

BRUGG LAUSANNE ZÜRICH DÜSSELDORF FRANKFURT A.M. FREIBURG I.BR. HAMBURG MUNICH STUTTGART VIENNA BASEL BERN



SQL versus NoSQL 3rd July 2014



Guido Schmutz

- Working for Trivadis for more than 17 years
- Oracle ACE Director for Fusion Middleware and SOA
- Co-Author of different books
- Consultant, Trainer Software Architect for Java, Oracle, SOA and Big Data / Fast Data
- Member of Trivadis Architecture Board
- Technology Manager @ Trivadis
- More than 25 years of software development experience
- Contact: guido.schmutz@trivadis.com
- Blog: <u>http://guidoschmutz.wordpress.com</u>
- Twitter: <u>gschmutz</u>

2014 © Trivadis SQL versus NoSQL 3rd July 2014











Oracle Service Bus 11g Service Oriented Architecture Development Cookbook An Integration Blueprint

IDACKTI enterprise³⁸ Guide Schmutz Das

Guido Schmutz Daniel Liebhart (PACKT) enterprise 8 Peter Welkenbach

trivadi

trivadis makes IT easier.

Philipp Salvisberg

- With Trivadis since April 2000
 - Senior Principal Consultant, Partner
 - Member of the Board of Directors
 - philipp.salvisberg@trivadis.com
 - www.salvis.com/blog
 - @phsalvisberg
- Member of the **trivadis**





- Main focus on database centric development with Oracle DB
 - Application Development
 - Business Intelligence
 - Application Performance Management
- Over 20 years experience in using Oracle products





1. Motivation

- 2. Overview of SQL and NoSQL Data Stores
- 3. Use Cases Let's Get Ready to Rumble
- 4. Recommended Reads
- 5. Core Messages







- 1. Motivation
- 2. Overview of SQL and NoSQL Data Stores
- 3. Use Cases Let's Get Ready to Rumble
- 4. Recommended Reads
- 5. Conclusion

trivadis makes IT easier.

SQL Data Stores

- Relational Model
- Standardized, SQL:2011 is the 7th major revision since SQL-86
 - 9 parts, more than 4000 pages
 - But no single database implements all standards/features
- Rich set of features
 - Incl. SQL/PSM, SQL/MED, SQL/XML, SQL/RPR, Temporal Features
 - Incl. User-defined Types and Collection Types (since SQL:1999)
- ACID Transactions
 - Atomicity: all or nothing
 - Consistency: from valid state to valid state considering constraints, triggers, ...
 - Isolation: result is not affected through concurrent execution
 - Durability: committed data stays available after crash, power loss or errors
- Good support by different languages, frameworks and tools
- Good understanding of basic concepts by IT professionals





NoSQL Definition

- Next Generation Databases mostly addressing some of the points:
 - being non-relational,
 - distributed,
 - open-source and
 - horizontally scalable.
- Often more characteristics apply such as:
 - schema-free,
 - easy replication support,
 - simple API,
 - eventually consistent / BASE (not ACID),
 - a huge amount of data
 - and more.

BASE

- Basically Available: Availability is more important than consistency
- **S**oft State: Higher availability results in an eventual consistent state
- Eventually Consistent: If no new updates are made to a given data item, eventually all accesses to that item will return the last updated value
- The misleading term "nosql" (the community now translates it mostly with "not only sql") should be seen as an alias to something like the definition above

Source: http://nosql-database.org



Brewer's CAP Theorem

Any networked shared-data system can have at most two of the three desirable properties:

- Consistency All of the nodes see the same data at the same time, regardless of where the data is stored
- Availability Node failures do not prevent survivors from continuing to operate
- Network Partition tolerance The system continues to operate despite arbitrary message loss

2014 © Trivadis SQL versus NoSQL 3rd July 2014



makes IT easier





Standardized Model, Tooling, Complexity



2014 © Trivadis SQL versus NoSQL 3rd July 2014



- 1. Motivation
- 2. Overview of SQL and NoSQL Data Stores
- 3. Use Cases Let's Get Ready to Rumble
- 4. Recommended Reads
- 5. Core Messages



