

The Amadeus NoSQL saga from 1992 to 2016: highlights

Abstract

- In this talk, Jeremy will lead you through the different steps of NoSQL within Amadeus - one of the largest European data processing centers, handling billions of database actions per day over massive critical data updated in real time.
- With intent to pragmatically assist you to assess what would be the right NoSQL strategy for your business, spotlight will be on showing the successive maturity steps and on sharing experiences with IT enthusiasts to take the most benefit from NoSQL - in case, it fits! Indeed, you will as well receive hints on how to spot if NoSQL actually deserves interest from you.
- To conclude, Jeremy will share the R&D Amadeus view for the coming years for NoSQL technologies within Amadeus IT group and within the IT industry.

<http://www.telecom-valley.fr/pagecmsblog/01-07-big-data-nosql-bi-securite-cloud>

<http://www.telecom-valley.fr/pagecmsblog/j%C3%A9r%C3%A9my-meyer>

amadeus

The Amadeus NoSQL saga

from 1992 to 2016:
highlights

Jeremy Meyer

Architecture, Quality and Governance
Data Store Middleware

Telecom Valley SophiaConf 2014
Sophia Antipolis, 1st of July 2014

1 Amadeus


© 2014 Amadeus IT Group SA

amadeus

Amadeus in a few words

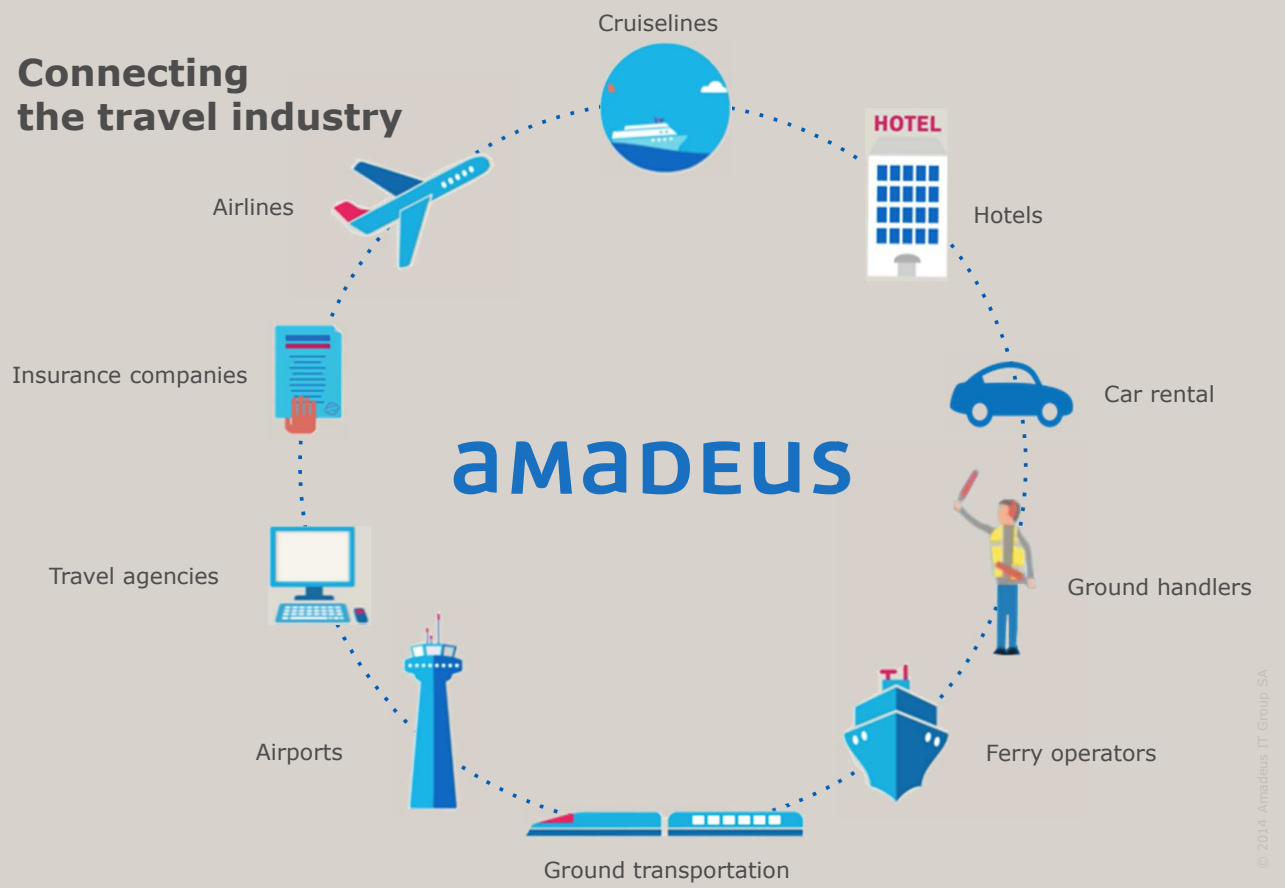
Amadeus is a technology company dedicated to the
global travel industry

