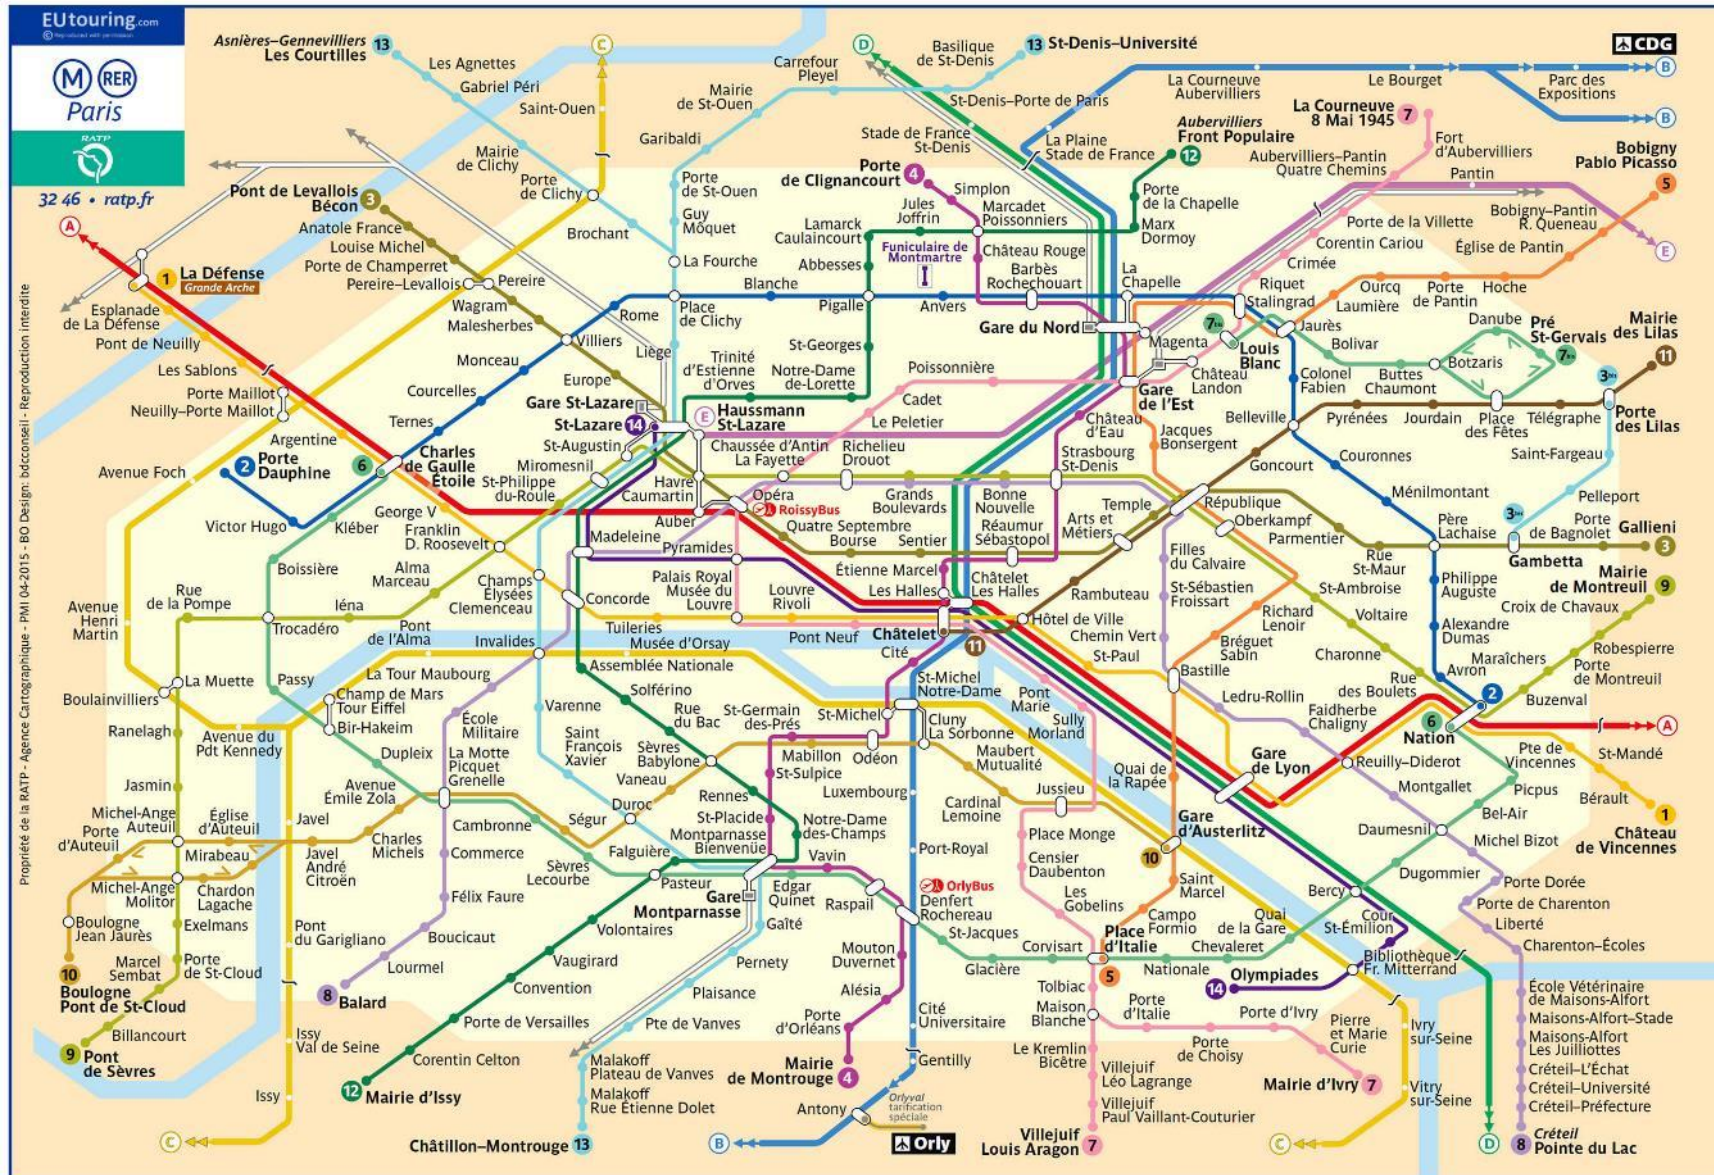
How did we get where we are with data?