Round 1 Smart Meter



2014 © Trivadis SQL versus NoSQL 3rd July 2014



Smart Meter – Introduction

- Data is captured from 2 million smart meters
 (1 main sensor, 10 optional sub-sensors, e.g. fridge, stove, dish-washer, wash-machine, TV, computer, ...)
 - One smart meter per household
 - Measures the energy consumption
 - Delivery interval between 1 second and 5 minutes
- Every sensor measures and delivers the energy consumption per second (kWh)
- AP Characteristics (Consistency is not so important)





Smart Meter – Customer Dashboard





2014 © Trivadis SQL versus NoSQL

3rd July 2014

Smart Meter – Use Cases

Smart Meter Store data of millions of sensors Household Dashboard

- Store sensor and its sub-sensor values
 - 2 Mio sensors, up to 10 sub-sensors
 - Energy consumption per second per sensor (kWh)
 - Delivery interval between 1 second and 5 minutes
- Query usage per sensor and its sub-sensors to visualize a time series on a customer dashboard
 - Available in different granularities, values are aggregated in
 - Minute
 - Quarter of hour (15-minutes)
 - Hour
 - Day
 - Responsive UI



Cassandra NoSQL Datastore



- Wide-Column Store
- Developed at Facebook
- Professional grade support from DataStax
- Main Features
 - Real-Time
 - Highly Distributed
 - Support for Multiple Data Center
 - Highly Scalable
 - No Single Point of Failure
 - Fault Tolerant
 - Tunable Consistency
 - CQL Cassandra Query Language







The Cassandra Way

Household	Bucket									
AFG10	MINUTE-2014/03/5	sensor	1	1	1		2		•••	
		at	24	າ * 60m * ∶	11 sensor	= 1	5'840 cols	6	•••	
		kwh	7.05	7.10	8.11		6.95	7.04		
AFG10	QHOUR-2014/03	sensor	1	1	1		2		•••	
		at	30d *	24h * 4q	* 11 senso	or =	31′680 co	'680 cols		
			5	0	5		5	C		
		kwh	105.78	104.73	102.29	•••	102.78	121.61	•••	
AFG10	HOUR-2014/03	sensor	30	d * 24h *	11 sensor	= 7	'920 cols		•••	
		at	5T11	5T10	5T09	•••	5T11	5 10	•••	
		kwh	423.00	410.33	395.99	• • •	598.32	572.12	•••	
AFG10	DAY-2014	sensor	nsor 365d * 11 sensor = 4'011 cols					•••		
		at	5T	3T	2T	•••	5T	4	•••	
		kwh	10100.2	9892.2	8987.4	•••	879.8	912,3		
GXK11	MINUTE-2014/03/5	sensor	1	1	1	• • •	2	2	•••	
		at	11:59	11:03	11:04	• • •	11:01	11:02	•••	
Growth		kwh	100.10	90.88	95.00	•••	92.50	88.50	•••	

7

2014 © Trivadis

SQL versus NoSQL 3rd July 2014



The Cassandra Way

Household	Bucket							
AFG10	MINUTE-2014/03/5	sensor	1	1	1	 2	2	
		at	11:59	11:58	11:57	 11:59	11:58	
		kwh	7.05	7.10	8.11	 6.95	7.04	
AFG10	QHOUR-2014/03	sensor	1	1	1	 2	2	
		at	5T11:45	5T11:30	5T11:15	 5T11:45	5T11:30	
		kiech	105 38	104 33	103 38	103 38	101 61	
		at	5111:45	5711:30	5111:15	 5T11:45	5111:30	