We are present in **195 countries**
with a worldwide team of more than **11,000** people



© 2014 Amadeus IT Group SA

Connecting the travel industry



5

amadeus

Our commitment to innovation

— Amadeus has invested **€2.9bn** in
Research & Development
since 2004



Robust global operations

_ Amadeus today



2 Our view on NoSQL

Our definition of NoSQL



NoSQL movement is about

- promoting *alternative* storage solutions to *complement relational DBs*
- turning them into commodity *products* and *services*

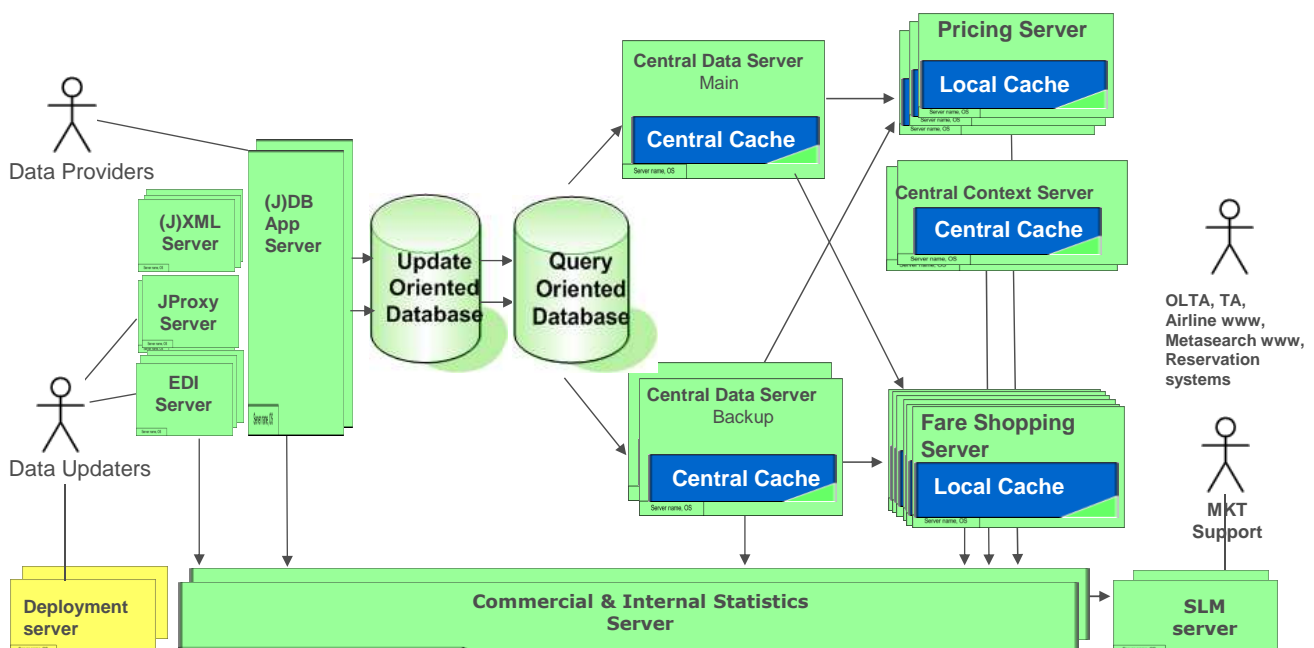
— NoSQL is not “new”