The evolving database landscape map - The 451 Research Group



Hard to find your way - Just like the Metro in Paris



© 2019 SAP SE or an SAP affiliate company. All rights reserved. | PUBLIC



Polyglot Persistence as in Polyglot Programming?



Polyglot Persistence

Posted on October 15, 2008 by Scott Leberknight

In late 2006 Neal Ford wrote about [Polyglot Programming](#) in the industry to use the right language for the specific job at hand, then warring over the many different available frameworks, job rather than just the right framework(s). For a while now polyglot programming, a relational database seems to be the due to the fact that organizations have standardized on RDBMS simply what we're used to doing, and possibly we don't even have [Google Bigtable](#), [Microsoft SQL Server Data Services \(SSDS\)](#) beginning of **Polyglot Persistence** in addition to polyglot programming.

PolyglotPersistence



Martin Fowler
16 November 2011

In 2006, my colleague Neal Ford coined the term **Polyglot Programming**, to express the idea that applications should be written in a mix of languages to take advantage of the fact that different languages are suitable for tackling different problems. Complex applications combine different types of problems, so picking the right language for the job may be more productive than trying to fit all aspects into a single language.



One of the interesting consequences of this is that we are gearing up for a shift to polyglot persistence [1] - where any decent sized enterprise will have a variety of different data storage technologies for different kinds of data. There will still be large amounts of it managed in relational stores, but increasingly we'll be first asking how we want to manipulate the data and only then figuring out what technology is the best bet for it.