CREATE TABLE mete	r_reading_timeunit (
household_id	uuid,
bucket_id	text,
at_timestamp	timestamp,
sensor_id	bigint,
kwh_consumed	counter,
PRIMARY KEY((hous	<pre>ehold_id, bucket_id), sensor_id, at_timestamp))</pre>
WITH CLUSTERING O	RDER BY (sensor_id ASC, at_timestamp DESC);

```
UPDATE meter_reading_timeunit
SET kwh_consumed = kwh_consumed + 10010
WHERE household_id = 2dc487f0-b271-11e3-a5e2-0800200c9a66
AND sensor_id = 1
AND bucket_id = 'MINUTE-2014/03/23/11'
AND at_timestamp = '2014-03-23T11:01:00';
```



The Cassandra Way

Household	Bucket							
AFG10	MINUTE-2014/03/5	sensor	1	1	1	 2	2	
		at	11:59	11:58	11:57	 11:59	11:58	
		kwh	7.05	7.10	8.11	 6.95	7.04	
AFG10	QHOUR-2014/03	sensor	1	1	1	 2	2	
		at	5T11:45	5T11:30	5T11:15	 5T11:45	5T11:30	
		kinch	105 38	104 33	103 38	103 38	101 61	
		at	5711:45	5111:30	5111:15	 5T11:45	5111:30	

- 288 nodes on EC2
- Over 1 Mio writes/sec => 60 Mio writes/min
- Rolling counters, always up to date

```
select household_id, bucket_id, at_timestamp, sensor_id, kwh_consumed
from meter_reading_timeunit
where household_id = 2dc487f0-b271-11e3-a5e2-0800200c9a66
and bucket_id = 'MINUTE-2014/03/23/11'
and sensor_id = 1
and at_timestamp > '2014-03-23T11:00:00'
order by sensor id, at timestamp DESC;
```



Relational Architecture



- Active Data Guard Configuration
- Global Data Services redirects requests based on
 - Server loads
 - Request type (read/write)
- Reader farm is geographically spread
- Failover/switchover to any node in the reader farm is possible
 - Read services are not affected
 - Write services are unavailable for a short period of time
- Scalability of the write services is the bottleneck of the system



2014 © Trivadis SQL versus NoSQL 3rd July 2014

Relational Data Model



- SENSOR_READINGS_...
 - Index-organized tables
 - Daily partitions
- JDBC Batch Merges
 - A transaction per sensor delivery
 - A single network roundtrip to merge 55 readings of a sensor delivery
 - Average between
 - 0.4 Mio tpm (delivery per 5 minutes)
 - 120 Mio tpm (delivery per second)
 - Top TPC-C Benchmark: 8.5 Mio tpm
- Batch job to aggregate readings every 15 minutes, avoiding intermediate results (updates)
 - Quarter of hour (5760 times a day)
 - Hour (24 times a day)
 - Day (once a day)



Query Sensor Data – The SQL Way

Use aggregate tables to change granularity (quarter of hours, hours, days)



2014 © Trivadis SQL versus NoSQL 3rd July 2014

Smart Meter

0 – 1







2014 © Trivadis SQL versus NoSQL

3rd July 2014

Round 2 Order Entry



2014 © Trivadis

SQL versus NoSQL 3rd July 2014

Order Entry – Introduction

- Orders are entered via GUI
 - Orders may be modified as long as the order status is "incomplete"
 - Order status is set to "complete" after a final approval
- A lot of data is available in the system
- Some sales volume analysis are wanted
- CA Characteristics (network partition tolerance is not so important)