- NoSQL in Amadeus since 1992 through IBM TPF mainframes for Bookings
- Home-made
 - two-layer object cache since 2000 for Pricing and Shopping services
 - in-memory GraphDB since 2001 for building routes of the Journeys
- Memcached widely used since 2009 for Flights Inventory

9

amadeus

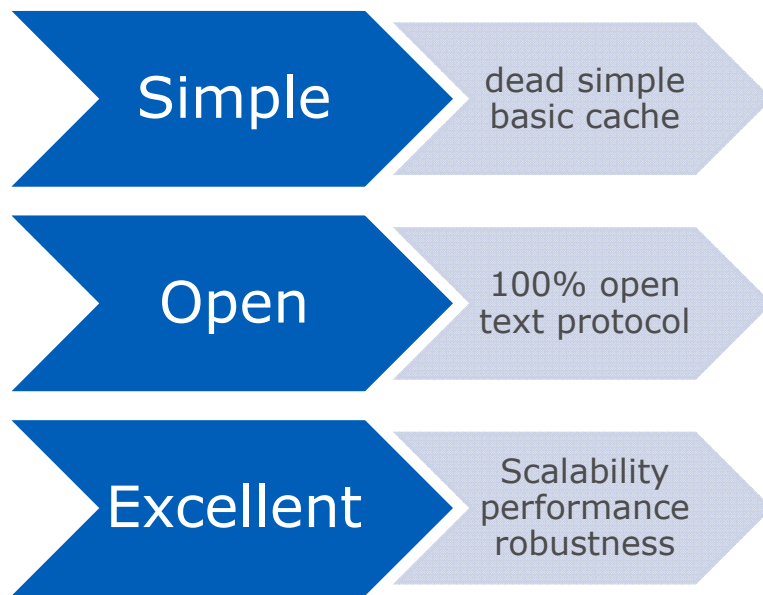
2001: Pricing and Shopping 2-level Key/Value object cache



10

2009: Memcached in Amadeus

First open-source NoSQL solution in Amadeus



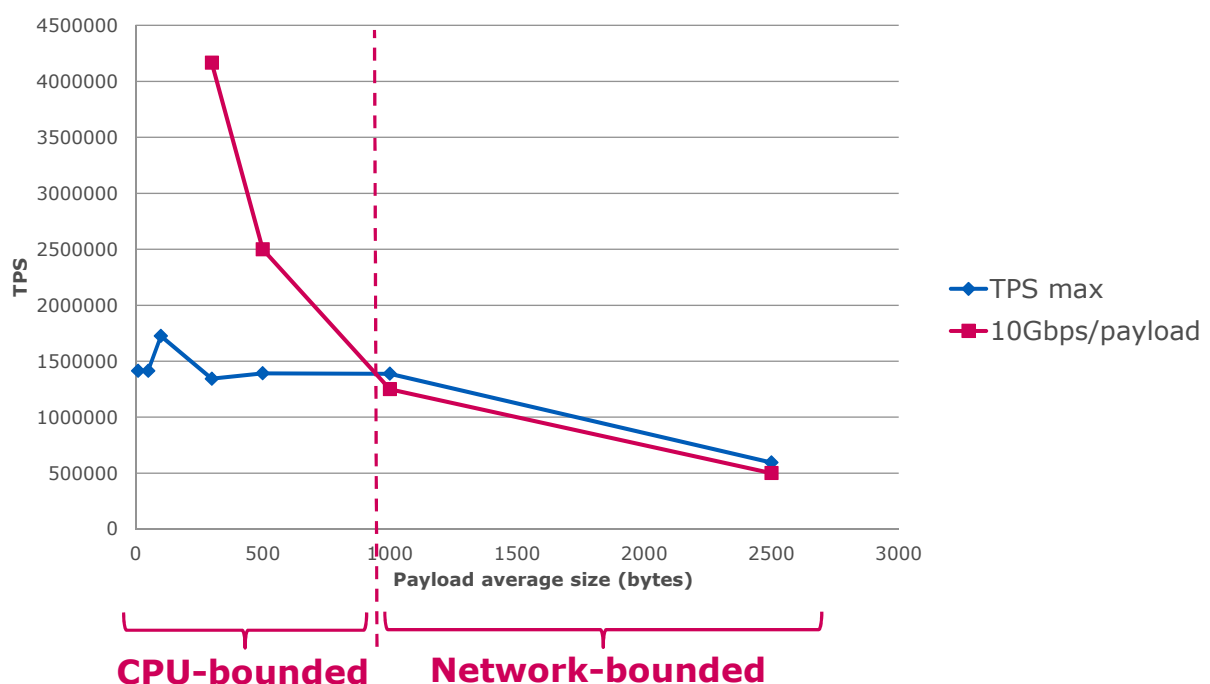
Accurate representative of NoSQL products !