This chapter is from the book



NoSQL Distilled: A Brief Guide to the Emerging World of Polyglot Persistence

[Learn More](#)

[Buy](#)

13.2. Polyglot Data Store Usage

Let's take our e-commerce example and use the polyglot persistence approach to see how some of these data stores can be applied (Figure 13.2). A key-value data store could be used to store the shopping cart data before the order is confirmed by the customer and also store the session data so that

13fig03_alt.jpg (550x320) - Google Chrome
Niezabezpieczona | ptgmedia.pearsoncmg.com/images/chap13_9780321826626/elementLinks/13fig0...

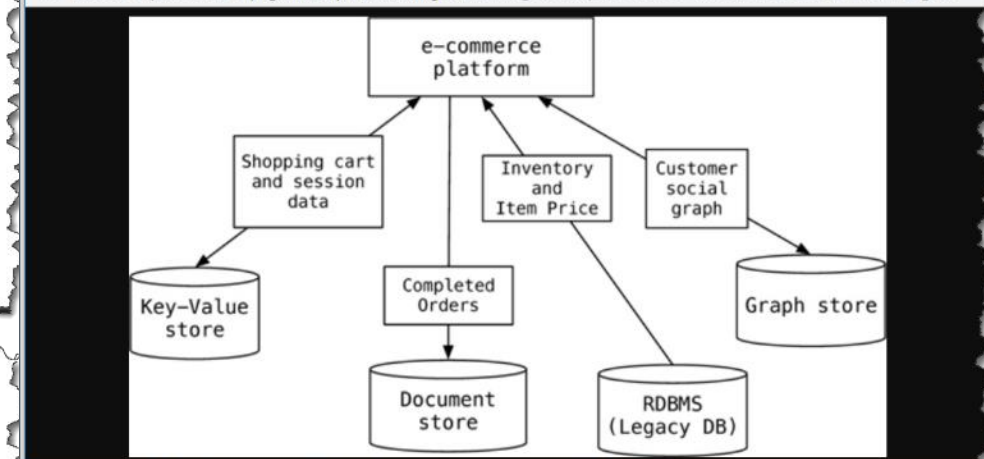


Figure 13.2: Use of key-value stores to speed session and shopping cart data storage

Is Polyglot Persistence the only option? What about multi-model database?

Multi-model database

From Wikipedia, the free encyclopedia

Most database management systems are organized around a single [data model](#) that determines how data can be organized, stored, and manipulated. In contrast, a **multi-model database** is designed to support multiple data models against a single, integrated backend.^[1] Document, graph, relational, and key-value models are examples of data models that may be supported by a multi-model database.



Multi-model database: „Neither fish nor fowl...”

Matthew Aslett - The 451 Group

One of the most complicated aspects of putting together our database landscape map was dealing with the growing number of (particularly NoSQL) databases that refuse to be pigeon-holed in any of the primary databases categories.

I have begun to refer to these as “multi-model databases” in recognition of the fact that they are able to take on the characteristics of multiple databases. In truth though there are probably two different groups of products that could be considered “multi-model”:

True multi-model databases that have been designed specifically to serve multiple data models and use-cases

General-purpose databases with multi-model options

OrientDB Community Edition

github.com/orientechnologies/orientdb



From Wikipedia, the free encyclopedia

OrientDB is an open source [NoSQL database management system](#) written in [Java](#). It is a [Multi-model database](#), supporting [graph](#), [document](#), [key/value](#), and [object](#) models,^[2] but the relationships are managed as in graph databases with direct connections between records. It supports schema-less, schema-full and schema-mixed modes. It has a strong security profiling system based on users and roles and supports querying with [Gremlin](#) along with [SQL](#) extended for graph traversal. OrientDB uses several indexing mechanisms based on [B-tree](#) and [Extendible hashing](#), the last one is known as "hash index", there are plans to implement [LSM-tree](#) and [Fractal tree index](#) based indexes. Each record has [Surrogate key](#) which indicates position of record inside of [Array list](#), links between records are stored either as single value of record's position stored inside of referrer or as [B-tree](#) of record positions (so-called record IDs or RIDs) which allows fast traversal (with $O(1)$ complexity) of one-to-many relationships and fast addition/removal of new links. OrientDB is the third most popular graph database according to the [DB-Engines](#) graph database ranking^[3], as of September 2017.