SQL versus NoSQL 3rd July 2014

2014 © Trivadis

Order Entry – Example

Accpac/portal60a/po	rtal.html	💌 🐓 🔀 Google	ρ-
File Edit View Favorites Tools Help			
🖕 Favorites 🛛 🍰 🙋 bcr505522-SageCRM 🤌 Sage Web Too	kit Widget Sh 🥍 Sage Accpac 🙋 vpn.na.sage 🙋 Sage Software Organization		
Sage Accpac		👌 • 🗟 · 🗖 🖷	• Page • Safety • Tools • 🔞 •
A Sage Accpac	Welcome A	Administrator This Session: 04/30/2020 SAMINC C	hange Password Help Sign Out
Tasks Reports Inquiry Administratio	n		
MY Shortcuts	E Order Entry		inancial entr
Home Order No. 14	QT00000000003 ► ► Q D No. of Shipments 0 Last S	hipment No.	
Add Snapshots Customer No. 1 Order Qustom Femplate Cod Order Date	200 Cartrive Cartrie Carteria	sted Source: Entered	2
 What's New in this Version? Customize Your Home Page Add a Snapshot to Your Home Create a Query Visit the Sage Accpac Commu Search the Help System 	Quote Expiration Date 06/30/2010 Uob Related on Quote Exp. Ship Date 05/31/2010 Image: Calc. Tax Quote 3 with header comments Reference Image: Calc. Tax e Q. Item No / Misc. Charge Q. Kt/BOM Description A1-105/0 Image: Calc. Tax Image: Calc. Tax A1-700/0 Calculator U	Project Invoicing Retainage	
🗃 Browse Video Tutorials			90
Aged Receivables Location 1	La.) 70 21 🖓 105 🖓	0 II" 70	
Aged As OF 04/30/2020 (By Due Date) All Locations	(Ea.) 378 29 🖓 500,535 🖓	0 🖤 378	
30,000	Components Ship All	Order Subtotal 225.54	? . ×
20,000 - Post	Delete History Prepayment	Close	
10,000 - Current 1.30 31-60 61-90 Over 90 Aging Periods (Days) Total: <u>28,826</u>	Income Statement For 4 Periods Ending 04/30/2020 Revenue	Payables Outstanding: 3 days Annual Purchases: 24,815 Ending A/P Balance: 183	×
Done		Second Intranet	🖓 • 🔍 100% • 🔡

trivadis makes IT easier.

Order Entry – Use Cases

- Update the quantity in stock of all ordered products
 - When order status changes from "incomplete" to "complete"
 - When order status changes from "complete" to "cancelled"



- Ensure that the quantity in stock is always correct (no lost updates or similar)
- Create a report to show all sales for a year per country
- Create a report for the 5 top-selling products for a year









trivadis

makes IT easier.

Change Quantity in Stock – The SQL Way

Single Transaction

```
UPDATE ORDERS
SET order_status = :p_value_for_complete
WHERE order_id = :p_order_id;
MERGE INTO PRODUCTS t
USING (SELECT product_id,
SUM(quantity) AS quantity
FROM order_items
WHERE order_id = :p_order_id
GROUP BY product_id) s
ON (t.product_id = s.product_id)
WHEN MATCHED THEN
UPDATE SET t.quantity_on_stock =
t.quantity_on_stock = s.quantity;
```

COMMIT;

2014 © Trivadis SQL versus NoSQL 3rd July 2014



Sales Volume per Country – The SQL Way

```
SELECT c.name as country_name,
    SUM(i.quantity * i.unit_price) AS sales_volume
FROM order_items i
INNER JOIN orders o
    ON o.order_id = i.order_id
INNER JOIN addresses a
    ON a.address_id = o.shipping_address_id
INNER JOIN countries c
    ON c.country_id = a.country_id
WHERE o.order_date >= DATE '2013-01-01'
    AND o.order_date < DATE '2014-01-01'
GROUP BY c.name
ORDER BY 2 DESC;
```

trivadis makes IT easier.

SQL versus NoSQL 3rd July 2014

2014 © Trivadis

■ 5 Top-Selling Products – The SQL Way

```
SELECT p.name AS product_name,
    SUM(i.quantity * i.unit_price) AS sales_volume
FROM order_items i
INNER JOIN orders o
    ON o.order_id = i.order_id
INNER JOIN products p
    ON p.product_id = i.product_id
WHERE o.order_date > DATE '2013-01-01'
    AND o.order_date <= DATE '2014-01-01'
    AND o.order_status = :p_value_for_complete
GROUP BY p. name
ORDER BY 2 DESC
FETCH FIRST 5 ROWS WITH TIES;
```

trivadis makes IT easier.