11



amadeus

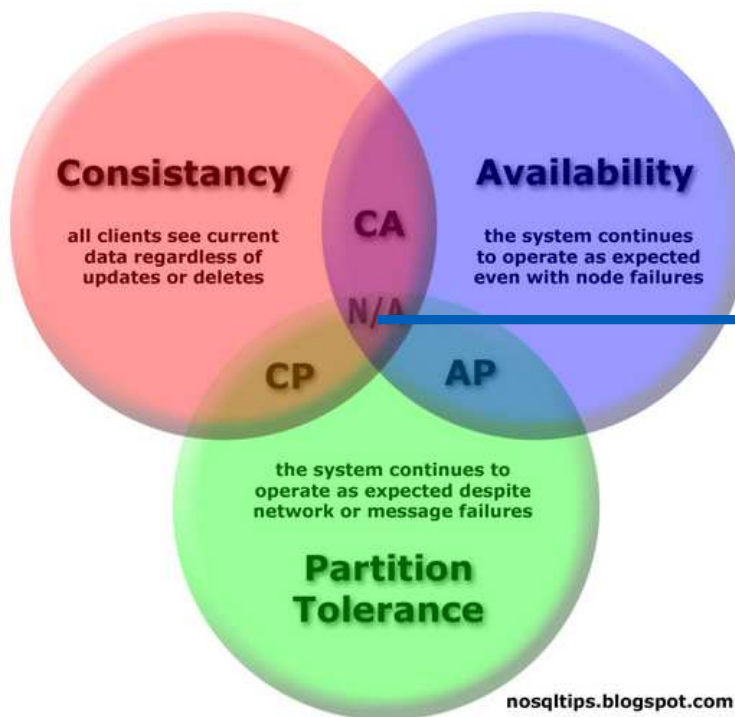
Max throughput of a single Memcached process over 10Gig network



amadeus

NoSQL roots

CAP theorem, E. Brewer, 2000*



→ Pick two!

13 *Symposium on Principles of Distributed Computing

amadeus

NoSQL roots



From academia

- Mike Stonebraker: *"One size does not fit all"**
- Ingres, Postgres, Informix Illustra, Streambase, Vertica, VoltDB

The big four of the web industry: the "GAFA"

- Pragmatic people with immediate & tangible problems
- RDBMSs do not scale-out for massive write operations