The development of OrientDB still relies on an open source community led by OrientDB LTD company created by its original author Luca Garulli. The project uses [GitHub](#)^[4] to manage the sources, contributors and versioning, [Google Group](#)^[5] and [Stack Overflow](#)^[6] to provide free support to the worldwide users. OrientDB also offers a free [Udemy](#) course^[7] for those hoping to learn the basics and get started with OrientDB.

orientechnologies / orientdb

Watch 304

Star 3,824

Fork 785

Code

Issues 1,319

Pull requests 40

Projects 1

Wiki

Insights

OrientDB is the most versatile DBMS supporting Graph, Document, Reactive, Full-Text, Geospatial and Key-Value models in one Multi-Model product. OrientDB can run distributed (Multi-Master), supports SQL, ACID Transactions, Full-Text indexing and Reactive Queries. OrientDB Community Edition is Open Source using a liberal Apache 2 license. <http://orientdb.com>

orientdb

graph-database

document-database

dbms

nosql

graph-store

multi-master

multi-model-dbms

database

performance

fast

sql

17,808 commits

71 branches

142 releases

121 contributors

Apache-2.0

Branch: develop

New pull request

Create new file

Upload files

Find File

Clone or download



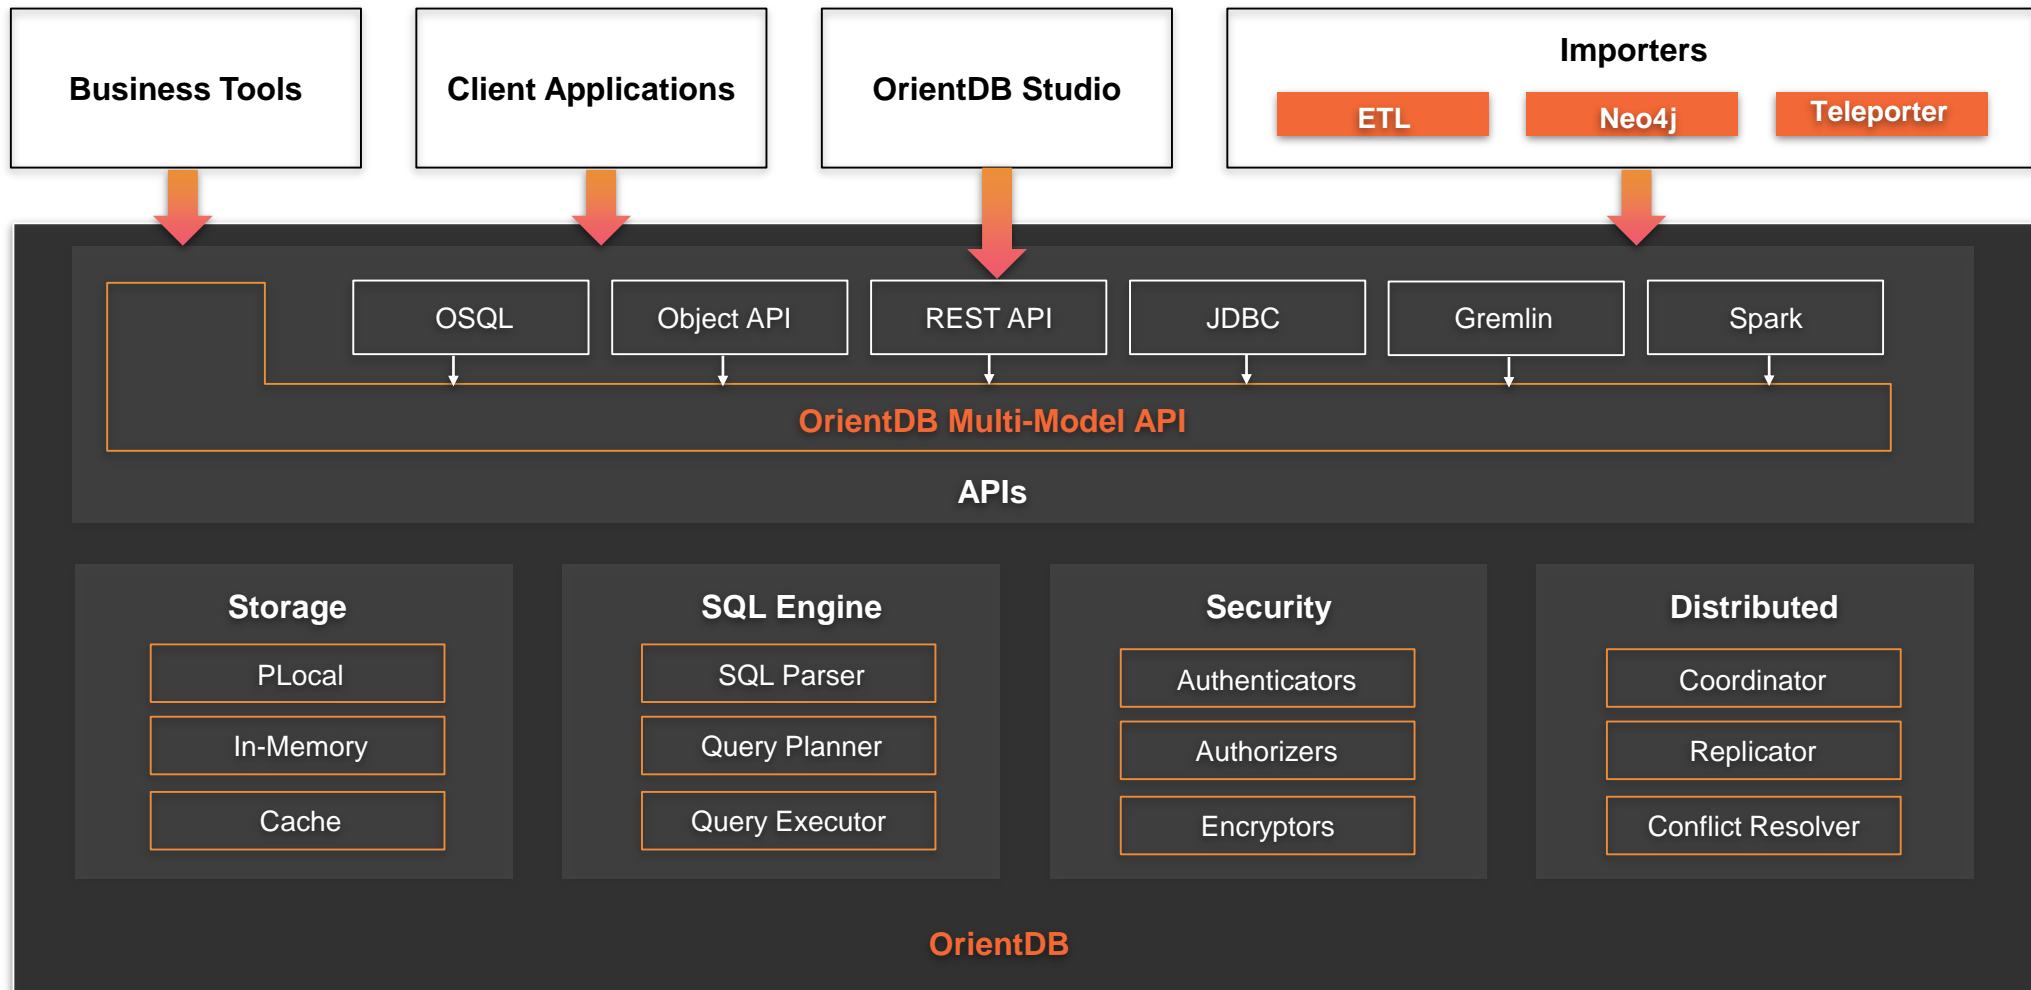
luigidellaquila Enhance quorum failure logging (distributed)

Latest commit 2491a38 3 days ago




Feature Matrix

	Community
License	Apache 2 license
Embedded, Remote and In-Memory Storage	✓
Studio Visual tool and Console tool	✓
High Availability Replication, Sharding and Fail-over	✓
Standards: SQL, Blueprints, Gremlin, Pipes, SparQL	✓
Teleporter – synchronization mode	
Metric Analysis (Monitor and tune your servers), Query Profiler (Locate 80% of the bottlenecks in your application), Configurable Alerts (Email alerts on executing HTTP requests), Scheduler for Full and Incremental Non-Stop Backups, Auditing to track all the access to your data	
10/5 & 24/7 Production Support Available	
Price	FREE for any use, even Commercial. No hidden restrictions.
	Download