2014 © Trivadis SQL versus NoSQL

3rd July 2014

MongoDB NoSQL Store



- Document Store
- Developed by 10gen, now MongoDB Inc.
- Professional grade support by MongoDB Inc.
- Main Features
 - JSON Data Model with Dynamic Schemas
 - Auto-Sharding for Horizontal Scalability
 - Built-In Replication for High Availability
 - Rich Secondary Indexes, including geospatial and TTL indexes
 - Text Search
 - Aggregation Framework & Native MapReduce





MongoDB Document Data Model (Aggregate Pattern)

Customer	Order	Product
		Category
Address	Billing Address	
		Suppliers
	Shipping Address	Supplier 1
		Address
	Order Items	Supplier 1
	Order Item 1	Address
	Order Item 2	
		Supplier n
	Order Item n	Address



Update Quantity in Stock – The MongoDB Way

Transaction 1

Read Operation

```
db.orders.find ( { orderId: 1} );
```

Transaction 2 .. n

```
ForEach orderItems.item {
    db.products.update( { productId : 101 },
        { $inc : { quantity: -10 } },
        { multi: false }
    );
}
```



Sales Volume per Country – The MongoDB Way

```
var mapFunction = function() {
    for (var idx = 0; idx < this.orderItems.length; idx++) {</pre>
        var value = this.orderItems[idx].unitPrice *
                          this.orderItems[idx].quantity;
        emit(this.shippingAddress.country, value);
    } };
var reduceFunction = function(name, valuesPrices) {
    return Array.sum(valuesPrices);
};
db.orders.mapReduce(mapFunction,
    reduceFunction,
    { out : {inline:1},
      query: { orderStatus: "COMPLETE",
                orderDate: { $qt: ISODate("2014-01-01"),
                             $lt: ISODate("2014-04-01") }
    });
```



5 Top-Selling Products – The MongoDB Way

```
db.orders.aggregate([
    { $match : {
           orderStatus: "COMPLETE",
           orderDate: { $qt: ISODate("2014-01-01"),
                        $lt: ISODate("2014-04-01") }
                       } },
    { $unwind : "$orderItems" },
    { $project : { id: 0,
           productId: "$orderItems.productId",
           total : { $multiply : ["$orderItems.quantity",
                                 "$orderItems.unitPrice"] }
    } },
    { $group : { _id: "$productId",
                total : { $sum : "$total"} } },
    { $sort : { total: -1 }},
    { $limit : 5 }
])
```

trivadis makes IT easier.







Round 3 Spotify



2014 © Trivadis SQL versus NoSQL 3rd July 2014



Spotify – Introduction

- Music service which is available worldwide
 - Over 20'000'000 music tracks available
 - Millions of users
 - Each user has dozens of playlists
- Many AP but also some CA use cases