Google



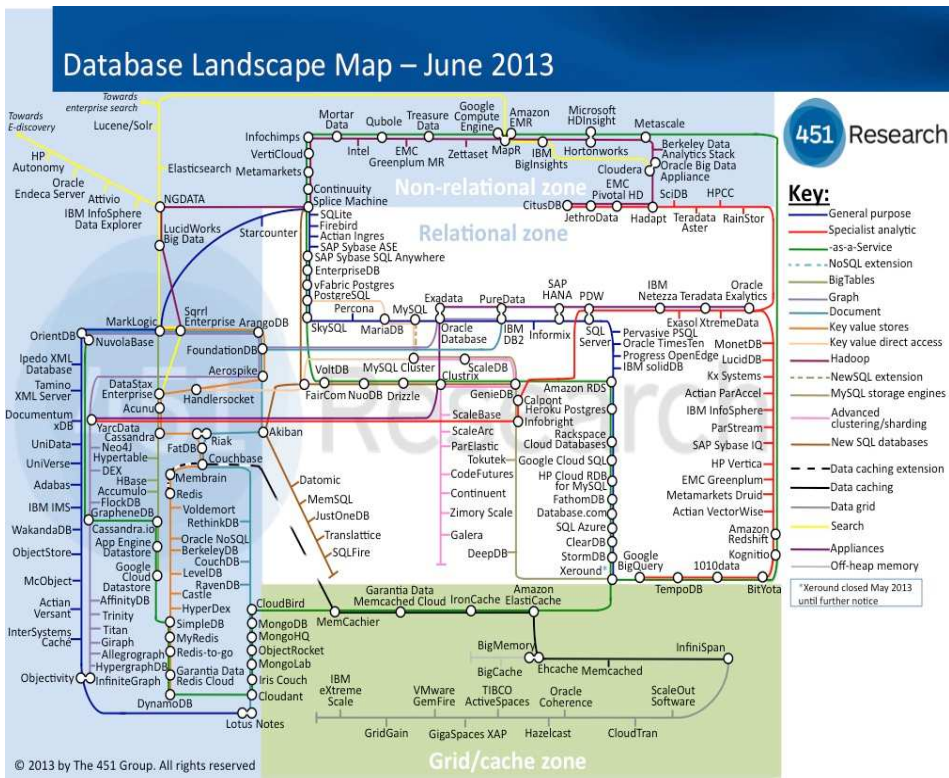
facebook

amazon

14 *21st International Conference on Data Engineering, 2005

amadeus

NoSQL roots



150+ products !

15

amadeus

With NoSQL, we got

- Scalability and Performance
- Eventual Consistency
 - with freedom of data-persistency choice
- Many nodes to administrate instead of a pair
 - + network infrastructure to be adapted
- Access
 - to source code of the solution
 - to active communities of enthusiastic people
- Technologies with clear can-do/cannot-do boundaries
 - Improve app designs



amadeus

With NoSQL, we lost

- Relational database core abilities
 - Large and flexible feature set
 - Convenient query language
 - ACID transactions
- Decades of investments spent in enterprise RDBMSs
 - Advanced admin/tuning/maintenance facilities
 - Mature security features
- Wealth of trained people all over the world



NoSQL

Enabler of web-scale apps

Not suitable for all purposes

For people who know “what they do”

3

NoSQL implemented in Amadeus

© 2014 Amadeus IT Group SA

amadeus

NoSQL & Amadeus: 2013

Setting-up our collaboration models with NoSQL providers

— Comparison of Key/Value store providers + choice

- Redis, JBoss DataGrid "Infinispan", Oracle Coherence
- Couchbase

— First experimentations with Mongo Document Store

— Definition of our way-of-working with software providers


- Major partners, driving how the products evolve
- Benefit from consulting, training, support services

— Strong relationship tied with Couchbase

- Gave talks to Couchbase UK conference of 2013 and 2014, *see youtube*
- Our Executive Committee members visited them in the SiliconValley

Couchbase throughput test

→ requested and got Couchbase enhancement to multithread write-to-disk



21

amADEUS

Our 3 NoSQL picks

- Sub-millisecond latency
- Extreme throughput
- Extreme scalability
- Basic query language

Recommended use cases
massive simple read/write
(millions per second)

Low-Latency
Key/Value
Store



- Rich query model
- Rich data modelling abilities
- Developer friendly
- Suitable for asymmetric workloads: Write << Read