OrientDB Community Edition



Learn More: OrientDB at Udemy


CategoriesUdemy for BusinessBecome an instructor

Log InSign Up

2 results for **orientdb**


Ratings ▾

Duration ▾

 All Filters

Sort ▾

Did you mean: [oriented](#)





A Complete Guide to OrientDB: A NoSQL Database

NEW 48 lectures • 3 hours • All Levels

Learn **OrientDB** from Beginning to Advanced | By Elementary Learners

zł34.99
~~zł609.99~~

 0.0
(0 ratings)




OrientDB - Getting Started

21 lectures • 2 hours • Beginner

A full introductory course to **OrientDB** 2.0, the Document-Graph database. Learn about NoSQL's latest trend: GraphDBs! | By **OrientDB** Team

Free




 4.2
(2,583 ratings)

Can't decide?

Every course comes with a **30-day money-back guarantee**

Top companies trust Udemy

Get your team access to Udemy's top 80,000+ courses



[Try Udemy for Business](#)

SAP HANA

Not Only Multi-Model!



From Wikipedia, the free encyclopedia

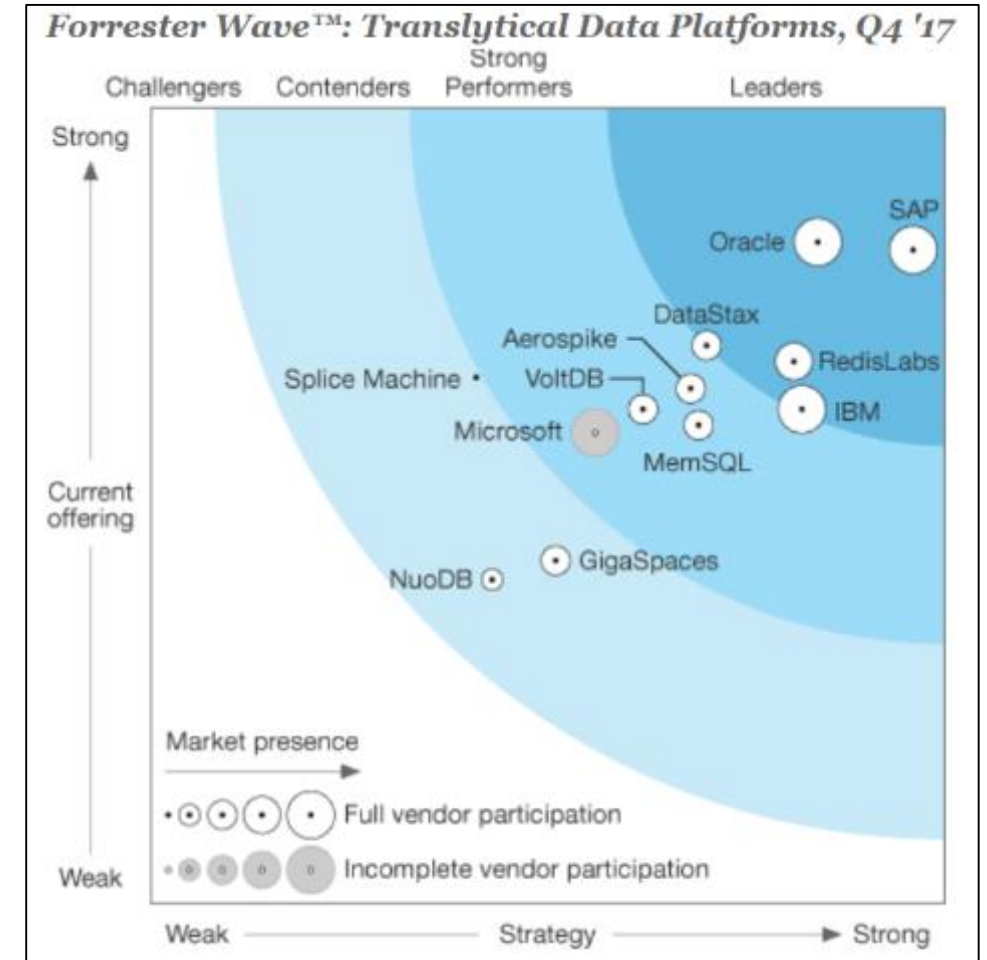
SAP HANA is an in-memory, column-oriented, relational database management system developed and marketed by SAP SE.^{[2][3]} Its primary function as a database server is to store and retrieve data as requested by the applications. In addition, it performs advanced analytics (predictive analytics, spatial data processing, text analytics, text search, streaming analytics, graph data processing) and includes extract, transform, load (ETL) capabilities as well as an application server.