Spotify – Example

000	Spotify Premium	2 M
♦ ► Q Search		Guido Schmutz 📡 .
🎜 ZeZé Di Camargo – Zezé di Camarg 🔺		Activity
🎵 Life On A Rock	Coldplay – Live 2012	Greg Stipkovich listened to
🎵 Tim McGraw - Two Lanes Of Freedom	Follower	S Runner's High
Bon Jovi – What About Now	Current Coldplay	The Hold Steady
Britt Nicole - Acoustic		
J Britt Nicole – Gold		Millow insteried to
JJ Zezé Di Camargo & Luciano – 2 💔 🌑	Available Offline 🛛 🔿 15 tracks (1h 6min)	Someday
J Emell Sande - Our Version of Events		Float Fail
Francesca Battistelli – My Paner Heart	Track Artist Time Album	Milow starred
Francesca Battistelli – Hundred Mor	Mylo Xyloto - Live Coldplay 0:58 Live 2012	Little Numbers
🎜 Travis – The Boy With No Name	🗠 Hurts Like Heaven – Live Executor Coldplay 4:16 Live 2012	BOY
🎵 Travis – 12 Memories	🖙 In My Place - Live Coldplay 3:55 Live 2012	Renata Lucia Schmutz listened to
🎵 Travis - The Man Who	🗠 Major Minus – Live EXELICIT Coldplay 3:40 Live 2012	Files An Mars
Reuno E Marrone - Sonhos Amores	Yellow - Live Coldplay 6:52 Live 2012	Jorge & Mateus
THE STATES THE AL	God Put a Smile Upon Your Executor Coldplay 5:22 Live 2012	
YOU RIN ADMS	Princess of China – Live Coldplay 3:49 Live 2012	Greg Stipkovich updated Best of 2014
M	Et Up in Elamos – Livo Coldplay 2:18 Livo 2012	
	Coldplay 5.16 Live 2012	Oaks
	Colopiay 4:58 Live 2012	INFAMS 9:02
	Charlie Brown - Live Coldplay 5:01 Live 2012	
	Paradise - Live Coldplay 5:32 Live 2012	Capitol Records listened to
	Us Against the World - Live Coldplay 3:52 Live 2012	Nude As The News
	Clocks - Live Coldplay 4:45 Live 2012	Cat Power
	Fix You - Live Coldplay 5:01 Live 2012	Milow created the playlist
Sector Chi	🗠 Every Teardrop Is a Waterfall - Live Coldplay 5:24 Live 2012	att Sa.
		Berlin
Coldolay		• • • • • • • • • • • • • • • • • • •
	C•	3:18 × C

0



Spotify – Use Cases

- Playlist, Showing Ads, Following Artists ... are all uses cases which have to be highly available, and accessible worldwide
 - Needs to be distributed to be fast
 - Service should be available even if a partition happen (due to network) failure/machine failure)
- First time subscription and subscription renewal must be absolutely consistent
 - Customer should only pay once!



makes IT easier.





Polyglot Persistence – SQL And NoSQL







Spotify 2 – 2

Draw!









- 1. Motivation
- 2. Overview of SQL and NoSQL Data Stores
- 3. Use Cases Let's Get Ready to Rumble
- 4. Recommended Reads
- 5. Core Messages





- <u>http://martinfowler.com/books/nosql.html</u>
- http://www.manning.com/mccreary/
- http://highlyscalable.wordpress.com
- <u>http://nosql-database.org</u>
- <u>http://db-engines.com/</u>





NoSQL Data Modeling Techniques

5.0 (g) details on an other compared by velocity and incident oriential and an analysis, performance, and consistency. This agrees of 2005g is well readed both in performance for the performance of the second second

a vession into the mains themes the production of the second state with a more set cause of the grammar. To explore data modeling techniques, we have to start with a more set less systematic view of NoSQL data models that preferably even it reads and and interconsections. The following figure depicts imaginary "evolution" of the major NoSQL system families, namely, Key-Valoe strees, high able with elastones, Decomment databases, Juli Testerori, Engine, and Orgals hardsectore







- 1. Motivation
- 2. Overview of SQL and NoSQL Data Stores
- 3. Use Cases Let's Get Ready to Rumble
- 4. Recommended Reads
- 5. Core Messages



Core Messages

- We will see a major consolation in the NoSQL area
- SQL is and stays important
- Polyglot persistence will be part of every solution design in the near future
- Enterprise capabilities are required
 - Tooling (monitoring, backup & recovery, data security, ...)
 - Organization, skills
 - Opportunity for cloud based solutions



2014 © Trivadis

SQL versus NoSQL 3rd July 2014

Questions and answers ...

Guido Schmutz Technology Manager

Philipp Salvisberg Senior Principal Consultant



BASEL BERN BRUGES LAUSANNE ZÜRICH DÜSSELDORF FRANKFURT A.M. FREIBURG I.BR. HAMBURG MUNICH STUTTGART VIENNA

L8 2014 © SQL ve 3rd July