Recommended use cases
Small to medium OLTP apps

Document
Store



- Huge amount of data (>10 TB)
- Strong data-crunching abilities

Recommended use cases
Batch processing
Data analytics
Long term archiving

BigData



NoSQL & Amadeus: 2014

NoSQL to enter Amadeus technology portfolio

- One pilot app for each selected NoSQL area
 - Couchbase pilot: [Air Availability](#)
 - MongoDB pilot: [Accounting](#)
 - Hadoop/MapR pilot: [Searching and Shopping Logging](#)
- Learn how to [deploy](#) and [operate](#) NoSQL
 - Non-Functional Requirements are [key](#)
- Deploy NoSQL farms with [guidelines defined/validated during pilots](#)
 - Ex: [we like big large servers](#), 4 sockets, 12-core CPU, FusionIO PCIE cards
- [100+ persons](#) involved on NoSQL in Amadeus
 - 20+ groups of R&D
 - 10+ groups of Operations

Page 23

amadeus


Air Availability

part of Searching and Shopping stack

The screenshot displays the Amadeus Air Availability interface. On the left, a sidebar shows flight details for 'mercredi 12 août 2009' and 'mercredi 19 août 2009'. The main area shows a list of flights from Montreal, Trudeau, QC (YUL) to Paris, Charles De Gaulle, FR (CDG). The flights are sorted by price (low to high), with the lowest price being \$1093+. The interface includes various filters such as 'Stops', 'Times', and 'Airports'. A calendar on the right shows the price for the selected dates. The bottom right corner features a red circular logo with a white 'G' inside.

Our latest in-house NoSQL product: ProxyD

Scalability: need to multiplex network connections



- 100 times less network connections
- Based on Twemproxy source code from [Twitter](#)
- We submitted [code back to community](#)

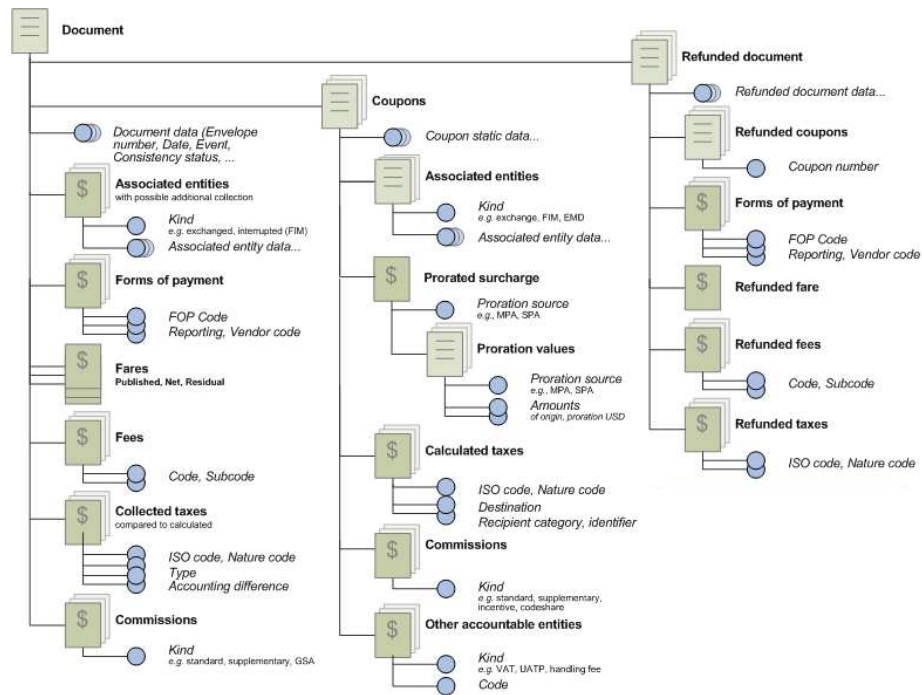
Accounting: ad-hoc searches

- [Datamart](#): fed in real time, 3 years history, 6TB
- [Web UI](#): chaining queries to drill data
- [Interactive](#) results
- [MassivelyParallelProcessing](#) way to deploy Mongo
 - [Maximization](#) of resource usage, 1 MongoDB per CPU core
 - [Multiple massive data transformations](#) through aggregation framework



Accounting: ad-hoc searches

Data model: hierarchical documents




1000+
different
document
"schemas" !