From Wikipedia, the free encyclopedia

Hybrid transaction/analytical processing (HTAP) is an emerging application architecture that "breaks the wall" between transaction processing and analytics. It enables more informed and "in business real time" decision making.^{[1][2]}

“...**Translytical data platforms**, an emerging technology, deliver faster access to business data to support various workloads and use cases. EA pros can use them to drive new business initiatives...”

“SAP HANA is a shared-nothing, in-memory data platform, the core of SAP's translytical platform, which supports many use cases, including real-time applications, analytics, translytical apps, systems of insight, and advanced analytics...”

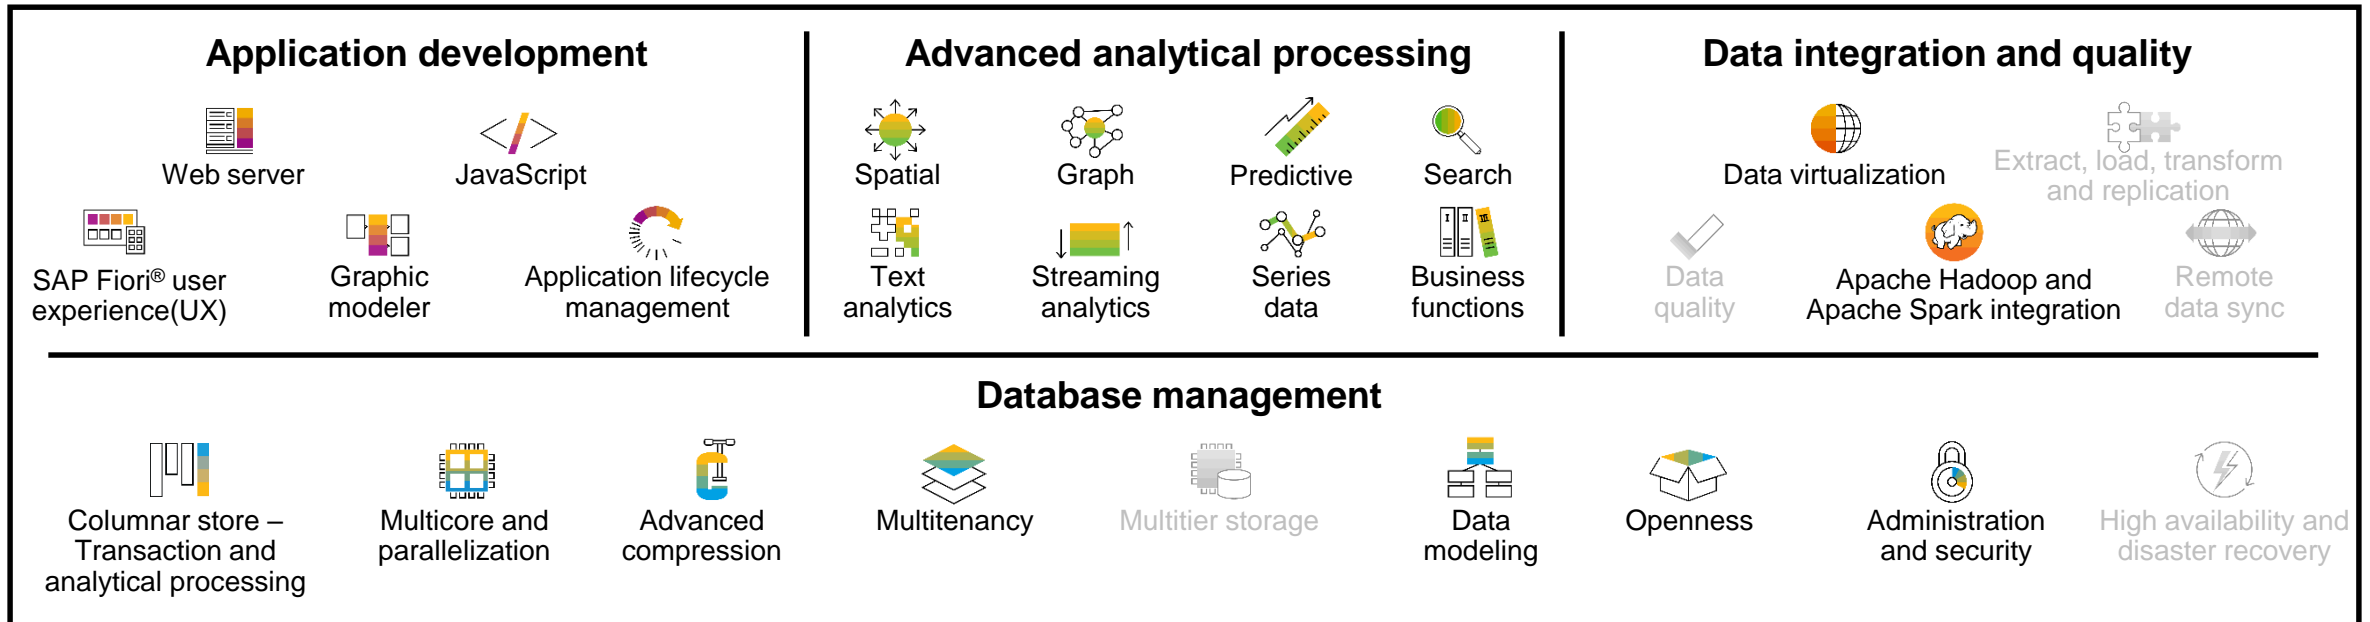


SAP HANA, express edition

developers.sap.com/sap-hana-express



- Free pre-packaged and self learning tutorials, use cases, and sample code available from SAP Developer Center
- Enable fast application development and deployment with essential SAP HANA features
- Free productive use up to 32 GB of RAM
- Comes as a binary installer, as a pre-configured virtual machine image (ova file), as a docker or cloud image (with launchers) - No certified HW needed
- Flexible access from laptop, personal computer, server, or cloud