Page 27

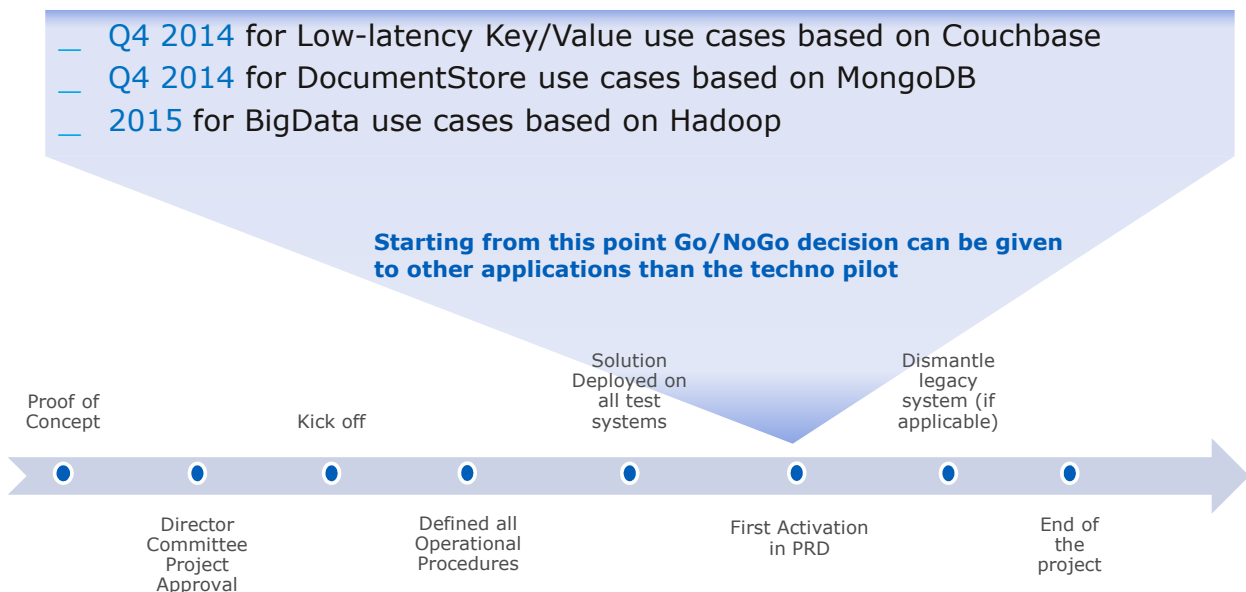
amadeus

BigData Logging Architecture



NoSQL standards in Amadeus

Rationale and first prod-activation timeline



NoSQL: what to look at before embracing?

- NoSQL Tradeoffs: what to relax?
 - Availability, Consistency, Tolerance to NW partitioning
 - Transactional needs?
- Data Storage
 - Structure and volume of your data?
 - Is NoSQL master? How is NoSQL fed? Migration steps?
- Data Access Paths: how data is accessed?
 - Workload: Read/Update/Insert mix? OLTP? Analytics?
 - Latency and throughput?
 - How to handle data inconsistencies?
- Business sensitivity of the data? Security regulations?
- Business Case: cost vs profit (money, stability, ...)
 - HW/SW cost of a NoSQL farm can be higher than expected
 - Needed to be estimated at the very beginning of your project

4 — What's next?


© 2014 Amadeus IT Group SA

amadeus

Ultra High Availability through NoSQL

Active/Active topology to achieve target of **max** 3mins of outage per year

__Active/Active Data Centers



__Eventual Consistency issues to be managed

BigData



Thomas H. Davenport



Hervé Couturier

[At the Big Data Crossroads: turning towards a smarter travel experience](#) report



5 Conclusion

NoSQL: a major growth enabler

Strong internal
momentum to
adopt NoSQL

**Wide and
controlled NoSQL
standardisation**
going-on

NoSQL is a **key
enabler** for
Amadeus to
Shape the **Future
of Travel**

© 2014 Amadeus IT Group SA

_____ Thank you

You can follow us on:
AmadeusITGroup



amadeus.com/blog
amadeus.com

© 2014 Amadeus IT Group SA