Learn More: SAP HANA at openSAP

Full-Text Search with SAP HANA

- The Basics: Full-Text Indexing
- Matching: Fuzzy Search and Search Rules
- Search Models
- API and Query Language
- SAPUI5 Search User Interface
- SAP HANA Enterprise Search

<https://open.sap.com/courses/hshesh1>

Analyzing Connected Data with SAP HANA Graph

- Connected Data and SAP HANA Graph
- The Basics: Nodes, Edges, and Workspaces
- Pattern Matching
- Built-In Algorithms
- GraphScript
- SAP HANA Hierarchies

<https://open.sap.com/courses/hsgra1>

Spatial Analysis with SAP HANA

- Introduction to SAP HANA Spatial
- Spatial Analysis
- Development of Spatial Applications

<https://open.sap.com/courses/hsgs1>

And many more!

Thank you!

Merci! Dzięk!

Visit the SAP Developer Center for free tools and resources:

 developers.sap.com

Follow us:



[@sapdevs](https://twitter.com/sapdevs) on Twitter



[SAP Developers](https://www.facebook.com/SAPDevelopers) on Facebook



[SAP Developers](https://www.youtube.com/SAPDevelopers) on YouTube



sap.github.io on GitHub

Abdel Dadouche



Vitaliy Rudnytskiy






Please

**Remember to
rate this session**

Thank you!



goto;
chicago



Click 'Rate Session'

Rate **5** sessions to get the
supercool GOTO reward